Clustering the Vertices of Sparse Edge-Weighted Graphs via Non-Backtracking Spectra
Marianna Bolla
BME Math. Inst.
marib@math.bme.hu
Summit280 Conference

July 8, 2024.

Outline

- Non-backtracking matrix of simple graphs.
- Sparse stochastic block model.
- Belief propagation.
- Inflation–deflation.
- K-means clustering with vertex representatives.
- Non-backtracking matrix of edge-weighted graphs.
- Edge percolation for simulated data.
- Application to real-world data.

<div align="center">

DENSE   SPARSE

Füredi-Komlós, Combinatorica   Percolated $G_n(p)$, $p = \frac{c}{n}$
(1981): $\mathbb{E}(a_{ij}) = \mu > 0 \quad (i \neq j)$ $(n \to \infty)$ Erdős-Rényi

B, Wiley (2013): $\mathbb{E}(a_{ij}) = c_{ab}$ Percolated SBM: $\mathbb{E}(a_{ij}) = \frac{c_{ab}}{n}$
$(i \in V_a, j \in V_b)$, $1 \leq a, b \leq k$ $(i \in V_a, j \in V_b)$, $1 \leq a, b \leq k$

$\lambda_i(\mathbf{A})$ is aligned with $\lambda_i(\mathbb{E}\mathbf{A})$, $\lambda_i(\mathbf{B})$ is aligned with $\lambda_i(\mathbb{E}\mathbf{A})$,
$i = 1, \ldots, k$ (F-K: $k = 1$) $i = 1, \ldots, k$ non-backtracking

Perturbation: Wigner-type matrix   deformed Wigner matrix
Subspace perturbation: Davis-Kahan   Bauer–Fike

</div>

Non-Backtracking (Hashimoto) matrix of simple graphs

$G = (V, E)$ simple graph, $|V| = n$, $|E| = m$;
The entries of the non-backtracking matrix $\mathbf{N} = (n_{ef})$ are indexed by the directed edges (bidirected edges of E), $|E^{\rightarrow}| = 2m$:

$$n_{ef} = \delta_{e \rightarrow f} \delta_{f \neq e^{-1}}, \quad n_{i \rightarrow j, \, s \rightarrow l} = \delta_{js}(1 - \delta_{il}),$$

where $e = \{i \rightarrow j\}$ and $f = \{s \rightarrow l\}$ are directed edges, and $e \rightarrow f$ with $e = (e_1, e_2)$ and $f = (f_1, f_2)$ means that $e_2 = f_1$; $e^{-1} = \{j \rightarrow i\}$.

Alon, Benjamini, Lubetzky, Sodin, Non-backtracking random walks mix faster, Commun. Contemp. Math. (2007).

Relation to line-graphs

If $\mathbf{N} = \begin{pmatrix} \mathbf{N}_{11} & \mathbf{N}_{12} \\ \mathbf{N}_{21} & \mathbf{N}_{22} \end{pmatrix}$, where the two (row/column) blocks
correspond to the edges and their inverses (in the same order), then

$$\mathbf{N}_{11}^* = \mathbf{N}_{22}, \quad \mathbf{N}_{22}^* = \mathbf{N}_{11}, \quad \mathbf{N}_{12}^* = \mathbf{N}_{12}, \quad \mathbf{N}_{21}^* = \mathbf{N}_{21}.$$

Further, $\mathbf{N}_{11} + \mathbf{N}_{12} + \mathbf{N}_{21} + \mathbf{N}_{22}$ is equal to the $m \times m$ adjacency
matrix of the line-graph of G.

In Lovász, Combinatorial Exercises it is proved that if the
line-graphs of two simple graphs, provided they both have
node-degrees at least 4, are isomorphic, then they are isomorphic
too. However, if the degree condition does not hold, it can happen
that two not isomorphic graphs have isomorphic line-graphs. For
example, a triangle and a star on 4 vertices.

Though **N** is not a normal matrix, even not always diagonalizable (the algebraic and geometric multiplicity of some of its eigenvalues may not be the same), it exhibits some symmetry: $n_{ef}^* = n_{e^{-1} f^{-1}}$.

With the notation $\breve{x}_e := x_{e^{-1}}$ for the coordinates of $\mathbf{x}, \breve{\mathbf{x}} \in \mathbb{R}^{2m}$: if $\mathbf{x} = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix}$, then $\breve{\mathbf{x}} = \begin{pmatrix} \mathbf{x}_2 \\ \mathbf{x}_1 \end{pmatrix}$ (swapping).

Consequently: if **x** is a right eigenvector of **N** corresponding to a real eigenvalue, then $\breve{\mathbf{x}}$ is a left eigenvector of **N** (and right eigenvector of $\mathbf{N}^*$) with the same real eigenvalue.

Eigenvalues of N (Ihara formula)

If G is connected, not a cycle and the minimum node-degree≥ 2, then $\mathbf{N}$ is irreducible. By Frobenius thm, it has a single real eigenvalue $\lambda_{max}(\mathbf{N}) > 0$ of largest modulus. Since the characteristic polynomial of $\mathbf{N}$ has real coefficients, its complex eigenvalues occur in conjugate pairs in the bulk, and there can be some other „structural" real eigenvalues out of the bulk.

