The graph alignment problem (for more than two graphs)

Research domain introduction

Louis Vassaux

Supervisor : Laurent Massoulié
Institut National de Recherche en Informatique et en Automatique

Contents

I would like to thank my supervisor Laurent Massoulié, as well as the entire ARGO research team, for welcoming me in so warmly to INRIA.

1 Introduction to graph alignment

1.1 The graph alignement problem

There are various statistical problems where we observe a network of linked objects which can be represented by a large graph; for instance, a network of protein-protein interactions or a social network. The (planted) **graph alignment problem**, in a broad sense, is the following: we observe the same underlying data twice, obtaining two **unlabeled** graphs which should be very similar, and attempt to retrieve the underlying isomorphism which pairs up any two nodes representing the same element.

Formally, the problem can be described as follows. Assume that we have an underlying pair of labeled random graphs (G, G') (say, with node set $[\![1, n]\!]$) which are meant to represent the same underlying data, possibly with some noise. Mathematically, this translates to the fact that, for any $i, j \in [\![1, n]\!]$, the random variables

$$\mathbb{1}\{i \sim j \text{ in } G\} \text{ and } \mathbb{1}\{i \sim j \text{ in } G'\} \tag{1.1}$$

are correlated. We then pick a permutation $\pi^* \in \mathcal{S}_n$ uniformly at random (independently from G, G'), obtaining a new pair of random graphs $(G, \pi^*(G'))$ where

$$\pi^*(i) \sim \pi^*(j) \text{ in } \pi^*(G') \Leftrightarrow i \sim j \text{ in } G' \tag{1.2}$$

Our aim is to reconstruct the permutation π^* from the observation of $(G, \pi^*(G'))$ ie. to build an estimator

$$\begin{aligned} \hat{\pi} : \{\text{graphs on } [\![1, n]\!]\}^2 &\to \mathcal{S}_n \\ (G, \pi^*(G')) &\mapsto \pi^* \end{aligned}$$

Ideally, we would like to guarantee that $\hat{\pi} = \pi^*$ all of the time, but this is a very difficult combinatorial problem. As a result, in recent years, a focus has instead been put on this probabilistic ("planted") version of the problem, where we try to have $\hat{\pi} = \pi^*$ with high probability (whp) as $n \to +\infty$. In fact, even this **exact alignment** condition is not necessarily required for practical purposes! We may simply wish to maximise the overlap between $\hat{\pi}$ and π^*:

$$ov(\hat{\pi}, \pi^*) = \frac{1}{n} \#\{u \in [\![1, n]\!], \hat{\pi}(u) = \pi^*(u)\} \tag{1.3}$$

The **almost exact alignment** problem consists of trying to guarantee that $\text{ov}(\pi^*, \hat{\pi}) = 1 - o(1)$ whp; the **partial alignment** problem is where we try to guarantee that $\text{ov}(\pi^*, \hat{\pi}) \geq \delta > 0$ whp. As we shall see, depending on the specific random graph model, we may expect different objectives to be feasible.

Finally, let us mention a couple of other relevant problems:

- when achieving partial alignment, it is reasonable to expect to know which nodes have been correctly aligned: this is the **one-sided** partial alignment problem, and this is often what interests us in practice. The same idea applies to almost exact alignment.
- we would also like to find polynomial-time algorithms, in order to actually be able to align graphs in practice. This is not always easy! As we shall see, there are situations where an estimator exists but we suspect that it is impossible to compute it in polynomial time. Searching for algorithms that work well, and run quickly, on real-world graphs is also its own (very interesting) field of study which we will not look into here.

1.2 Aligning two Erdős–Rényi graphs

A well-studied toy model is the case where (G, G') are correlated Erdős–Rényi graphs. Let $0 \leq p \leq 1$ be a density parameter, and $0 \leq s \leq 1$ be a correlation parameter; then:

1. We generate an Erdős–Rényi graph $H \sim G(n, p)$.
2. We subsample H to obtain G by keeping any given edge of H with probability s.
3. We subsample H independently in the same way to obtain the graph G'.

