Settling the Sharp Reconstruction Thresholds of Random Graph Matching

Yihong Wu
Department of Statistics and Data Science
Yale University
New Haven CT 06511, USA
Email: yihong.wu@yale.edu

Jiaming Xu and Sophie H. Yu
The Fuqua School of Business
Duke University
100 Fuqua Drive, Durham, NC 27708, USA
Email: {jx77, haoyang.yu}@duke.edu

2021 IEEE International Symposium on Information Theory (ISIT) | 978-1-5386-8209-8/21/$31.00 ©2021 IEEE | DOI: 10.1109/ISIT45174.2021.9518266

***Abstract*—This paper studies the problem of recovering the hidden vertex correspondence between two edge-correlated random graphs. We focus on the Gaussian model where the two graphs are complete graphs with correlated Gaussian weights and the Erdős-Rényi model where the two graphs are subsampled from a common parent Erdős-Rényi graph $\mathcal{G}(n,p)$. For dense graphs with $p = n^{-o(1)}$, we prove that there exists a sharp threshold, above which one can correctly match all but a vanishing fraction of the vertices and below which correctly matching any positive fraction is impossible, a phenomenon known as the "all-or-nothing" phase transition. Even more strikingly, in the Gaussian setting, above the threshold all vertices can be exactly matched with high probability. In contrast, for sparse Erdős-Rényi graphs with $p = n^{-\Theta(1)}$, we show that the all-or-nothing phenomenon no longer holds and we determine the thresholds up to a constant factor. Along the way, we also derive the sharp threshold for exact recovery, sharpening the existing results in Erdős-Rényi graphs [1], [2].**

The proof of the negative results builds upon a tight characterization of the mutual information based on the truncated second-moment computation in [3] and an "area theorem" that relates the mutual information to the integral of the reconstruction error. The positive results follows from a tight analysis of the maximum likelihood estimator that takes into account the cycle structure of the induced permutation on the edges.

A full version of this paper is accessible at: https://arxiv.org/pdf/2102.00082.pdf

I. Introduction

The problem of graph matching (or network alignment) refers to finding the underlying vertex correspondence between two graphs on the sole basis of their network topologies. Going beyond the worst-case intractability of finding the optimal correspondence (quadratic assignment problem [4], [5]), an emerging line of research is devoted to the average-case analysis of graph matching under meaningful statistical models, focusing on either information-theoretic limits [1]–[3], [6]–[8] or computationally efficient algorithms [9]–[15]. Despite these recent advances, the sharp thresholds of graph matching remain not fully understood especially for approximate reconstruction. The current paper aims to close this gap.

Following [11], [16], we consider the following probabilistic model for two random graphs correlated through a hidden vertex correspondence. Let the ground truth π be a uniformly random permutation on $[n]$. We generate two random weighted graphs on the common vertex set $[n]$ with (weighted) adjacency vectors $A = (A_{ij})_{1\le i<j\le n}$ and $B = (B_{ij})_{1\le i<j\le n}$ such that $\left(A_{\pi(i)\pi(j)}, B_{ij}\right)$ are i.i.d. pairs of correlated random variables with a joint distribution P, which implicitly depends on n. Of particular interest are the following two special cases:

- (Gaussian model): $P = \mathcal{N}\left(\binom{0}{0}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right)$ is the joint distribution of two standard Gaussian random variables with correlation coefficient $\rho > 0$. In this case, we have $B = \rho A^\pi + \sqrt{1-\rho^2} Z$, where A and Z are independent standard normal vectors and $A^\pi_{ij} = A_{\pi(i)\pi(j)}$.
- (Erdős-Rényi random graph): P denotes the joint distribution of two correlated $\mathrm{Bern}(q)$ random variables X and Y such that $\mathbb{P}\{Y = 1 \mid X = 1\} = s$, where $q \le s \le 1$. In this case, A and B are the adjacency vectors of two Erdős-Rényi random graphs $G_1, G_2 \sim \mathcal{G}(n,q)$, where G_1^π (with the adjacency vector A^π) and G_2 are independently edge-subsampled from a common parent graph $G \sim \mathcal{G}(n,p)$ with $p = q/s$.

Given the observation A and B, the goal is to recover the latent vertex correspondence π as accurately as possible. More specifically, given two permutations $\pi, \widehat{\pi}$ on $[n]$, denote the fraction of their overlap by

$$\mathsf{overlap}(\widehat{\pi}, \pi) \triangleq \frac{1}{n} \left|\{i \in [n] : \pi(i) = \widehat{\pi}(i)\}\right|.$$

Definition 1: We say an estimator $\widehat{\pi}$ of π achieves, as $n \to \infty$,

- *partial* recovery, if $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}, \pi) \ge \delta\} = 1 - o(1)$ for some constant $\delta \in (0,1)$;
- *almost exact* recovery, if $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}, \pi) \ge \delta\} = 1 - o(1)$ for any constant $\delta \in (0,1)$;
- *exact* recovery, if $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}, \pi) = 1\} = 1 - o(1)$.