Ihara formula: If G is not a tree, $\mathbf{N}$ has $m - n$ eigenvalues equal to 1 and $m - n$ eigenvalues equal to -1, whereas its further eigenvalues are those of the $2n \times 2n$ matrix

$$\mathbf{K} = \begin{pmatrix} \mathbf{O} & \mathbf{D_A} - \mathbf{I}_n \\ -\mathbf{I}_n & \mathbf{A} \end{pmatrix},$$

where $\mathbf{A}$ is the adjacency- and $\mathbf{D_A}$ is the degree-matrix of the graph (diagonal, contains the node-degrees=row-sums of $\mathbf{A}$).

$\mathbf{K}$ always has at least one eigenvalue 1, the geometric multiplicity of which is equal to the number of the connected components of G and $\lambda_{max}(\mathbf{K}) = \lambda_{max}(\mathbf{N}) \leq \lambda_{max}(\mathbf{A})$.

Real eigenvalues and eigenvectors of N

Two auxiliary matrices are introduced: the $2m \times n$ matrix **End** has entries $end_{ei} = 1$ if i is the end-node of the (directed) edge e and 0, otherwise; the $2m \times n$ matrix **Start** has entries $start_{ei} = 1$ if i is the start-node of the (directed) edge e and 0, otherwise. Then for any vector $\mathbf{u} \in \mathbb{R}^n$ and for any directed edge $\{i \to j\}$:

$$(\mathbf{End\,u})_{i \to j} := u_j \quad \text{and} \quad (\mathbf{Start\,u})_{i \to j} := u_i.$$

Consequently, **End u** is the $2m$-dimensional inflated version of the n-dimensional vector $\mathbf{u}$, where the coordinate u_j of $\mathbf{u}$ is repeated as many times, as many edges have end-node j; likewise, in the $2m$-dimensional inflated vector **Start u**, the coordinate u_i of $\mathbf{u}$ is repeated as many times, as many edges have start-node i.
As each edge is considered in both possible directions, these multiplicities are the node-degrees d_j and d_i, respectively.

End* **End** = **Start*** **Start** = $\operatorname{diag}(d_1, \ldots, d_n)$ = $\mathbf{D_A}$

For any vector $\mathbf{x} \in \mathbb{R}^{2m}$, define

$$x_i^{out} := \sum_{j:\, j \sim i} x_{i \to j} \quad \text{and} \quad x_i^{in} := \sum_{j:\, j \sim i} x_{j \to i} \quad (i = 1, \ldots, n).$$

These become the coordinates of the n-dimensional (column) vectors $\mathbf{x}^{in}$ and $\mathbf{x}^{out}$: $\mathbf{x}^{out} = \mathbf{Start}^*\mathbf{x}$ and $\mathbf{x}^{in} = \mathbf{End}^*\mathbf{x}$.

If $\mathbf{x}$ is a (right) eigenvector of $\mathbf{N}^*$ with structural real eigenvalue μ, then the vector $\begin{pmatrix} \mathbf{x}^{out} \\ \mathbf{x}^{in} \end{pmatrix}$ is a right eigenvector of $\mathbf{K}$ with the same eigenvalue μ, where $\mathbf{x}^{out} = \frac{1}{\mu}(\mathbf{D_A} - \mathbf{I}_n)\mathbf{x}^{in}$.

Edge percolation

The edge percolation threshold for the giant component to appear in a sparse simple graph is $\beta > \frac{1}{\lambda_{max}(\mathbf{N})}$, where β is the edge retention probability, see Newman, M. E. J., Message passing methods on complex networks, Proc. R. Soc. London A (2023). The proof uses the method of Belief Propagation (BP) (when the so-called message passing system of the approximating linear equations has a non-trivial and unstable trivial solution).

In the dense case, it happens at $\frac{1}{\lambda_{max}(\mathbf{A})} \le \frac{1}{\lambda_{max}(\mathbf{N})}$, see Bollobás, B., Borgs, C., Chayes, J., and Riordan, O., Percolation on dense graph sequences, Ann. Probab. (2010).

More generally, we are looking for the number k, so that k strongly connected clusters (communities) can be detected (within the giant component) in a graph coming from the sparse stochastic block model. We are also looking for the clusters themselves.

The Erdős–Rényi graph $G_n(p)$ is a special case with $k = 1$, where the edges of the complete graph on n vertices are retained independently, with probability $\beta = p$.

The sparse stochastic block model SBM_k

The $k \times k$ *probability matrix* **P** of the random graph $G_n \in SBM_k$ has entries $p_{ab} = \frac{c_{ab}}{n}$, where the $k \times k$ symmetric affinity matrix **C** $= (c_{ab})$ stays constant as $n \to \infty$. An edge between $i < j$ comes into existence, independently of the others, with probability p_{ab} if $i \in V_a$ and $j \in V_b$, where $(V_1, \ldots, V_k)$ is a partition of the node-set V into k disjoint clusters; $a_{ji} := a_{ij}$. It can be extended to the $i = j$ case when self-loops are allowed, or else, the diagonal entries of the adjacency matrix are zeros.

Ā: the $n \times n$ inflated matrix of the $k \times k$ **P**: $\bar{a}_{ij} = p_{ab}$ if $i \in V_a$ and $b \in V_b$. When loops are allowed, then $\mathbb{E}(a_{ij}) = \bar{a}_{ij}$ for all $1 \leq i, j \leq n$. In the loopless case, the expected adjacency matrix $\mathbb{E}$**A** differs from **Ā** with respect to the the main diagonal, but the diagonal entries are negligible.