Ultimately, this means that any given edge is in $G \cap G'$ with probability ps^2; is in $G \setminus G'$ with probability $ps(1-s)$; is in $G' \setminus G$ with probability $ps(1-s)$, and is absent otherwise.

The previously mentioned problems have more or less been solved for this model. In particular, whether or not alignment is possible depends mostly upon the value of ps^2, ie. the parameter of the Erdős–Rényi graph $G \cap G'$.

Theorem 1.1 (Various authors, 2017-22)**.** *Let (G, G') be as before, and assume some mild sparsity conditions (for instance, $p \leq n^{-\varepsilon}$ for some $\varepsilon > 0$). Then:*

- *exact alignment feasibility has a hard cutoff at $ps^2 = \frac{\log n}{n}$ (the threshold below which $G \cap G'$ starts having isolated vertices) [2];*
- *almost exact alignment is possible iff $ps^2 = \omega(\frac{1}{n})$ ie. if $G \cap G'$ has $o(n)$ isolated vertices [3];*
- *partial alignment feasibility has a hard cutoff at $ps^2 = \frac{1}{n}$ ([11], [4]) (the threshold below which $G \cap G'$ no longer has a giant component).*

Showing that these conditions are necessary is generally pretty tractable; showing that they are sufficient, on the other hand, can be quite difficult. In particular, almost exact and partial alignment involve working over certain specific, highly connected subgraphs of G and G' (respectively, their k-cores and their maximally dense subgraphs): these are the subgraphs which we know we can align.

It is also unclear exactly how much of the graphs we can align in both cases; this is an interesting but difficult topic, particularly in the case of partial alignment.

Let us finally mention the question of finding a polynomial-time algorithm for alignment. The relevant threshold here seems to be $s^2 = \alpha \simeq 0.34$, where α is Otter's constant (the number of unlabeled trees of size n is roughly α^{-n}). Specifically (again under mild sparsity conditions),

- if $s^2 > \alpha$, exact and almost exact recovery are feasible in polynomial time iff they are information-theoretically feasible [13]. This may also be the case for partial recovery, which can be done in polynomial time over a large part of the region $s^2 > \alpha$ ([13], [12]).
- if $s^2 < \alpha$, recent work by Ding et al. [5] seems to point to poly-time algorithms not existing, even if alignment is theoretically possible. Proving this seems very difficult since even the graph isomorphism problem is not known to be NP-complete.

1.3 Aligning two Gaussian graphs

The setup here is slightly different than that described previously, since the graphs here are weighted. Here, we simply require that, for any $i, j \in [\![1, n]\!]$, the weights G_{ij} and G'_{ij} are correlated. (This generalises the unweighted graph case if we consider an unweighted graph to be a weighted graph with weights in $\{0, 1\}$.)

Naturally, the weights chosen here are Gaussian: for any $i < j$, (G_{ij}, G'_{ij}) are standard centred Gaussians with covariance $\rho \in [0, 1]$. This toy model has again been thoroughly investigated, and the information-theoretical limits for alignment are known.

Theorem 1.2 (Wu et al, Ganassalli (independently); 2020)**.** *Let*

$$\rho_0 = \sqrt{\frac{4 \log n}{n}} \tag{1.4}$$

Then, if $\rho > (1 + \varepsilon)\rho_0$ for some $\varepsilon > 0$, exact alignment is possible; if $\rho < (1 - \varepsilon)\rho_0$, then even partial alignment is impossible.

This all-or-nothing phenomenon may seem very different from the Erdős–Rényi case, but this is in fact not necessarily the case. This mostly stems from the fact that the Gaussian model resembles a very dense Erdős–Rényi model, which was excluded from the previous propositions in many cases (because real-life unweighted graphs tend to be sparser).

As before, we may also ask for which values of ρ we can find polynomial-time algorithms to align Gaussian graphs. A lot of work has been put into spectral alignment methods in this case: aligning the graphs so as to make the eigenvectors of their adjacency matrices as close as possible. The Gaussian model is particularly well-suited to this sort of analysis since its adjacency matrix is essentially a GOE matrix. Another direction currently being explored (see [9], [14]) is a a convex relaxation of the problem, where we solve a related convex optimisation problem over the convex hull of the set of permutation matrices in $\mathcal{M}_n(\mathbb{R})$, and hope to extract from this optimum the correct alignment.

