Fundamental Limits of Community Detection in Contextual Multi-Layer Stochastic Block Models

Zhangsong Li

School of Mathemaitcal Sciences, Peking University

joint with

Shuyang Gong
(PKU → Yale)

Dong Huang
(THU)

Recovering communities from networks

- Networks often hide groups that are not directly labeled.
- Communities can reveal social circles, functional modules, disease subtypes, or coordinated behavior.
- The statistical question is whether the observed links contain enough signal to recover those groups.

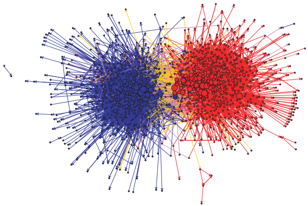

Figure: Political blogs in US prior to 2004 elections

Recovering communities from networks

- Networks often hide groups that are not directly labeled.
- Communities can reveal social circles, functional modules, disease subtypes, or coordinated behavior.
- The statistical question is whether the observed links contain enough signal to recover those groups.

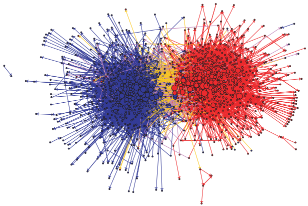

Figure: Political blogs in US prior to 2004 elections

A useful benchmark

(1) How much signal is necessary for recovering the group structure?
(2) Can an efficient algorithm reach that limit?

Stochastic Block Model (SBM)

[Holland-Laskey-Leinhardt 83]

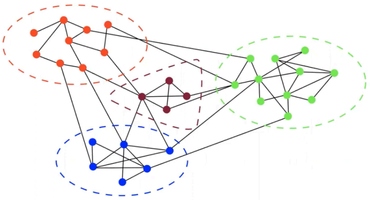

This talk: two balanced communities

- n nodes;
- $\sigma_i \in \{-1, +1\}$ uniform community labels;
- given $\{\sigma_i\}$, edges drawn independently:
  - if $\sigma_i = \sigma_j$, then $i \sim j$ with probability p;
  - if $\sigma_i \neq \sigma_j$, then $i \sim j$ with probability q.

Many works in physics, statistics, probability, CS, info theory . . . including:

- [Decelle-Krzakala-Moore-Zdeborová 11]
- [Mossel-Neeman-Sly 15,16,18]
- [Massoulié 14]
- [Abbé-Bandeira-Hall 14]
- [Abbé-Sandon 15,18]
- [Bordenave-Lelarge-Massoulié 15]
- [Abbé 17]
- [Mossel-Sly-Sohn 24, 25]
- . . .

Qustion: given the graph without community labels, can we recover the communities?

- partial (weak) recovery?
- almost-exact recovery?
- exact recovery?

Multiple correlated networks

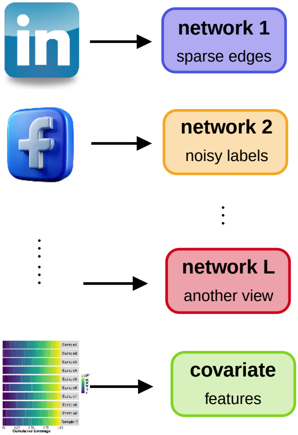

Question: can we synthesis information from multiple correlated networks/contextual features to better recover communities?

STOCHASTIC BLOCKMODELS: FIRST STEPS *

Paul W. HOLLAND

*Educational Testing Service ***

Kathryn Blackmond LASKEY and Samuel LEINHARDT

Carnegie-Mellon University †

lowercase letters. If $\boldsymbol{X}$ is a random adjacency array for g nodes and m relations, then the probability distribution of $\boldsymbol{X}$ is called a *stochastic multigraph*. We will denote the probability distribution of $\boldsymbol{X}$ by $p(\boldsymbol{x}) = \Pr(\boldsymbol{X} = \boldsymbol{x})$.