Special cases

Sometimes $c_{ab} = c_{in}$ is the within-cluster ($a = b$) and $c_{ab} = c_{out}$ is the between-cluster ($a \neq b$) affinity. The network is called assortative if $c_{in} > c_{out}$, and disassortative if $c_{in} < c_{out}$. Of course, remarkable difference is needed between the two, to recognize the clusters.

The cluster sizes are $n_1, \ldots, n_k$ ($\sum_{i=1}^{k} n_i = n$), so the $k \times k$ diagonal matrix $\mathbf{R} := \mathrm{diag}(r_1, \ldots, r_k)$, where $r_a = \frac{n_a}{n}$ is the relative size of cluster a ($a = 1, \ldots, k$), is also a model parameter ($\sum_{a=1}^{k} r_a = 1$). It is nearly kept fixed as $n \to \infty$.

The model SBM_k is called symmetric if $r_1 = \cdots = r_k = \frac{1}{k}$ and all diagonal entries of the affinity matrix are equal to c_{in}, whereas the off-diagonal ones to c_{out}.

The average degree of a real world graph on m edges and n nodes is $\frac{2m}{n}$. The expected average degree of the random graph $G_n \in SBM_k$ is

$$c = \frac{1}{n} \sum_{a=1}^{k} \sum_{b=1}^{k} n_a n_b p_{ab} = \frac{1}{n^2} \sum_{a=1}^{k} \sum_{b=1}^{k} n_a n_b c_{ab} = \sum_{a=1}^{k} r_a c_a,$$

where $c_a = \sum_{b=1}^{k} r_b c_{ab}$ is the average degree of cluster a. It is valid only if self-loops are allowed. Otherwise, c_a and c should be decreased with a term of order $\frac{1}{n}$, but it will not make too much difference in the subsequent calculations.

Kesten–Stigum threshold

In Bordenave, C., Lelarge, M., Massoulié, L., Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs, Ann. Probab. (2018), the case when $c_a = c$ for all a is considered. (This is the hardest case, as otherwise the clusters could be distinguished by sorting the node-degrees.) In this case $\frac{1}{c}\bar{\mathbf{A}}$ is a stochastic matrix, and so, the spectral radius of $\bar{\mathbf{A}}$ is c.

In the symmetric case, $c = \frac{c_{in}+(k-1)c_{out}}{k}$ and the separation of the clusters only depends on the c_{in}, c_{out} relation. If c_{in} is „close" to c_{out}, then the groups cannot be distinguished. The detectability Kesten–Stigum threshold in the symmetric case is

$$|c_{in} - c_{out}| > k\sqrt{c} \Longleftrightarrow \mu_2 = \cdots = \mu_k > \sqrt{c},$$

where $c = \mu_1 > \mu_2 = \cdots = \mu_k$ are the leading (real) eigenvalues of $\mathbf{N}$.

BP in the general sparse SBM_k model

Given the observed graph G on n nodes, assume that our neighbors are independent of each other, when conditioned on our own state. This can be modeled by having each node j send a „message" to i, which is an estimate of j's marginal if i were not there. Therefore, the conditional probability

$$\psi_{j \to i}^a := \mathbb{P}(j \text{ is in cluster } a \text{ when } i \text{ is not present})$$

can be computed through neighbors of j that are different from i:

$$\psi_{j \to i}^a = C_a^{ij} r_a \prod_{l \sim j, \, l \neq i} \sum_{b=1}^{k} \psi_{l \to j}^b p_{ab}, \quad a = 1, \ldots, k,$$

where C_a^{ij} is a normalizing factor.

This BP (belief propagation) system of equations ($2mk$ non-linear equations with the same number of unknowns) can be solved by initializing messages randomly, then repeatedly updating them.

Calculations

In Moore, C., The computer science and physics of community detection: Landscapes, phase transitions, and hardness, Bull. EATCS (2017), the symmetric case is treated, when BP has a trivial fixed point $\psi_{j\to i}^a = \frac{1}{k}$, for $a = 1, \ldots, k$. If it gets stuck there, then BP does no better than chance. It happens when this trivial fixed point of this discrete dynamical system is asymptotically stable.

In the generic case, we have an unstable fixed point via linearization:

$$\psi_{j\to i}^a := r_a + \varepsilon_{j\to i}^a.$$

We substitute it in the original BP system and expand it to first order in ε (vector of $2mk$ coordinates $\varepsilon_{j\to i}^a$'s).

The linear dynamical system approximating the above system of difference equations is $\boldsymbol{\varepsilon} = (\mathbf{N} \otimes \mathbf{T})\boldsymbol{\varepsilon}$, where $\mathbf{T} = \mathbf{GRC}$ is the transmission matrix with $\mathbf{G} = \operatorname{diag}(\frac{1}{c_1}, \ldots, \frac{1}{c_k})$.

The fixed point $\mathbf{0}$ of $\boldsymbol{\varepsilon}^{(t+1)} = (\mathbf{N} \otimes \mathbf{T})\boldsymbol{\varepsilon}^{(t)}$ is unstable, if the spectral radius $\rho(\mathbf{N} \otimes \mathbf{T}) > 1$.