The information-theoretic threshold of exact recovery has been determined for the Erdős-Rényi graph model [2] in the regime of $p = O\left(\log^{-3}(n)\right)$ and more recently for the Gaussian model [8]; however, the results and proof techniques in [2] do not hold for denser graphs. In contrast, approximate recovery are far less well understood. Apart from the sharp condition for almost exact recovery in the sparse regime $p = n^{-\Omega(1)}$ [6], only upper and lower bounds are known for partial recovery [7]. See Section III for a detailed review of these previous results.

In this paper, we characterize the sharp reconstruction thresholds for both the Gaussian and dense Erdős-Rényi graphs with $p = n^{-o(1)}$. Specifically, we prove that there exists a sharp threshold, above which one can estimate all but a vanishing fraction of the latent permutation and below which recovering any positive fraction is impossible, a phenomenon known as the "all-or-nothing" phase transition [17]. This phenomenon is even more striking in the Gaussian model, in the sense that above the threshold the hidden permutation can be estimated error-free with high probability. In contrast, for sparse Erdős-Rényi graphs with $p = n^{-\Theta(1)}$, we show that the all-or-nothing phenomenon no longer holds. To this end, we determine the threshold for partial recovery up to a constant factor and show that it is order-wise smaller than the almost exact recovery threshold found in [6].

Along the way, we also derive a sharp threshold for exact recovery, sharpening existing results in [1], [2]. As a byproduct, the same technique yields an alternative proof of the result in [8] for the Gaussian model.

II. Main results

Throughout the paper, we let $\epsilon > 0$ denote an arbitrarily small but fixed constant. Let $\widehat{\pi}_{\mathsf{ML}}$ denote the maximum likelihood estimator (MLE), which reduces to

$$\widehat{\pi}_{\mathsf{ML}} \in \arg\max_{\pi'} \langle A^{\pi'}, B \rangle \,. \tag{1}$$

Since the prior distribution of the latent permutation π is uniform, the MLE coincides with the maximum a posteriori (MAP) estimator in [2] for the Erdős-Rényi graphs. Note that the MLE may not maximize $\mathbb{E}\left[\mathsf{overlap}(\pi, \widehat{\pi})\right]$.

Theorem 1 (Gaussian model): If

$$\rho^2 \geq \frac{(4+\epsilon)\log n}{n}, \tag{2}$$

then $\mathbb{P}\left\{\mathsf{overlap}\left(\widehat{\pi}_{\rm ML}, \pi\right) = 1\right\} = 1 - o(1)$.

Conversely, if

$$\rho^2 \leq \frac{(4-\epsilon)\log n}{n}, \tag{3}$$

then for any estimator $\widehat{\pi}$ and any fixed constant $\delta > 0$, $\mathbb{P}\left\{\mathsf{overlap}\left(\widehat{\pi}, \pi\right) \leq \delta\right\} = 1 - o(1)$.

Theorem 1 implies that in the Gaussian setting, the recovery of π exhibits a sharp phase transition in terms of the limiting value of $\frac{n\rho^2}{\log n}$ at threshold 4, above which exact recovery is possible and below which even partial recovery is impossible. The positive part of Theorem 1 was first shown in [8]. Here we provide an alternative proof that does not rely on the Gaussian property and works for Erdős-Rényi graphs as well.

The next result determines the sharp threshold for the Erdős-Rényi model in terms of the key quantity nps^2, the average degree of the intersection graph $G_1 \wedge G_2$ (whose edges are sampled by both G_1 and G_2).

Theorem 2 (Erdős-Rényi graphs, dense regime): Assume that p is bounded away from 1 and that $p = n^{-o(1)}$. If

$$nps^2 \geq \frac{(2+\epsilon)\log n}{\log\frac{1}{p} - 1 + p}, \tag{4}$$

then for any constant $\delta < 1$, $\mathbb{P}\left\{\mathsf{overlap}\left(\widehat{\pi}_{\rm ML}, \pi\right) \geq \delta\right\} = 1 - o(1)$. Conversely, if

$$nps^2 \leq \frac{(2-\epsilon)\log n}{\log\frac{1}{p} - 1 + p}, \tag{5}$$

then for any estimator $\widehat{\pi}$ and any constant $\delta > 0$, $\mathbb{P}\left\{\mathsf{overlap}\left(\widehat{\pi}, \pi\right) \leq \delta\right\} = 1 - o(1)$.

Theorem 2 implies that analogous to the Gaussian model, in dense Erdős-Rényi graphs, the recovery of π exhibits an "all-or-nothing" phase transition in terms of the limiting value of $\frac{nps^2\left(\log\frac{1}{p} - 1 + p\right)}{\log n}$ at threshold 2, above which almost exact recovery is possible and below which even partial recovery is impossible. However, as we will see in Theorem 4, unlike the Gaussian model, this threshold differs from that of exact recovery.