A *stochastic blockmodel* is a special case of a stochastic multigraph which satisfies the following requirements.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1, +1\}^n$: the ground-truth community label.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1, +1\}^n$: the ground-truth community label.
- $\boldsymbol{\sigma}_\ell \in \{-1, +1\}^n$: the community in the ℓ'th network.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1, +1\}^n$: the ground-truth community label.
- $\boldsymbol{\sigma}_\ell \in \{-1, +1\}^n$: the community in the ℓ'th network.
  - Different layers need not share exactly the same labels; they may only be correlated.
  - Suppose that $\boldsymbol{\sigma}(i) = \boldsymbol{\sigma}_\ell(i)$ with probability $\frac{1+\rho}{2}$, independently for all $1 \leq \ell \leq L$ and $1 \leq i \leq n$.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1, +1\}^n$: the ground-truth community label.
- $\boldsymbol{\sigma}_\ell \in \{-1, +1\}^n$: the community in the ℓ'th network.
  - Different layers need not share exactly the same labels; they may only be correlated.
  - Suppose that $\boldsymbol{\sigma}(i) = \boldsymbol{\sigma}_\ell(i)$ with probability $\frac{1+\rho}{2}$, independently for all $1 \le \ell \le L$ and $1 \le i \le n$.
- The ℓ'th network is a stochastic block model with community structure $\boldsymbol{\sigma}_\ell$, in-community density $\frac{(1+\epsilon_\ell)\lambda_\ell}{n}$, and out-community density $\frac{(1-\epsilon_\ell)\lambda_\ell}{n}$.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1, +1\}^n$: the ground-truth community label.
- $\boldsymbol{\sigma}_\ell \in \{-1, +1\}^n$: the community in the ℓ'th network.
  - Different layers need not share exactly the same labels; they may only be correlated.
  - Suppose that $\boldsymbol{\sigma}(i) = \boldsymbol{\sigma}_\ell(i)$ with probability $\frac{1+\rho}{2}$, independently for all $1 \le \ell \le L$ and $1 \le i \le n$.
- The ℓ'th network is a stochastic block model with community structure $\boldsymbol{\sigma}_\ell$, in-community density $\frac{(1+\epsilon_\ell)\lambda_\ell}{n}$, and out-community density $\frac{(1-\epsilon_\ell)\lambda_\ell}{n}$.
- Covariates add a high-dimensional feature view $\boldsymbol{Y} = \sqrt{\frac{\mu}{n}} \boldsymbol{\sigma} u^\top + \boldsymbol{Z} \in \mathbb{R}^{n \times p}$.

Contextual multi-layer SBM

- Considered in [Deshpande-Montanari-Mossel-Sen 18] [Ma-Nandy 23], [Yang-Lin-Sen 25].
- $\boldsymbol{\sigma} \in \{-1,+1\}^n$: the ground-truth community label.
- $\boldsymbol{\sigma}_\ell \in \{-1,+1\}^n$: the community in the ℓ'th network.
  - Different layers need not share exactly the same labels; they may only be correlated.
  - Suppose that $\boldsymbol{\sigma}(i) = \boldsymbol{\sigma}_\ell(i)$ with probability $\frac{1+\rho}{2}$, independently for all $1 \le \ell \le L$ and $1 \le i \le n$.
- The ℓ'th network is a stochastic block model with community structure $\boldsymbol{\sigma}_\ell$, in-community density $\frac{(1+\epsilon_\ell)\lambda_\ell}{n}$, and out-community density $\frac{(1-\epsilon_\ell)\lambda_\ell}{n}$.
- Covariates add a high-dimensional feature view $\boldsymbol{Y} = \sqrt{\frac{\mu}{n}}\boldsymbol{\sigma} u^\top + \boldsymbol{Z} \in \mathbb{R}^{n\times p}$.
- Our focus: sparse ($\lambda_\ell = O(1)$) + high-dimensional ($p \propto n$).