Remark:

- The eigenvalues of $\mathbf{N} \otimes \mathbf{T}$ are the products of the eigenvalues of $\mathbf{N}$ and $\mathbf{T}$.
- If $r_1 = \cdots = r_k$, then $\mathbf{T}$ is a stochastic matrix, so $\rho(\mathbf{T}) = 1$.
- If G is a connected graph which is not a tree or cycle and $d_{\min}(G) \geq 2$, then $\rho(\mathbf{N}) > 1$.
- $\rho(\mathbf{N} \otimes \mathbf{T}) > 1$ is a necessary condition for the existence of a non-trivial fixed point. If it exists, then a sparse SBM_k model can be fitted. How the eigenvalues of $\mathbf{N}$ and $\mathbf{T}$ are related?

Special cases

If $c_1 = \cdots = c_k = c$, then for

$$\lambda(\mathbf{N} \otimes (c\mathbf{T})) = \lambda(\mathbf{N})\,\lambda(\mathbf{RC}) > c$$

it suffices that $\lambda(\mathbf{N}) > \sqrt{c}$, as the eigenvalues of $\mathbf{N}$ and $\mathbf{RC}$ are aligned, see

Bordenave, C., Lelarge, M., Massoulié, L., Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs, Ann. Prob. (2018).

They prove that w.h.p.

$$\mu_i = \nu_i + o(1) \quad (i = 1, \ldots k_0) \quad \text{and} \quad \mu_i < \sqrt{c} + o(1) \quad (i > k_0),$$

where μ_i's and ν_i's $(i = 1, \ldots, k_0)$ are the structural eigenvalues of $\mathbf{N}$ and $\mathbf{RC}$, respectively, whereas $k_0 \leq k$ is the positive integer for which $\nu_i^2 \geq \nu_1$ $(i = 1, \ldots k_0)$ and $\nu_{k_0+1}^2 < \nu_1$ holds.

In particular, in the SBM_1 (Erdős–Rényi) model, $\mu_1 = c + o(1)$ and $\mu_2 \leq \sqrt{c} + o(1)$.

Even if the average degrees of the clusters are not the same, in the next (Inflation–Deflation) slide we will show that the non-zero eigenvalues of $\bar{\mathbf{A}}$ are the same as those of $\mathbf{RC}$, so they are in the neighborhood of the leading eigenvalues of $\mathbf{N}$ within a factor between u and v, where

$$u = \min_a \frac{c}{c_a} \quad \text{and} \quad v = \max_a \frac{c}{c_a}.$$

However, the leading eigenvalues of $\bar{\mathbf{A}}$ and $\mathbf{A}$ are farther apart, seemingly contradicting to the laws of large numbers.
Also see the theory of deformed Wigner matrices: Capitaine, M., Donati-Martin, C., Féral, D., The largest eigenvalues of finite rank deformation of large Wigner matrices,..., Ann. Prob. (2009).

Inflation–deflation

Proposition

The matrix $\bar{\mathbf{A}}$ has rank k and its non-zero eigenvalues (ν's) with unit norm eigenvectors ($\mathbf{u}$'s) satisfy $\bar{\mathbf{A}}\mathbf{u} = \nu\mathbf{u}$, where $\mathbf{u}$ is the inflated vector of $\tilde{\mathbf{u}} = (u(1), \ldots, u(k))^*$ with block-sizes $n_1, \ldots, n_k$. With the notation $\mathbf{R} = \frac{1}{n}\mathrm{diag}(n_1, \ldots, n_k) = \mathrm{diag}(r_1, \ldots, r_k)$, the deflated equation is equivalent to

$$\mathbf{R}^{\frac{1}{2}}\mathbf{C}\mathbf{R}^{\frac{1}{2}}\mathbf{v} = \nu\mathbf{v},$$

where $\mathbf{v} = \sqrt{n}\mathbf{R}^{\frac{1}{2}}\tilde{\mathbf{u}}$. Further, if $\mathbf{u}_1, \ldots \mathbf{u}_k$ is the set of orthonormal eigenvectors of $\bar{\mathbf{A}}$, then $\mathbf{v}_i = \sqrt{n}\mathbf{R}^{\frac{1}{2}}\tilde{\mathbf{u}}_i$ $(i = 1, \ldots, k)$ is the set of orthonormal eigenvectors of $\mathbf{R}^{\frac{1}{2}}\mathbf{C}\mathbf{R}^{\frac{1}{2}}$. Also, $\mathbf{R}^{\frac{1}{2}}\mathbf{v}_i = \sqrt{n}\mathbf{R}\tilde{\mathbf{u}}_i$ are right eigenvectors of $\mathbf{R}\mathbf{C}$ and $\mathbf{R}^{-\frac{1}{2}}\mathbf{v}_i = \sqrt{n}\tilde{\mathbf{u}}_i$ are left eigenvectors of $\mathbf{R}\mathbf{C}$ with the same eigenvalues ν_i, for $i = 1, \ldots, k$.

Proposition (Based on Theorem 1 of Stephan, L., Massoulié, Non-backtracking spectra of inhomogeneous random graphs, Mathematical Statistics and Learning (2022).)