Remark 1 (Entropy interpretation of the thresholds): The sharp thresholds in Theorem 1 and Theorem 2 can be interpreted from an information-theoretic perspective. Suppose an estimator $\widehat{\pi} = \widehat{\pi}(A, B)$ achieves almost exact recovery with $\mathbb{E}[\mathsf{overlap}(\widehat{\pi}, \pi)] = 1 - o(1)$, which, by a rate-distortion computation, implies that $I(\pi; \widehat{\pi})$ must be close to the full entropy of π, that is, $I(\pi; \widehat{\pi}) = (1 - o(1)) n \log n$. On the other hand, by the data processing inequality, we have $I(\pi; \widehat{\pi}) \leq I(\pi; A, B)$. The latter can be bounded simply as (see Section IV-A)

$$I(\pi; A, B) \leq \binom{n}{2} I(P),$$

where $I(P)$ denote the mutual information between a pair of random variables with joint distribution P. For the Gaussian model, we have

$$I(P) = \frac{1}{2}\log\frac{1}{1-\rho^2}.$$

For the correlated Erdős-Rényi graph,

$$I(P) = q d(s\|q) + (1-q) d(\eta\|q), \tag{6}$$

where $q = ps$, $\eta \triangleq \frac{q(1-s)}{1-q}$, and $d(s\|q) \triangleq D(\mathrm{Bern}(s)\|\mathrm{Bern}(q))$ denotes the binary KL divergence. By Taylor expansion, we have $I(P) = s^2 p\left(p - 1 + \log\frac{1}{p}\right)(1 - o(1))$ when $q = o(1)$. Combining these with $\binom{n}{2} I(P) \geq (1 - o(1)) n \log n$ shows that almost exact recovery is impossible under the conditions (3) and (5). The impossibility of partial recovery under the same conditions takes more effort to show, which we do in Section IV-A.

Theorem 3 (Erdős-Rényi graphs, sparse regime): Assume $p = n^{-\Omega(1)}$. If

$$nps^2 \geq (2+\epsilon)\max\left\{\frac{\log n}{\log(1/p)}, 2\right\}, \tag{7}$$

then there exists a constant $\delta > 0$ such that $\mathbb{P}\left\{\mathsf{overlap}\left(\widehat{\pi}_{\rm ML}, \pi\right) \geq \delta\right\} = 1 - o(1)$. Conversely, assuming $np = \omega(\log^2 n)$, if

$$nps^2 \leq 1 - \epsilon, \tag{8}$$

then for any estimator $\widehat{\pi}$ and any constant $\delta > 0$, $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}, \pi) \leq \delta\} = 1 - o(1)$.

Theorem 3 implies that for sparse Erdős-Rényi graphs with $p = n^{-\alpha}$ for a constant $\alpha \in (0, 1)$, the information-theoretic thresholds for partial recovery is at $nps^2 \asymp 1$, which is much lower than the almost exact recovery threshold $nps^2 = \omega(1)$ as established in [6]. Hence, interestingly the all-or-nothing phenomenon no longer holds for sparse Erdős-Rényi graphs. Note that the conditions (7) and (8) differ by a constant factor. Determining the sharp threshold for partial recovery in the sparse regime remains an open question.

Finally, we address the exact recovery threshold in the Erdős-Rényi graph model. For ease of notation, we consider a general correlated Erdős-Rényi graph model specified by the joint distribution $P = (p_{ab} : a, b \in \{0, 1\})$, so that $\mathbb{P}\left\{A_{\pi(i)\pi(j)} = a, B_{ij} = b\right\} = p_{ab}$ for $a, b \in \{0, 1\}$. In this general Erdős-Rényi model, $\widehat{\pi}_{\mathsf{ML}}$ is again given by the maximization problem (1) if $p_{11}p_{00} > p_{01}p_{10}$ (positive correlation) and changes to minimization if $p_{11}p_{00} < p_{01}p_{10}$ (negative correlation). The subsampling model is a special case with positive correlation, where

$$p_{11} = ps^2, \quad p_{10} = p_{01} = ps(1-s), \quad p_{00} = 1 - 2ps + ps^2. \tag{9}$$

Theorem 4 (Erdős-Rényi graphs, exact recovery): Under the subsampling model (9), if

$$n\left(\sqrt{p_{00}p_{11}} - \sqrt{p_{01}p_{10}}\right)^2 \geq (1+\epsilon)\log n, \tag{10}$$

then $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}_{\mathrm{ML}}, \pi) = 1\} = 1 - o(1)$.

Conversely, if

$$n\left(\sqrt{p_{00}p_{11}} - \sqrt{p_{01}p_{10}}\right)^2 \leq (1-\epsilon)\log n, \tag{11}$$

then for any estimator $\widehat{\pi}$, $\mathbb{P}\{\mathsf{overlap}(\widehat{\pi}, \pi) = 1\} = o(1)$.