Single-view baseline

Sparse SBM

[Mossel-Neeman-Sly 15,18]: For one graph with average degree λ and edge contrast ϵ, the sharp weak-recovery threshold is the KS threshold

$$\epsilon^2 \lambda = 1\,.$$

Spiked matrix

[Baik-Ben Arous-Peché 05]: For high-dimensional covariates with $p/n = 1/\gamma$, the sharp BBP threshold is

$$\frac{\mu^2}{\gamma} = 1\,.$$

- Below these numbers, each individual view looks statistically weak.

Single-view baseline

Sparse SBM

[Mossel-Neeman-Sly 15,18]: For one graph with average degree λ and edge contrast ϵ, the sharp weak-recovery threshold is the KS threshold

$$\epsilon^2\lambda = 1\,.$$

Spiked matrix

[Baik-Ben Arous-Peché 05]: For high-dimensional covariates with $p/n = 1/\gamma$, the sharp BBP threshold is

$$\frac{\mu^2}{\gamma} = 1\,.$$

- Below these numbers, each individual view looks statistically weak.
- The main question is what happens when the weak views are correlated: can we recover the community label even when
  - Each network is below the KS threshold;
  - The covariate is below the BBP threshold?

Our results

Theorem (Gong-Huang-L. 26+)

Suppose that $L = o(n)$, $\lambda_\ell = o(n)$, and $p = n/\gamma$. Define

$$F = \max\left\{ \frac{\mu^2}{\gamma},\ \max_\ell \epsilon_\ell^2 \lambda_\ell,\ \frac{\mu^2}{\gamma} + \sum_{\ell=1}^{L} \frac{\rho^4 \epsilon_\ell^2 \lambda_\ell}{1 - (1-\rho^4)\epsilon_\ell^2 \lambda_\ell} \right\}.$$

(1) *When $F < 1$, weak recovery is information-theoretically impossible.*

(2) *When $F > 1$, weak recovery can be achieved by an $n^{O(\log L)}$ time algorithm.*

- The third term is the new gain from combining covariates with correlated network layers.

Sketch of lower bound

- Some natural ideas of proving the information-theoretic lower bound:

Sketch of lower bound

- Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);

Sketch of lower bound

- Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);

Sketch of lower bound

▷ Some natural ideas of proving the information-theoretic lower bound:
- [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
- [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
- seems analytically challenging in the multi-viewed + sparse setting.

Sketch of lower bound

- ▷ Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
  - seems analytically challenging in the multi-viewed + sparse setting.
- ▷ Our approach (motivated by [Mossel-Neeman-Sly 15]): consider the simpler detection problem between $\mathbb{P}$ (original model) and $\mathbb{Q}$ (independent noise).

Sketch of lower bound

- ▷ Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
  - seems analytically challenging in the multi-viewed + sparse setting.
- ▷ Our approach (motivated by [Mossel-Neeman-Sly 15]): consider the simpler detection problem between $\mathbb{P}$ (original model) and $\mathbb{Q}$ (independent noise).
  - We are able to show $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$.

Sketch of lower bound

- ▷ Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
  - seems analytically challenging in the multi-viewed + sparse setting.
- ▷ Our approach (motivated by [Mossel-Neeman-Sly 15]): consider the simpler detection problem between $\mathbb{P}$ (original model) and $\mathbb{Q}$ (independent noise).
  - We are able to show $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$.
  - Thus, if $L = O(1)$ detection is impossible; then recovery should be impossible (if we believe recovery is harder then detection).

Sketch of lower bound

- ▷ Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
  - seems analytically challenging in the multi-viewed + sparse setting.
- ▷ Our approach (motivated by [Mossel-Neeman-Sly 15]): consider the simpler detection problem between $\mathbb{P}$ (original model) and $\mathbb{Q}$ (independent noise).
  - We are able to show $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$.
  - Thus, if $L = O(1)$ detection is impossible; then recovery should be impossible (if we believe recovery is harder then detection).
  - **Question:** what if $L = \omega(1)$?