Let $\mathbb{E}\mathbf{A}$ be the expected adjacency matrix of a random simple graph. Assume that $k = \mathrm{rank}\,(\mathbb{E}\mathbf{A}) = n^{o(1)}$, the graph is sparse enough, and the eigenvectors corresponding to the non-zero eigenvalues of the matrix $\mathbb{E}\mathbf{A}$ are sufficiently delocalized. Let k_0 denote the number of eigenvalues of $\mathbb{E}\mathbf{A}$ whose absolute value is larger than $\sqrt{\rho}$, where ρ is the spectral radius of $\mathbb{E}\mathbf{A}$: these are $\nu_1 \geq \cdots \geq \nu_{k_0}$ with corresponding eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_{k_0}$ (they form an orthonormal system as $\mathbb{E}\mathbf{A}$ is a real symmetric matrix). Then for $i \leq k_0 \leq k$, the ith largest eigenvalue μ_i of $\mathbf{N}$ is asymptotically (as $n \to \infty$) equals to ν_i and all the other eigenvalues of $\mathbf{N}$ are constrained to the circle (in the complex plane) of center 0 and radius $\sqrt{\rho}$.

Proposition continued (eigenvectors of N)

Further, if $i \leq k_0$ is such that ν_i is a sufficiently isolated eigenvalue of $\mathbb{E}\mathbf{A}$, then the standardized eigenvector of $\mathbf{N}$ corresponding to μ_i has inner product close to 1 with the standardized inflated version of $\mathbf{u}_i$, namely, with $\frac{\mathbf{End}\,\mathbf{u}_i}{\|\mathbf{End}\,\mathbf{u}_i\|}$.

Let $\mathbf{x}$ be a unit-norm eigenvector of $\mathbf{N}$, corresponding to the eigenvalue μ that is close to the eigenvalue ν of the expected adjacency matrix, with corresponding eigenvector $\mathbf{u} \in \mathbb{R}^n$. If our graph is from the SBM_k model, then (without knowing its parameters) we know that $\mathbf{u}$ is a step-vector with at most k different coordinates. Then by the above Proposition,

$$\left\langle \mathbf{x}, \frac{\mathbf{End}\,\mathbf{u}}{\|\mathbf{End}\,\mathbf{u}\|} \right\rangle \geq \sqrt{1 - \varepsilon} \geq 1 - \frac{1}{2}\varepsilon,$$

where ε can be arbitrarily „small" with increasing n.

Continued

$$\left\| \mathbf{x} - \frac{\mathbf{End}\,\mathbf{u}}{\|\mathbf{End}\,\mathbf{u}\|} \right\|^2 \leq 2 - 2(1 - \frac{1}{2}\varepsilon) = \varepsilon$$

and by $\mathbf{x}^{in} = \mathbf{End}^*\mathbf{x}$ and $\mathbf{End}^*\mathbf{End} = \mathbf{D_A}$,

$$\left\| \mathbf{End}^*\mathbf{x} - \mathbf{End}^* \frac{\mathbf{End}\,\mathbf{u}}{\|\mathbf{End}\,\mathbf{u}\|} \right\|^2 = \left\| \mathbf{x}^{in} - \mathbf{D_A} \frac{\mathbf{u}}{\|\mathbf{End}\,\mathbf{u}\|} \right\|^2.$$

Consequently,

$$\left\| \mathbf{D_A}^{-1}\mathbf{x}^{in} - \frac{\mathbf{u}}{\|\mathbf{End}\,\mathbf{u}\|} \right\|^2 \leq \|\mathbf{D_A}^{-1}\mathbf{End}^*\|^2 \varepsilon \leq \varepsilon$$

as $\|\mathbf{D_A}^{-1}\mathbf{End}^*\|^2 \leq \max_i \frac{1}{d_i} = \frac{1}{\min_i d_i} \leq 1$.

Summarizing

Assume that the expected adjacency matrix of the underlying random graph on n nodes and m edges has rank k with k single eigenvalues and corresponding unit-norm eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_k \in \mathbb{R}^n$. Assume that the non-backtracking matrix $\mathbf{N}$ of the random graph has k structural eigenvalues (aligned with those of the expected adjacency matrix) with eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^{2m}$ such that

$$\left\langle \mathbf{x}_j, \frac{\mathbf{End}\,\mathbf{u}_j}{\|\mathbf{End}\,\mathbf{u}_j\|} \right\rangle \geq \sqrt{1-\varepsilon}, \quad j = 1, \ldots, k.$$

Then for the transformed vectors $\mathbf{D}_\mathbf{A}^{-1}\mathbf{x}_j^{in} \in \mathbb{R}^n$, the relation $\sum_{j=1}^{k} \left\| \mathbf{D}_\mathbf{A}^{-1}\mathbf{x}_j^{in} - \frac{\mathbf{u}_j}{\|\mathbf{End}\,\mathbf{u}_j\|} \right\|^2 \leq k\varepsilon$ holds.

Consequences

Corollary: If $\mathbf{u}_j$'s are step-vectors on k steps (e.g., if our graph comes from the SBM_k model), then the k-variance (sum of the inner variances of the clusters) of the node representatives (objective function of the k-means algorithm)

$$(\frac{1}{d_i}x_{1i}^{in}, \dots, \frac{1}{d_i}x_{ki}^{in}), \quad i = 1, \dots, n$$

is estimated from above with $k\varepsilon$ too.

Remark: In case of a simple graph, the n-dimensional vectors $\mathbf{x}_j^{in}$ ($j = 1, \dots, k$) are the first segments of the right eigenvectors of the matrix $\mathbf{K}$. So, we have to perform the spectral decomposition of a $2n \times 2n$ matrix only instead of a $2m \times 2m$ one, which fact has further computational benefit (except for trees, $n \le m$, but usually n is much smaller than m).