1.4 Aligning more than two graphs

The graph alignment problem can be fairly naturally generalised to 3 graphs or more. Here, instead of building an estimator $\hat{\pi}$ for some permutation $\pi^* : G_2 \to G_1$, we want an estimator $(\hat{\pi_2}, \hat{\pi_3}, \dots, \hat{\pi_p})$ for some underlying permutation $(\pi_2^*, \pi_3^*, \dots, \pi_p^*) : G_2 \times \dots G_p \to G_1$. Until a year ago, nothing was known about this problem from a theoretical perspective.

Last year, the exact alignment problem for multiple Erdos-Renyi graphs was solved in [1]. Specifically:

Theorem 1.3 (Ameen & Hajek, 2024)**.** *Assume some sparsity conditions. Then, exact alignment for p graphs $G_1, G_2, \dots, G_p$ is possible iff we can exactly align G_1 and $G_2 \cup \dots \cup G_p$. (This means that we have a hard cutoff at $\lambda s(1-(1-s)^{p-1}) = \frac{\log n}{n}$.)*

One direction is fairly clear: if we can align the G_i, then we can align G_1 and $G_2 \cup \dots \cup G_p$. We will discuss later on why the other implication is a fairly reasonable sufficient condition.

2 Main results

2.1 Aligning p Gaussian graphs

The precise mathematical question is the following. We have p weighted graphs $(G^1, G^2, \dots, G^p)$ over $[\![1, n]\!]$ such that, for $1 \leq i < j \leq n$, $(G_{ij}^{(1)}, \dots, G_{ij}^{(p)})$ are iid (centred) Gaussian vectors, such that

$$\operatorname{Var} G_{ij}^{(k)} = 1; \operatorname{Cov}(G_{ij}^{(k)}, G_{ij}^{(l)}) = \rho \tag{2.1}$$

for $1 \leq k \neq l \leq p$. Can we align all of our graphs?

The result for 2 graphs suggests that there is a hard cutoff, above which exact alignment is possible, and below which even partial alignment is impossible. Indeed, this is the case in general.

Theorem 2.1 (-, Massoulié; 2025)**.** *Let*

$$\rho_0 = \sqrt{\frac{c_p \log n}{n}}, \text{ where } c_p = \frac{8}{p} \tag{2.2}$$

Then, if $\rho > (1+\varepsilon)\rho_0$ for some $\varepsilon > 0$, exact alignment is possible; if $\rho < (1-\varepsilon)\rho_0$, then even partial alignment is impossible.

We will outline the proof of this theorem later on; the full proof is in [15].

This result shows that increasing the number of graphs should *a priori* make it significantly easier to align them, at least from an information-theoretical point of view.

2.2 Aligning many Erdős–Rényi graphs

Here, we generate $p \geq 3$ correlated Erdős–Rényi graphs using the same procedure as in Section 1.2; instead of subsampling H twice, we subsample it p times to obtain a p-uple of graphs $(G_1, \dots, G_p)$. As already noted, the exact alignment problem has already been solved in this; we focus on the partial alignment problem, in the sparse regime where the graphs have $O(n)$ edges. As in the two-graph problem, the sparse case seems to be more complex than the dense case.

A first note is that there are two different partial alignment problems which we can ask here:

- given the data (G_1, G_2, G_3), is there a subset X of size at least εn on which we can align all three graphs?

- given the data (G_1, G_2, G_3), are there subsets X_1 (resp. X_2, X_3) on which we can align G_2 and G_3 (resp. G_1 and G_3, G_1 and G_2)?

It is clear that solving the first problem implies that we can solve the second one. It also *seems* clear that solving the second problem should mean that we can solve the first one: indeed, even if the subsets X_1, X_2, X_3 were independent in some sense, they should have a nontrivial intersection, and in practice we actually expect them to be correlated. Nevertheless, it has not been proven that the two problems are equivalent and there does not appear to be a simple argument. In the rest of this paper, we focus on answering the first question (since it is mathematically easier) with the understanding that it is likely equivalent to the second one.