Sketch of lower bound

- ▷ Some natural ideas of proving the information-theoretic lower bound:
  - [Mossel-Neeman-Sly 15, Mossel-Sly-Sohn 25]: relating to a "local" problem (reconstruction on trees);
  - [Deshpande-Abbé-Montanari 17, Yang-Lin-Sen 25]: interpolate with a suitable Gaussian model (spiked random matrix);
  - seems analytically challenging in the multi-viewed + sparse setting.
- ▷ Our approach (motivated by [Mossel-Neeman-Sly 15]): consider the simpler detection problem between $\mathbb{P}$ (original model) and $\mathbb{Q}$ (independent noise).
  - We are able to show $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$.
  - Thus, if $L = O(1)$ detection is impossible; then recovery should be impossible (if we believe recovery is harder then detection).
  - **Question:** what if $L = \omega(1)$?
  - In fact, when $L = \omega(1)$ detection is possible for some $F < 1 \ldots$

Sketch of lower bound

Key observation:

- The bound $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$ implies that $\mathrm{Adv}(f) := \frac{\mathbb{E}_{\mathbb{P}}[f]}{\sqrt{\mathbb{E}_{\mathbb{Q}}[f^2]}} \leq e^{O(L)}$ for any measurable function f.

Sketch of lower bound

Key observation:

- The bound $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$ implies that $\mathsf{Adv}(f) := \frac{\mathbb{E}_{\mathbb{P}}[f]}{\sqrt{\mathbb{E}_{\mathbb{Q}}[f^2]}} \le e^{O(L)}$ for any measurable function f.
- However, if weak recovery is information theoretically possible, we can construct a measurable function f such that $\mathsf{Adv}(f) = e^{\Theta(n)} \Longrightarrow$ contradiction!

Sketch of lower bound

Key observation:

- The bound $\chi^2(\mathbb{P}\|\mathbb{Q}) = e^{O(L)}$ implies that $\mathsf{Adv}(f) := \frac{\mathbb{E}_{\mathbb{P}}[f]}{\sqrt{\mathbb{E}_{\mathbb{Q}}[f^2]}} \leq e^{O(L)}$ for any measurable function f.
- However, if weak recovery is information theoretically possible, we can construct a measurable function f such that $\mathsf{Adv}(f) = e^{\Theta(n)} \Longrightarrow$ contradiction!
- A diagram explaining the construction of f: a cross-validation trick (inspired by [Ding-Hua-Slot-Steurer 25]).

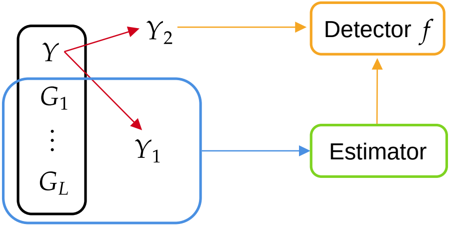

Algorithm heuristics

- A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way ...

Algorithm heuristics

- A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way ...
  - sparse SBM: algorithms based on counting self-avoiding/non-backtracking paths are used to achieved weak recovery up to KS threshold [Mossel-Neeman-Sly 18, Massoulié 14, Bordenave-Lelarge-Massoulié 15, Abbé-Sandon 15,18];

Algorithm heuristics

- ▷ A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way . . .
  - sparse SBM: algorithms based on counting self-avoiding/non-backtracking paths are used to achieved weak recovery up to KS threshold [Mossel-Neeman-Sly 18, Massoulié 14, Bordenave-Lelarge-Massoulié 15, Abbé-Sandon 15,18];
  - spiked rectangular matrix: spectral algorithm are used to achieved weak recovery up to BBP threshold [Baik-Ben Arous-Peché 05].