Assume that p is bounded away from 1. Then Theorem 4 implies that the exact recovery threshold is given by $\lim_{n\to\infty} \frac{nps^2(1-\sqrt{p})^2}{\log n} = 1$. Since $\log\frac{1}{p} - 1 + p \geq 2(1-\sqrt{p})^2$, with equality if and only if $p = 1$, the threshold of exact recovery is strictly higher than that of almost exact recovery in the Erdős-Rényi graph model, unlike the Gaussian model.

III. Comparisons to Prior work

A. Exact recovery

The information-theoretic thresholds for exact recovery have been determined for the Gaussian model and the general Erdős-Rényi graph model in certain regimes. In particular, for the Gaussian model, it is shown in [8] that if $n\rho^2 \geq (4+\epsilon)\log n$, then the MLE given in (1) achieves exact recovery; if instead $n\rho^2 \leq (4-\epsilon)\log n$, then exact recovery is impossible. Theorem 1 significantly strengthens this negative result by showing that under the same condition even partial recovery is impossible.

Analogously, for Erdős-Rényi random graphs, it is shown in [1], [2] that the MLE achieve exact recovery when $nps^2 = \log n + \omega(1)$ under the additional restriction that $p = O(\log^{-3}(n))$. Conversely, exact recovery is shown in [1] to be information-theoretically impossible provided that $nps^2 \leq \log n - \omega(1)$, based on the fact the intersection graph $G_1 \wedge G_2 \sim \mathcal{G}(n, ps^2)$ has many isolated nodes with high probability. In comparison, Theorem 4 implies that the precise exact recovery threshold is instead given by $\lim_{n\to\infty} \frac{nps^2\left(\sqrt{1-2ps+ps^2} - \sqrt{p}(1-s)\right)^2}{\log n} = 1$. In particular, deriving the tight condition (11) requires more than eliminating isolated nodes. See the discussions in Section IV-C for details.

B. Partial and almost exact recovery

Compared to exact recovery, the understanding of partial and almost exact recovery is less precise. It is shown in [6] that in the sparse regime $p = n^{-\Omega(1)}$, almost exact recovery is information-theoretically possible if and only if $nps^2 = \omega(1)$ The more recent work [7] further investigates partial recovery. It is shown that if $nps^2 \geq C(\delta)\max\left\{1, \frac{\log n}{\log(1/p)}\right\}$, then there exists an exponential-time estimator $\widehat{\pi}$ that achieves $\mathsf{overlap}(\widehat{\pi}, \pi) \geq \delta$ with high probability, where $C(\delta)$ is some large constant that tends to ∞ as $\delta \to 1$; conversely, if $\frac{I(P)}{\delta} = o\left(\frac{\log(n)}{n}\right)$ with $I(P)$ given in (6), then no estimator can achieve $\mathsf{overlap}(\widehat{\pi}, \pi) \geq \delta$ with positive probability. These conditions do not match in general and are much looser than the results in Theorems 2 and 3.

IV. Proof Sketch

A. Impossibility of partial recovery

As a first step, we characterize the asymptotic value of the mutual information $I(A, B; \pi)$. By definition,

$$\begin{aligned} I(A, B; \pi) &\triangleq \mathbb{E}\left[D\left(\mathcal{P}_{A,B|\pi} \| \mathcal{P}_{A,B}\right)\right] \\ &= \mathbb{E}\left[D\left(\mathcal{P}_{A,B|\pi} \| \mathcal{Q}_{A,B}\right)\right] - D\left(\mathcal{P}_{A,B} || \mathcal{Q}_{AB}\right) \end{aligned}$$

for any joint distribution $\mathcal{Q}_{A,B}$ of (A, B) such that $D\left(\mathcal{P}_{A,B} || \mathcal{Q}_{AB}\right) < \infty$. Note that $\mathcal{P}_{A,B|\pi}$ factorizes into a product distribution $\prod_{i<j} \mathcal{P}_{A_{\pi(i)\pi(j)}, B_{ij}} = P^{\otimes\binom{n}{2}}$, where P is the joint distribution of $(A_{\pi(i)\pi(j)}, B_{ij})$. Thus, to exploit the tensorization property of the KL-divergence, we choose $\mathcal{Q}_{A,B}$ to be a product distribution under which A and B are independent and (A_{ij}, B_{ij}) are i.i.d. pairs of independent random variables with a joint distribution Q with the same marginals as P. (We shall refer to this $\mathcal{Q}_{A,B}$ as the null model.) In particular, for the Gaussian (resp. Erdős-Rényi) model, Q is the joint distribution of two independent standard normal (resp. $\mathrm{Bern}(q)$) random variables. Under this choice, we have $D\left(\mathcal{P}_{A,B|\pi} \| \mathcal{Q}_{A,B}\right) = \binom{n}{2} D(P\|Q) = \binom{n}{2} I(P)$ and hence

$$I(A, B; \pi) = \binom{n}{2} I(P) - D\left(\mathcal{P}_{A,B} || \mathcal{Q}_{AB}\right).$$

By the non-negativity of the KL divergence, we have $I(A, B; \pi) \leq \binom{n}{2} I(P)$. This bound turns out to be tight, as made precise by the following proposition. To prove this result, we leverage the previous truncated second moment computation in [3] to conclude that $\mathcal{D}(\mathcal{P}_{A,B} \| \mathcal{Q}_{A,B})$ is negligible under the desired conditions.