Possible workflow for the unweighted case

- Find k (initial number of clusters) based on the spectral gap in **N** (suggested by BP).
- Run the EM (Expectation–Maximization) algorithm to estimate the parameters of the k-cluster model. Investigate hypotheses on the model fit (likelihood ratio test, information theoretical criteria).
- In case of „good" fit, find the clusters with the help of the vertex representatives (deflated **N**-eigenvectors), segments of **K**.

Non-Backtracking matrix of edge-weighted graphs

Let $G = (V, E)$ be the skeleton of an edge-weighted graph, $|V| = n$, $|E| = m$; the weight of edge $e = \{i, j\}$ is $W_e = w_{ij} = w_{ji} > 0$, where the remaining entries of the the $n \times n$ symmetric edge weight matrix $\mathbf{W}$ are zeros (including the diagonal).

Let the $2m \times 2m$ diagonal matrix $\mathbf{D}$ contain the positive edge-weights in its main diagonal (the first m diagonal entries are the same as the second m ones as $W_e = W_{e^{-1}}$). With them,

$$\mathbf{B} = \mathbf{ND} \quad \text{and} \quad \mathbf{B}^* = \mathbf{DN}^*.$$

The general entry of the $2m \times 2m$ non-backtracking matrix $\mathbf{B}$ is

$$b_{ef} = W_f \delta_{e \to f} \delta_{f \neq e^{-1}}.$$

Notation

We assume that there are constants C_1 and C_2 (independent of n):

$$C_1 \leq w_{ij} \leq C_2, \quad \text{for} \quad w_{ij} \neq 0.$$

Further, we assume that the skeleton's node degrees

$$d_i = |\{i : w_{ij} > 0, \quad j = 1, \ldots, n\}|, \quad i = 1, \ldots, n$$

are of constant order (it is the case in the k-cluster stochastic block model (SBM_k)).

Let $\mathbf{D^W}$ denote the $n \times n$ diagonal matrix of diagonal entries

$$d_i^{\mathbf{W}} = \sum_{j=1}^{n} w_{ij}, \quad i = 1, \ldots, n,$$

that are the so-called generalized degrees. In the unweighted case (0-1 weights), $d_i^{\mathbf{W}} = d_i$ and $C_1 = C_2 = 1$; in general,

$$C_1 d_i \leq d_i^{\mathbf{W}} \leq C_2 d_i, \quad i = 1, \ldots, n.$$

Start- and End-matrices, in- and out-vectors

The **End** and **Start** matrices are defined as in the unweighted case:

$$\textbf{End}^* \, \textbf{D} \, \textbf{End} = \textbf{Start}^* \, \textbf{D} \, \textbf{Start} = \textbf{D}^{\textbf{W}} \quad \text{and} \quad \textbf{Start}^* \, \textbf{D} \, \textbf{End} = \textbf{W}.$$

For any vector $\mathbf{x} \in \mathbb{R}^{2m}$, the following n-dimensional vectors are introduced:

$$\mathbf{x}^{out} := \textbf{Start}^* \textbf{D} \mathbf{x} \quad \text{and} \quad \mathbf{x}^{in} := \textbf{End}^* \textbf{D} \mathbf{x}.$$

Coordinatewise, for $i = 1, \ldots, n$,

$$x_i^{out} = \sum_{j:\, j \sim i} w_{ij} x_{i \to j} = \sum_{e:\, e_1 = i} W_e x_e, \quad x_i^{in} = \sum_{j:\, j \sim i} w_{ij} x_{j \to i} = \sum_{e:\, e_2 = i} W_e x_e.$$

Tracing back the problem to lower order matrices

No counterpart of matrix **K** works here, but if we know a real eigenvalue μ of **B**, we are able to find a linear system of equations for the *out*-transform of the corresponding eigenvector that is necessary for spectral clustering. With a Laplacian type equation, μ can also be concluded.

Proposition: Let **x** be a (right) eigenvector of **B** corresponding to a single positive real eigenvalue μ such that $\mu \neq w_{ij}$, $\forall i, j \in \{1, \ldots, n\}$. Then $\mathbf{y} = \mathbf{x}^{out}$ satisfies the homogeneous system of linear equations

$$[\mathbf{I}_n - \tilde{\mathbf{A}}(\mu) + \tilde{\mathbf{D}}(\mu)]\mathbf{y} = \mathbf{0}$$

with a Laplacian type coefficient matrix, where

$$\tilde{\mathbf{A}}(\mu)_{ij} = \frac{\mu w_{ij}}{\mu^2 - w_{ij}^2} \quad \text{and} \quad \tilde{\mathbf{D}}(\mu)_{ii} = \sum_{j=1}^{n} \frac{w_{ij}^2}{\mu^2 - w_{ij}^2},$$

with the understanding that $w_{ij} = 0$ whenever $i \not\sim j$.

Consequences

The above homogeneous system of linear equations for the coordinates of **y** must have a non-trivial solution, so

$$|\mathbf{I}_n - \tilde{\mathbf{A}}(\mu) + \tilde{\mathbf{D}}(\mu)| = 0.$$

This is not a polynomial (characteristic) equation, but it is a rational function of μ. By the assumptions of the Proposition, the denominators are not zeros, so we can multiply the determinant equations with them, and we obtain an high-degree (higher than n) polynomial of μ.

The leading positive real solutions $\mu_1 \geq \cdots \geq \mu_k$ are the same as the structural eigenvalues of **B**. Their number will be denoted by k. The corresponding $\mathbf{y}_1, \ldots, \mathbf{y}_k$ can be obtained by solving the system of the above homogeneous linear equations (with only an $n \times n$ coefficient matrix).