Algorithm heuristics

- A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way ...
  - sparse SBM: algorithms based on counting self-avoiding/non-backtracking paths are used to achieved weak recovery up to KS threshold [Mossel-Neeman-Sly 18, Massoulié 14, Bordenave-Lelarge-Massoulié 15, Abbé-Sandon 15,18];
  - spiked rectangular matrix: spectral algorithm are used to achieved weak recovery up to BBP threshold [Baik-Ben Arous-Peché 05].
- The first thing we need to do: try to either

Algorithm heuristics

- A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way ...
  - sparse SBM: algorithms based on counting self-avoiding/non-backtracking paths are used to achieved weak recovery up to KS threshold [Mossel-Neeman-Sly 18, Massoulié 14, Bordenave-Lelarge-Massoulié 15, Abbé-Sandon 15,18];
  - spiked rectangular matrix: spectral algorithm are used to achieved weak recovery up to BBP threshold [Baik-Ben Arous-Peché 05].
- The first thing we need to do: try to either
  - understand the path-counting approach in terms of spectral algorithms;

Algorithm heuristics

- A successful algorithm should at least deal with a single sparse SBM/a single spiked rectangular matrix, but currently they are solved in different way . . .
  - sparse SBM: algorithms based on counting self-avoiding/non-backtracking paths are used to achieved weak recovery up to KS threshold [Mossel-Neeman-Sly 18, Massoulié 14, Bordenave-Lelarge-Massoulié 15, Abbé-Sandon 15,18];
  - spiked rectangular matrix: spectral algorithm are used to achieved weak recovery up to BBP threshold [Baik-Ben Arous-Peché 05].
- The first thing we need to do: try to either
  - understand the path-counting approach in terms of spectral algorithms;
  - understand the spectral approach in terms of path counting.

Algorithmic heuristics

Let us focus on the spiked model $\boldsymbol{Y} = \frac{\lambda}{\sqrt{n}} x u^\top + \boldsymbol{W}$. PCA: determine spiked/non-spiked by $\|\boldsymbol{Y}\|_{\mathrm{op}}$.

Algorithmic heuristics

Let us focus on the spiked model $\boldsymbol{Y} = \frac{\lambda}{\sqrt{n}} x u^\top + \boldsymbol{W}$. PCA: determine spiked/non-spiked by $\|\boldsymbol{Y}\|_{\mathrm{op}}$. Trace method:

$$\begin{aligned} &\|\boldsymbol{Y}\|_{\mathrm{op}}^{2k} - \|\boldsymbol{W}\|_{\mathrm{op}}^{2k} \overset{k \gg (\log n)}{\approx} \mathrm{tr}(\boldsymbol{Y}^{2k}) - \mathrm{tr}(\boldsymbol{W}^{2k}) \\ &= \sum_{i_1, \ldots, i_{2k}} \boldsymbol{Y}_{i_1, i_2} \ldots \boldsymbol{Y}_{i_{2k}, i_1} - \boldsymbol{W}_{i_1, i_2} \ldots \boldsymbol{W}_{i_{2k}, i_1} \quad (\star) \end{aligned}$$

Algorithmic heuristics

Let us focus on the spiked model $\boldsymbol{Y} = \frac{\lambda}{\sqrt{n}} x u^\top + \boldsymbol{W}$. PCA: determine spiked/non-spiked by $\|\boldsymbol{Y}\|_{\mathrm{op}}$. Trace method:

$$\|\boldsymbol{Y}\|_{\mathrm{op}}^{2k} - \|\boldsymbol{W}\|_{\mathrm{op}}^{2k} \overset{k \gg (\log n)}{\approx} \mathrm{tr}(\boldsymbol{Y}^{2k}) - \mathrm{tr}(\boldsymbol{W}^{2k})$$
$$= \sum_{i_1,\ldots,i_{2k}} \boldsymbol{Y}_{i_1,i_2} \ldots \boldsymbol{Y}_{i_{2k},i_1} - \boldsymbol{W}_{i_1,i_2} \ldots \boldsymbol{W}_{i_{2k},i_1} \quad (\star)$$