Proposition 1: It holds that

$$I(A,B;\pi)=\binom{n}{2}I(P)-\zeta_n,$$

where

- $\zeta_n=o(1)$ in the Gaussian model with $\rho^2\le\frac{(4-\epsilon)\log n}{n}$;
- $\zeta_n=o(1)$ in the dense Erdős-Rényi graphs with $p=n^{-o(1)}$ and $nps^2\left(\log(1/p)-1+p\right)\le(2-\epsilon)\log(n)$;
- $\zeta_n=O(\log n)$ in the sparse Erdős-Rényi graphs with $p=n^{-\Omega(1)}$ and $np=\omega(1)$ and $nps^2\le 1-\epsilon$.

Given the tight characterization of the mutual information in Proposition 1, we now relate it to the Bayes risk. Using the chain rule, we have

$$I(A,B;\pi)=I(B;\pi\mid A)=I(B;A^\pi\mid A),$$

where the second equality follows from the fact that $A\to A^\pi\to B$ forms a Markov chain. The intuition is that conditioned on A, B is a noisy observation of A^π (which is random owning to π). In such a situation, the mutual information can typically be related to an integral of the reconstruction error of the signal A^π. To make this precise, we first introduce a parametric model $\mathcal{P}_\theta$ that interpolates between the planted model $\mathcal{P}$ and the null model $\mathcal{Q}$ as θ varies. We write $\mathbb{E}_\theta$ to indicate expectation taken with respect to the law $\mathcal{P}_\theta$.

For the Gaussian model, let $\mathcal{P}_\theta$ denote the model under which $B=\sqrt{\theta}A^\pi+\sqrt{1-\theta}Z$, where A,Z are two independent Gaussian matrices and $\theta\in[0,1]$. Then $\theta=\rho^2$ corresponds to the planted model $\mathcal{P}$ while $\theta=0$ corresponds to the null model $\mathcal{Q}$. As θ increases from 0 to ρ^2, $\mathcal{P}_\theta$ interpolates between $\mathcal{Q}$ and $\mathcal{P}$. Let

$$\mathrm{mmse}_\theta(A^\pi)\triangleq\mathbb{E}_\theta[\|A^\pi-\mathbb{E}_\theta[A^\pi|A,B]\|^2]$$

denote the minimum mean-squared error (MMSE) of estimating A^π based on (A,B) distributed according to $\mathcal{P}_\theta$. The following proposition follows from the celebrated I-MMSE formula [18].

Proposition 2 (Gaussian model):

$$I(A,B;\pi)=\frac{1}{2}\int_0^{\rho^2}\frac{\mathrm{mmse}_\theta(A^\pi)}{(1-\theta)^2}d\theta\,.$$

The correlated Erdős-Rényi graph model requires more effort. Let us fix $q=ps$ and consider the following coupling P_θ between two $\mathrm{Bern}(q)$ random variables with joint probability mass function p_θ, where $p_\theta(11)=q\theta$, $p_\theta(01)=p_\theta(10)=q(1-\theta)$, and $p_\theta(00)=1-(2-\theta)q$, with $\theta\in[q,s]$. Let $\mathcal{P}_\theta$ denote the interpolated model under which $\left(A_{\pi(i)\pi(j)},B_{ij}\right)$ are i.i.d. pairs of correlated random variables with joint distribution P_θ. As θ increases from q to s, $\mathcal{P}_\theta$ interpolates between the null model $\mathcal{Q}=\mathcal{P}_q$ and the planted model $\mathcal{P}=\mathcal{P}_s$. Following [19], [20], we prove the following area theorem that relates $I(A,B;\pi)$ to the MMSE of A^π.

Proposition 3 (Erdős-Rényi random graph): It holds that

$$\begin{aligned}I(A,B;\pi)&\le\binom{n}{2}I(P)+\binom{n}{2}qs^2\\&+\int_q^s\frac{\theta-q}{s(1-q)^2}\left(\mathrm{mmse}_\theta(A^\pi)-\binom{n}{2}q(1-q)\right)d\theta.\end{aligned}$$

The above two steps together imply that the MMSE of A^π given the observation (A,B) is asymptotically equal to the estimation error of the trivial estimator $\mathbb{E}\left[A^\pi\right]$, which further asymptotically equals $\binom{n}{2}q=\mathbb{E}\left[\|A\|^2\right]$ when $q=o(1)$. Finally, we connect the MMSE of A^π to the Hamming loss of reconstructing π, concluding the impossibility of the partial recovery.