Clustering

Proposition of Stephan, L., Massoulié, Non-backtracking spectra of inhomogeneous random graphs, Mathematical Statistics and Learning (2022) is applicable to the edge-weighted case too.

At the instance, when the number of vertices is n: $\mathbf{P} = (p_{ij})$ is the $n \times n$ symmetric probability matrix of the edges and $\mathbb{W} = (W_{ij})$ is the $n \times n$ symmetric matrix of random weights of the edges. The proposition is applicable if $\bar{\mathbf{A}} := \mathbf{P} \circ \mathbb{EW}$ is a low rank matrix and the so obtained graph is sparse enough. A constant average degree can be guaranteed if, in the instance of n nodes, the p_{ij}'s are proportional to $\frac{1}{n}$. The authors only require for the average degree to be of order $o(\log n)$. In the classical literature, for the average degrees, the order $o(n)$ is considered as sparse. Nowadays the notion of intermediate density is introduced, e.g., for $\log n$ or $Poly(\log n)$ order average degrees, which is the case in the subsequent quantum chemistry examples.

The $SBM_k^{\mathbf{W}}$ model

The stochastic block models discussed before are special cases, where the weights are constantly 1.

In many practical situations, the entry W_{ij} of the $n \times n$ random weighted adjacency matrix $\mathbb{W}$ is $\tilde{w}_{ij}$ times a Bernoulli distributed random variable with parameter p_{ij}, for $1 \leq i < j \leq n$; these entries above the diagonal are independent of each other, while those below the diagonal are identical to them.

So the parameters of this distribution are contained in the symmetric matrices $\mathbf{P}$ and $\widetilde{\mathbf{W}} = (\tilde{w}_{ij})$ of real entries in (0,1]. In this way, the expected adjacency matrix is $\bar{\mathbf{A}} = \mathbf{P} \circ \widetilde{\mathbf{W}}$, with approximate matrix of variances $\mathbf{P} \circ \widetilde{\mathbf{W}} \circ \widetilde{\mathbf{W}}$ (if the entries of $\mathbf{P}$ are decreasing with n), and so, the above theory is applicable to it. However, for given (large) n, we only observe a realization $\mathbf{W}$ from the the distribution $\mathbb{W}$, in the non-zero positions of which the entries are equal to those of $\widetilde{\mathbf{W}}$.

Even if the observed $\mathbf{W}$ is sparse, the expected weighted adjacency matrix $\bar{\mathbf{A}}$ is not sparse (it is full) $\Longrightarrow$ dense matrix techniques.

Clustering in the $SBM_k^{\mathbf{W}}$ model

Let $\mathbf{x}$ be a unit-norm eigenvector of $\mathbf{B}$, corresponding to the eigenvalue μ that is close to the eigenvalue ν of the low-rank matrix $\bar{\mathbf{A}} = \mathbf{P} \circ \widetilde{\mathbf{W}}$, with corresponding eigenvector $\mathbf{u} \in \mathbb{R}^n$. Even if $\bar{\mathbf{A}}$ is not of low rank, but can be approximated with a low-rank matrix, the eigenvectors corresponding to its structural eigenvalues can also be approximated by step-vectors $\mathbf{v}$'s:

$$\left\langle \mathbf{x}, \frac{\mathbf{End}\,\mathbf{v}}{\|\mathbf{End}\,\mathbf{v}\|} \right\rangle \geq \sqrt{1 - \varepsilon} \geq 1 - \frac{1}{2}\varepsilon,$$

where ε can be arbitrarily „small" with increasing n.

$$\left\| \mathbf{x} - \frac{\mathbf{End\,v}}{\|\mathbf{End\,v}\|} \right\|^2 \le 2 - 2(1 - \frac{1}{2}\varepsilon) = \varepsilon$$

and

$$\left\| \mathbf{x}^{out} - \mathbf{W}\frac{\mathbf{v}}{\mathbf{Endv}} \right\|^2 = \left\| \mathbf{Start}^*\mathbf{D}\left(\mathbf{x} - \frac{\mathbf{Endv}}{\|\mathbf{Endv}\|} \right) \right\|^2 \le \|\mathbf{Start}^*\mathbf{D}\|^2\varepsilon.$$

Consequently,

$$\left\| \mathbf{W}^{-1}\mathbf{x}^{out} - \frac{\mathbf{v}}{\|\mathbf{End\,v}\|} \right\|^2 \le \|\mathbf{W}^{-1}\mathbf{Start}^*\mathbf{D}\|^2\varepsilon \le \left(\frac{C_2}{C_1} \right)^2 \varepsilon.$$

Consequences

Assume that the non-backtracking matrix $\mathbf{B}$ has k structural (real) eigenvalues with eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^{2m}$. Then for the transformed vectors $\mathbf{W}^{-1}\mathbf{x}_j^{out} \in \mathbb{R}^n$:

$$\sum_{j=1}^{k} \left\| \mathbf{W}^{-1}\mathbf{x}_j^{out} - \frac{\mathbf{v}_j}{\|\mathbf{End}\,\mathbf{v}_j\|} \right\|^2 \leq k\varepsilon \frac{C_2^2}{C_1^2}.$$

If $\mathbf{v}_j$'s are step-vectors on k steps, then the left-hand side estimates from above the k-variance of the node representatives that are row vectors of

$$(\mathbf{W}^{-1}\mathbf{x}_1^{out}, \ldots, \mathbf{W}^{-1}\mathbf{x}_k^{out}).$$

To get the $\mathbf{x}_j^{out} \in \mathbb{R}^n$ vectors we do not need the $2m$-dimensional eigenvectors $\mathbf{x}_j$'s of $\mathbf{N}$, but the previous calculations can be used.