Observation: $\gamma = (i_1, \ldots, i_{2k}, i_1)$ forming a "self-avoided" cycle is the main contribution of $(\star)$.

$$(\star) \approx \sum_{\gamma \text{ self-avoided}} \boldsymbol{Y}_\gamma - \boldsymbol{W}_\gamma$$
$$\|\boldsymbol{Y}\|_{\mathrm{op}}^{2k} \approx \sum_{\gamma \text{ self-avoided}} \boldsymbol{Y}_\gamma$$

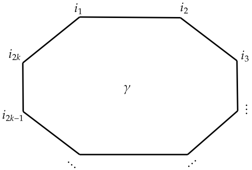

Detection to recovery: self-avoided cycle $\Longrightarrow$ self-avoided paths.

Algorithmic heuristics

Natural idea: paths $\Longrightarrow$ decorated (edge-colored) paths.

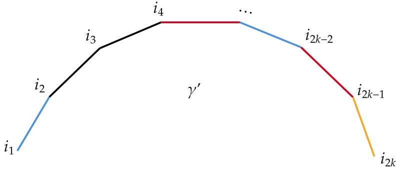

$$\phi_{\gamma'} := \boldsymbol{G}^{(1)}_{i_1,i_2} \boldsymbol{Y}_{i_2,i_3} \boldsymbol{Y}_{i_3,i_4} \ldots \boldsymbol{G}^{(2)}_{i_{2k-2},i_{2k-1}} \boldsymbol{G}^{(3)}_{i_{2k-1},i_{2k}}$$

Algorithmic heuristics

Natural idea: paths $\Longrightarrow$ decorated (edge-colored) paths.

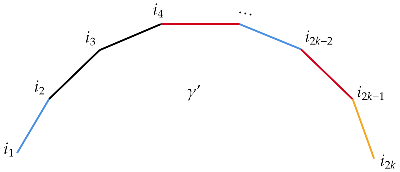

$$\phi_{\gamma'} := \boldsymbol{G}^{(1)}_{i_1,i_2} \boldsymbol{Y}_{i_2,i_3} \boldsymbol{Y}_{i_3,i_4} \ldots \boldsymbol{G}^{(2)}_{i_{2k-2},i_{2k-1}} \boldsymbol{G}^{(3)}_{i_{2k-1},i_{2k}}$$

Approach: consider a suitable **(weighted)** linear combination

$$\sum_{\gamma':\text{decorated path}} \Lambda(\gamma') \cdot \phi_{\gamma'} \,.$$

Summary and future perspectives

Take-home message:

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.
- Our information-theoretic lower bound use a recovery-to-detection reduction combined with a cross validation trick.

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.
- Our information-theoretic lower bound use a recovery-to-detection reduction combined with a cross validation trick.

Open problems:

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.
- Our information-theoretic lower bound use a recovery-to-detection reduction combined with a cross validation trick.

Open problems:

- Practical algorithms (perhaps spectral)?

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.
- Our information-theoretic lower bound use a recovery-to-detection reduction combined with a cross validation trick.

Open problems:

- Practical algorithms (perhaps spectral)?
- More than two communities? ([L. 26+]: evidence that statistical-computational gap will arise)

Summary and future perspectives

Take-home message:

- For the contextual multi-layer SBM with two symmetric groups, we determine the fundamental limits for weak recovery.
- Our algorithmic approach exploits the power of counting decorated paths.
- Our information-theoretic lower bound use a recovery-to-detection reduction combined with a cross validation trick.

Open problems:

- Practical algorithms (perhaps spectral)?
- More than two communities? ([L. 26+]: evidence that statistical-computational gap will arise)

Reference:

Shuyang Gong, Dong Huang, Zhangsong Li. Fundamental Limits of Community Detection in Contextual Multi-Layer Stochastic Block Models. arXiv:2602.08173.