Proposition 4: In both the Gaussian and Erdős-Rényi graph model, if $\mathrm{mmse}_\theta(A^\pi)\ge\mathbb{E}\left[\|A\|^2\right](1-\xi)$ for some $\xi>0$, then for any estimator $\widehat{\pi}=\widehat{\pi}(A,B)$,

$$\mathbb{E}_\theta[\mathsf{overlap}(\widehat{\pi},\pi)]\le O\left(\xi^{1/4}+\left(\frac{n\log n}{\mathbb{E}\left[\|A\|^2\right]}\right)^{1/4}\right).$$

B. Possibility of partial recovery

Let $\mathcal{S}_n$ denote the set of permutations on the node set $[n]$. For any two permutations $\pi,\pi'\in\mathcal{S}_n$, let $d(\pi,\pi')$ denote the number of non-fixed points in the $\pi'\circ\pi^{-1}$. The following proposition provides sufficient conditions for $\widehat{\pi}_{\mathsf{ML}}$ defined in (1) to achieve the partial recovery and almost exact recovery in Erdős-Rényi graphs.

Proposition 5: Let $ps\le\frac{1}{2}$. [1] Suppose that

- if $p=1-o(1)$,
$$\frac{ns^2(1-p)^2}{(1-ps)^2}\ge(4+\epsilon)\log n\,;\qquad(12)$$
- if $p\le 1-c_0$ for some constant c_0,
$$nps^2\ge\begin{cases}\frac{(2+\epsilon)\log n}{\log(1/p)-1+p}&\text{if }p\ge n^{-\frac{1}{2}}\\4+\epsilon&\text{if }p<n^{-\frac{1}{2}}\end{cases}.\qquad(13)$$

Then there exists a constant $\delta=\delta(\epsilon,c_0)<1$ such that

$$\mathbb{P}\left\{d\left(\widehat{\pi}_{\mathsf{ML}},\pi\right)<\delta n\right\}\ge 1-n^{-1+o(1)}.$$

If in addition $nps^2(1-p)^2=\omega(1)$, then for any constant $\delta>0$,

$$\mathbb{P}\left\{d\left(\widehat{\pi}_{\mathsf{ML}},\pi\right)<\delta n\right\}\ge 1-n^{-1+o(1)}.$$

We remark that in the dense regime of $n^{-o(1)}\le p\le 1-c_0$, (13) already implies that $nps^2=\omega(1)$ and hence the MLE achieves almost exact recovery provided $nps^2\ge\frac{(2+\epsilon)\log n}{\log(1/p)-1+p}$; this proves the positive part of Theorem 2. In contrast, in the sparse regime of $p=n^{-\Omega(1)}$, the MLE achieves the almost exact recovery provided that $nps^2=\omega(1)$, which is in fact needed for any estimator to succeed [6].

The proof of Proposition 5 relies on the following key lemma, which bounds the probability that the ML estimator (1) makes a given number of errors.

Lemma 1: Let $\epsilon\in(0,1)$ be an arbitrary constant and $ps\le\frac{1}{2}$.

[1] Note that the existence/absence of an edge is a matter of representation and they are mathematically equivalent. As a consequence, by flipping 0 and 1, the model with parameter (n,p,s) is equivalent to that with parameter (n,p',s') for an appropriate choice of p' and s' such that $p's'=1-ps$. Thus we can assume $ps\le 1/2$ without loss of generality.

- For the Erdős-Rényi model, suppose that either (12) or (13) holds. Then there exists some constant $0 < \delta < 1$ such that for any $k \geq \delta n$,

$$\mathbb{P}\left\{d\left(\widehat{\pi}_{\mathsf{ML}}, \pi\right)=k\right\} \leq 2 \exp\left(-n h\left(\frac{k}{n}\right)\right) \mathbf{1}_{\{k \leq n-1\}} + e^{-2\log n}\mathbf{1}_{\{k=n\}} + \exp\left(-\frac{1}{64}\epsilon k \log n\right), \qquad (14)$$

  where $h(x) = -x\log x - (1-x)\log(1-x)$ is the binary entropy function.
  If in addition $nps^2(1-p)^2 = \omega(1)$, then (14) holds for any constant $0<\delta<1$ and all $k \geq \delta n$.
- For the Gaussian model, suppose that $n\rho^2 \geq (4+\epsilon)\log n$. Then (14) holds for any constant $0 < \delta < 1$ and all $k \geq \delta n$.