The SBM_k^β model: edges are retained with probability β

The $k \times k$ probability matrix is $\frac{\beta \mathbf{C}}{n}$: $\mathbf{C}$ and c is multiplied by β, but $\mathbf{T} = \mathbf{GRC}$ remains unchanged.

So we consider $\beta c \mathbf{T}$ as for the model side, but the underlying graph and its $\mathbf{N}$ is the same as before. Therefore, the eigenvalues of $\mathbf{N} \otimes \beta c \mathbf{T} = \beta c (\mathbf{N} \otimes \mathbf{T})$ should be greater than c if a non-stable solution is required:

$$\beta \, \lambda(\mathbf{N} \otimes c\mathbf{T}) > c.$$

If the eigenvalues of $\mathbf{N}$ and $c\mathbf{T}$ are aligned, then this gives that $\lambda(\mathbf{N}) > \frac{\sqrt{c}}{\sqrt{\beta}}$ is needed for detectability; equivalently,

$$\beta > \frac{c}{\lambda^2(\mathbf{N})} = \left(\frac{\sqrt{c}}{\lambda(\mathbf{N})} \right)^2.$$

Remarks

- This is in accord with the fact, that in the $k = 1$ case (Erdős–Rényi model), $\lambda(\mathbf{N}) = c$ and $\beta = \frac{1}{c}$; special case of the percolation threshold $\frac{1}{\mu_1}$, see Newman, M. E. J., Message passing methods on complex networks, Proc. R. Soc. London A (2023), when the complete graph is edge-percolated.
- In the multiclass scenario, $\beta_i := \frac{c}{\mu_i^2}$ are further phase transitions, leading to i clusters, for $i = 1, \ldots, k_0$ until $\mu_{k_0} \geq \sqrt{c}$, but $\mu_{k_0+1} < \sqrt{c}$.
- This has relevance only if $\lambda_{max}(\mathbf{N}) > \sqrt{c}$, so eigenvalues of $\mathbf{N}$ greater than $\sqrt{c}$ give the phase transitions.
- Since $\mu_1 \geq \mu_2 \geq \ldots$, with larger β, larger number of clusters can be detected.
- $SBM_k^{\mathbf{W}}$ is a generalization of the SBM_k^{β} model, where the edges may have different edge-retention probability ($0 < w_{ij} \leq 1$ for the connected vertex-pair i, j); e.g., transmission of the infection depends on the randomly coupled individuals. Then the eigenvalues of $\mathbf{B} = \mathbf{ND}$ are used.

Simulation

A random graph was generated on $n = 900$ nodes, with parameter matrices $\mathbf{R} = \mathrm{diag}(\frac{35}{107}, \frac{42}{107}, \frac{30}{107})$ and

$$\mathbf{C} = \begin{pmatrix} 30 & 11.28 & 7.728 \\ 11.28 & 25 & 10.36 \\ 7.728 & 10.36 & 35 \end{pmatrix}$$

constructed so that the average degrees of the clusters be the same, i.e., $c_a = \sum_{b=1}^{3} r_b c_{ab}$ is the same for $a = 1, 2, 3$.

In the symmetric case, $\beta_2 = \cdots = \beta_k$, so all these phase transitions occur at the same time from the giant cluster to the k-cluster scenario.

Multiple transitions are spectacular if the eigenvalues of $\mathbf{B}$ greater than $\sqrt{c}$ are separated from each other and from $\sqrt{c}$.

(a) Adjacency matrix.

(b) Spectrum of matrix **K**.

(c) Adjacency matrix.

(d) Spectrum of matrix **K**.

(e) Adjacency matrix.


(f) Spectrum of matrix **K**.

(g) Adjacency matrix.

(h) Spectrum of matrix **K**.

Quantum chemistry application: 22 real eigenvalues of the the $2n \times 2n$ matrix **K** of the water molecule's skeleton, $n = 133$

Nodes: wave functions=electrons (Ferenc Krausz, Nobel-prize).

11 clusters obtained by the k-means algorithm applied to the in-vectors corresponding to the leading K-eigenvectors (skeleton graph, but darker colors correspond to larger edge-weights)
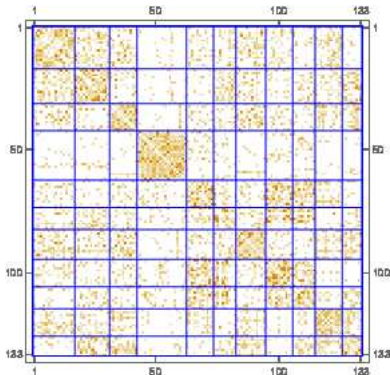
Leading 50 real eigenvalues of the the $2m \times 2m$ matrix B, $m = 3032$ (weighted case)

7 clusters obtained by the k-means algorithm applied to the out-vectors corresponding to the 7 leading B-eigenvectors

4 clusters obtained by the k-means algorithm applied to the out-vectors corresponding to the 4 leading B-eigenvectors

Acknowledgements