As in the dense case, there is a fairly natural necessary condition for graph alignment being possible:

Proposition 2.2. *If we cannot align G_1 and $G_2 \cup \ldots \cup G_p$, then we cannot align $(G_1, \ldots, G_p)$.*

This means that, setting $p = \frac{\lambda}{n}$, alignment is impossible when $\lambda s(1-(1-s)^{p-1}) = 1$. (Note that, due to the asymmetry between the two graphs, the fact that this is the true cutoff requires some extra work; this is done in [15].) In the dense case, as we saw, the converse of Proposition 2.2 is true; it seems plausible that this is also the case for the sparse case.

3 A general overview of our methods

There are various different approaches used in the literature to construct estimators for the true alignment, or to show that no such estimators are accurate. In this section, we sketch the ideas which we employed to show Theorems 2.1 and 2.2, which we believe to be quite general; however, they are by no means the only methods which can be used.

3.1 Maximum a posteriori estimation

Recall that we wish to show that exact alignment is possible when $\rho > (1+\varepsilon)\rho_0$ (resp. that partial alignment is impossible otherwise).

For exact alignment, it is known that there is (theoretically) a single optimal estimator: the **maximum a posteriori** estimator, ie.

$$\hat{\pi}_{MAP} = \text{argmax}_\pi \, \mathbb{P}(\pi^* = \pi | \text{ all observable variables}) \tag{3.1}$$

We say theoretically because in practical applications this may be completely impractical to evaluate; however, this is not our concern here since we are looking for theoretical bounds on graph alignment feasibility.

One may also ask whether this notion of maximum a posteriori estimation can be adapted to handle the problems of almost exact, or partial, alignment. It turns out that this is the case: the following is a specialisation of a more general result from Bayesian decision theory.

Proposition 3.1. *Consider, for $r \geq 0$, the (closed) ball $B(0,r) \subseteq \mathcal{S}_n$ for the distance function $d = 1 - ov$ (see (1.3)). We define the random variable*

$$C_n(r) = \max_{\pi \in \mathcal{S}_n} \mathbb{P}(B(\pi, r) | \textit{ all observable variables}) \tag{3.2}$$

which measures to what extent the random measure $\mathbb{P}_{post}$ concentrates around a single point. Then,

1. *a partial estimator exists iff, for some $0 \leq r < 1$, $C_n(r) \xrightarrow[\mathbb{P}]{} 1$.*
2. *an almost exact estimator exists iff, for any $\varepsilon > 0$, $C_n(\varepsilon) \xrightarrow[\mathbb{P}]{} 1$.*

This follows from the fact that, much like the MAP estimator, the estimator

$$\hat{\pi} = \arg\max_{\pi \in \mathcal{S}_n} \mathbb{P}(B(\pi, r) | \text{ all observable variables}). \tag{3.3}$$

is "optimal" for alignment up to distance r (in a certain precise sense).

The previous proposition also generalises effortlessly to our multi-graph case: here, the observable variables are the graphs $G_1, \ldots, G_p$; the overlap function also needs to be chosen suitably (this is linked to the different notions of alignment discussed in section 2.2).

3.2 The Gaussian model: a Gibbs measure

Consider the multi-graph Gaussian model defined previously. We may explicitly compute:

$$\mathbb{P}(\boldsymbol{\pi}^* = \boldsymbol{\sigma} | G_1, \ldots, G_p) = \frac{1}{\mathcal{Z}} e^{-\beta \mathcal{H}(\boldsymbol{\sigma})} \tag{3.4}$$

where $\beta, \mathcal{Z}$ are independent from σ, and

$$-\mathcal{H}(\boldsymbol{\sigma}) = \sum_{1 \leq u < v \leq n} \sum_{1 \leq i \neq j \leq p} G^{(i)}_{\pi_i^*(uv)} G^{(j)}_{\sigma_j \sigma_i^{-1} \pi_i^*(uv)} \tag{3.5}$$

Note the appearance of what is known as a Gibbs measure: the form of this random probability measure mirrors that of other random probability measures studied in the context of spin-glass systems. In fact, as in the spin-glass context, the key to understanding our measure lies in computing the so-called quenched free energy $\mathbb{E}[\log \mathcal{Z}]$. However, tackling the problem directly through this lens seems quite difficult: not only are rigourous results on these systems quite difficult to obtain, but in our case we have the extra difficulties of having an unusual state space as well as a nonisotropic Hamiltonian.