The proof of Lemma 1 uses the cycle structure of permutations (cf. [3, Section 3.1] for more details and examples). For each $\sigma \in \mathcal{S}_n$, let σ^{E} denote the induced permutation of σ on the edge set $\binom{[n]}{2}$ of unordered pairs, according to $\sigma^{\mathsf{E}}((i,j)) \triangleq (\sigma(i), \sigma(j))$. We shall refer to σ and σ^{E} as a node permutation and edge permutation. Each permutation can be decomposed as disjoint cycles known as orbits. Orbits of σ (resp. σ^{E}) are referred as *node orbits* (resp. edge orbits). Let F be the set of fixed points of σ. Let $\mathcal{O}$ be the collection of all edge orbits of σ^{E}. Denote $\mathcal{O}_1 = \binom{F}{2} \subset \mathcal{O}$, which is a subset of fixed points of edge permutation σ^{E},

Lemma 1 follows from a large deviation analysis of the maximum likelihood estimator (1). A crucial observation is that the difference of the objective function in (1) evaluated at a given permutation π' and the ground truth π can be decomposed across the edge orbits of $\sigma \triangleq \pi^{-1} \circ \pi'$ as

$$\left\langle A^{\pi'} - A^{\pi}, B\right\rangle = \sum_{O \in \mathcal{O}\backslash\mathcal{O}_1} X_O - \sum_{O \in \mathcal{O}\backslash\mathcal{O}_1} Y_O \triangleq X - Y,$$

where $X_O \triangleq \sum_{(i,j)\in O} A_{\pi'(i)\pi'(j)} B_{ij}$, and $Y_O \triangleq \sum_{(i,j)\in O} A_{\pi(i)\pi(j)} B_{ij}$, are independent across edge orbits O. Crucially, Y depends on π' only through its fixed point set F, which has substantially fewer choices than π' itself when $n - |F| \asymp n$. Therefore, for the purpose of applying the union bound it is beneficial to separately control X and Y. Indeed, we show that Y is highly concentrated on its mean. Hence, it remains to analyze the large-deviation event of X exceeding $\mathbb{E}[Y]$, which is accomplished by a careful computation of the moment generation function (MGF) $M_{|O|} \triangleq \mathbb{E}\left[\exp\left(tX_O\right)\right]$ and proving that

$$M_{|O|} \leq M_2^{|O|/2}, \text{ for } |O| \geq 2. \qquad (15)$$

Intuitively, it means that the contribution of longer edge orbits can be effectively bounded by that of the 2-edge orbits. Capitalizing on this key finding and applying the Chernoff bound together with a union bound over π' yield Lemma 1.

C. Exact recovery

Building upon the almost exact recovery results in the preceding section, we now analyze the MLE for exact recovery. In parallel with Lemma 1, the following lemma gives a tighter bound on the probability that the MLE makes a small number of errors.

Lemma 2: Suppose that for any constant $0 < \epsilon < 1$,

- for general Erdős-Rényi random graphs, if $n\left(\sqrt{p_{00}p_{11}} - \sqrt{p_{01}p_{10}}\right)^2 \geq (1+\epsilon)\log n$;
- for Gaussian model, if $n\rho^2 \geq (4+\epsilon)\log n$;

then for any $k \in [n]$ such that $k \leq \frac{\epsilon}{16}n$,

$$\mathbb{P}\left\{d\left(\widehat{\pi}_{\mathsf{ML}}, \pi\right) = k\right\} \leq \exp\left(-\frac{\epsilon}{8} k \log n\right). \qquad (16)$$

The positive result in Theorem 4 readily follows from the union bound on k by applying (14) for large k and (16) for small k.

To prove Lemma 2, we need to consider π' that is close to π, i.e., $d(\pi, \pi') \leq \epsilon n/16$. In this regime, the number of choices of F is comparable to that of π'. Hence, instead of separately bounding X and Y, it is more favorable to directly applying the Chernoff bound to the difference $X - Y$. Crucially, the moment generation function $\mathbb{E}\left[\exp\left(t(X_O - Y_O)\right)\right]$ continues to satisfy the relation (15) and the bottleneck for exact recovery happens at $|F| = n-2$, where π' differs from π by a 2-cycle (transposition).

Prompted by this observation, we prove a matching necessary condition of exact recovery in Theorem 4 by considering permutations $\sigma \triangleq \pi^{-1} \circ \pi'$ that consists of $n-2$ fixed points and a 2-node orbit (i,j), for some $i,j \in [n]$, in which case,

$$\left\langle A^{\pi'} - A^{\pi}, B\right\rangle = -\sum_{k \in [n]\backslash\{i,j\}} \underbrace{\left(A^{\pi}_{ik} - A^{\pi}_{jk}\right)\left(B_{ik} - B_{jk}\right)}_{\triangleq X_{ij,k}}.$$

Note that $X_{ij,k} \overset{\text{i.i.d.}}{\sim} a\delta_1 + b\delta_{-1} + (1-a-b)\delta_0$ with $a = 2p_{00}p_{11}$ and $b = 2p_{01}p_{10}$ for all $k \neq i,j$. There remains two key challenges to conclude the existence of many choices of (i,j) for which $\langle A^{\pi'}, B\rangle \geq \langle A^{\pi}, B\rangle$. First, to derive a tight impossibility condition, we need to obtain a tight large-deviation lower estimate for this event. Second, the RHS of the above equation is correlated for different pairs of (i,j). This dependency is addressed by fixing a subset $T \subset [n]$ with $|T| = o(n)$ and breaking the summation in the above equation into two terms X_{ij} and Y_{ij}, where $X_{ij} \triangleq \sum_{k \in T^c\backslash\{i,j\}} X_{ij,k}$. Then we further partition T as $T_1 \cup T_2$ and separately bound X_{ij} and Y_{ij} for any $i \in T_1$ and $j \in T_2$. Since $|T| = o(n)$, a simple application of Markov's inequality ensures that $Y_{ij} \leq \tau$ for a relatively small threshold τ. Moreover, by construction, crucially X_{ij} and $X_{i'j'}$ are pairwise independent for any $i \neq i'$ in T_1 and $j \neq j'$ in T_2. Thus we apply a second-moment calculation to show that there exist many choices of $(i,j) \in T_1 \times T_2$ such that $X_{ij} \leq -\tau$.

Acknowledgment

This research was supported in part by NSF Grant IIS-1838124, CCF-1850743, CCF-1856424, CCF-1900507, an NSF CAREER award CCF-1651588, and an Alfred Sloan fellowship.

REFERENCES

[1] D. Cullina and N. Kiyavash, "Improved achievability and converse bounds for Erdős-Rényi graph matching," *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, no. 1, pp. 63–72, 2016.

[2] ——, "Exact alignment recovery for correlated Erdős-Rényi graphs," *arXiv preprint arXiv:1711.06783*, 2017.

[3] Y. Wu, J. Xu, and S. H. Yu, "Testing correlation of unlabeled random graphs," *arXiv preprint arXiv:2008.10097*, 2020.

[4] P. M. Pardalos and H. Wolkowicz, *Quadratic Assignment and Related Problems: DIMACS Workshop, May 20-21, 1993*. American Mathematical Soc., 1994, vol. 16.

[5] R. E. Burkard, E. Cela, P. M. Pardalos, and L. S. Pitsoulis, "The quadratic assignment problem," in *Handbook of combinatorial optimization*. Springer, 1998, pp. 1713–1809.

[6] D. Cullina, N. Kiyavash, P. Mittal, and H. V. Poor, "Partial recovery of Erdős-Rényi graph alignment via k-core alignment," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 3, no. 3, pp. 1–21, 2019.

[7] G. Hall and L. Massoulié, "Partial recovery in the graph alignment problem," *arXiv preprint arXiv:2007.00533*, 2020.

[8] L. Ganassali, "Sharp threshold for alignment of graph databases with gaussian weights," *arXiv preprint arXiv:2010.16295*, 2020.

[9] S. Feizi, G. Quon, M. Recamonde-Mendoza, M. Médard, M. Kellis, and A. Jadbabaie, "Spectral alignment of networks," *arXiv preprint arXiv:1602.04181*, 2016.

[10] V. Lyzinski, D. Fishkind, M. Fiori, J. Vogelstein, C. Priebe, and G. Sapiro, "Graph matching: Relax at your own risk," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 38, no. 1, pp. 60–73, 2016.

[11] J. Ding, Z. Ma, Y. Wu, and J. Xu, "Efficient random graph matching via degree profiles," *Probability Theory and Related Fields*, pp. 1–87, Sep 2020.

[12] B. Barak, C.-N. Chou, Z. Lei, T. Schramm, and Y. Sheng, "(Nearly) efficient algorithms for the graph matching problem on correlated random graphs," in *Advances in Neural Information Processing Systems*, 2019, pp. 9186–9194.

[13] Z. Fan, C. Mao, Y. Wu, and J. Xu, "Spectral graph matching and regularized quadratic relaxations I: The Gaussian model," *arxiv preprint arXiv:1907.08880*, 2019.

[14] ——, "Spectral graph matching and regularized quadratic relaxations II: Erdős-Rényi graphs and universality," *arxiv preprint arXiv:1907.08883*, 2019.

[15] L. Ganassali and L. Massoulié, "From tree matching to sparse graph alignment," in *Conference on Learning Theory*. PMLR, 2020, pp. 1633–1665.

[16] P. Pedarsani and M. Grossglauser, "On the privacy of anonymized networks," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2011, pp. 1235–1243.

[17] G. Reeves, J. Xu, and I. Zadik, "The all-or-nothing phenomenon in sparse linear regression," in *Conference on Learning Theory*. PMLR, 2019, pp. 2652–2663.

[18] D. Guo, S. Shamai, and S. Verdú, "Mutual information and minimum mean-square error in gaussian channels," *IEEE Trans. on Information Theory*, vol. 51, 2005.

[19] C. Méasson, A. Montanari, T. J. Richardson, and R. Urbanke, "The generalized area theorem and some of its consequences," *IEEE Transactions on Information Theory*, vol. 55, no. 11, pp. 4793–4821, 2009.

[20] Y. Deshpande, E. Abbe, and A. Montanari, "Asymptotic mutual information for the balanced binary stochastic block model," *Information and Inference: A Journal of the IMA*, vol. 6, no. 2, pp. 125–170, 2017.