Thankfully, here we can tackle our random measure through more elementary methods: what we really care about here are joint tail bounds on the $(\mathcal{H}(\boldsymbol{\sigma}))_{\boldsymbol{\sigma} \in \mathcal{S}_n^{p-1}}$, and these can be obtained using standard Chernoff-style bounds. For the impossiblity of partial alignemnt, the situation is significantly more challenging as we need to use some forms of anti-concentration while paying extra attention to the correlation structure of the Hamiltonians over $\mathcal{S}_n^{p-1}$, but the general idea remains similar.

3.3 The sparse case with 3 Erdős–Rényi graphs

The situation here is very much less nice than for the dense Gaussian model - or even the dense version of the Erdős–Rényi model. As we mentioned, there is a nice topological condition for the dense case: it is shown in [1] that we can align $G_1, G_2, \ldots, G_p$ iff we can

align G_1 with $G_2 \cup \ldots \cup G_p$. We will begin by explaining why this also seems like a reasonable candidate for the cutoff, in the case of sparse Erdős–Rényi graphs.

To begin with, let us understand why we can align G_1, G_2 only if their intersection graph has a giant component. In order to prove that alignment is impossible, the easiest way would be to create "fake" alignments which look just like the genuine alignment to anyone observing the unlabeled graphs, but which are equal to the real alignment almost nowhere. In [11], the authors do just that: if you take any automorphism σ of $G_1 \cap G_2$, then it is quite likely that the alignements π^* and $\sigma \circ \pi^*$ are entirely indistinguishable from the point of view of the observer. When $G_1 \cap G_2$ has no giant component, there are many such σ: in particular, there are many such σ which have almost no fixed points, hence the result.

Now, when $G_1 \cap G_2$ does indeed have a giant component, we cannot really flip this argument around: though there does not seem to be a reasonable way to create a fake alignment in a similar fashion, it would not be entirely implausible that there exist some magical fake alignments, which appear simply because there are so many permutations to choose from that one of them may look reasonable - purely by random chance. As it turns out, this is not the case - if there is indeed a giant component, then the densest subgraph of $G_1 \cap G_2$ can be identified and aligned properly [4].

Can we apply the same sort of reasoning to p graphs? Well, if the real alignment is $\boldsymbol{\pi}^* = (\pi_2, \pi_3, \ldots, \pi_p)$ and σ is an automorphism of $H_1 = G_1 \cap (\pi_2(G_2) \cup \pi_3(G_3) \cup \ldots \cup \pi_p(G_p))$, then $\sigma \circ \boldsymbol{\pi}^* = (\sigma \circ \pi_2, \sigma \circ \pi_3, \ldots, \sigma \circ \pi_p)$ is a plausible fake alignment. The reason for our choice of H_1 is that, informally, it is the minimal graph we can build from the $\pi_i(G_i)$, such that any automorphism σ of H_1 will (more or less) preserve the total number of edges of the union graph $\bigcup_{i=1}^n \pi_i(G_i)$. Using this construction, we can indeed show Theorem 2.2, regarding the impossiblity of partial alignment. We will also note that Proposition 3.1 allows us to sidestep many of the technical details from Ganassalli's paper [11]: an elaborate Poisson-related argument is needed there to build very rigid automorphisms, whereas the framework above allows us to relax this and only build approximate automorphisms.

Unfortunately, to prove feasibility of partial alignment, we cannot just hope to align the G_i by aligning the densest subgraphs of G_1 and $\bigcup_i \pi_i^*(G_i)$, since we do not a priori have access to the second graph! The general methods from [7] do not work here either since they require decomposing the alignment into disjoint cycles; this breaks down when we have multiple mappings whose cycle decompositions may me incompatible. We still believe that this condition is the true cutoff for partial alignment, for two reasons.

1. Clearly, this would fit rather well into the pattern given by the results for two graphs and the known result for exact alignment for more than two graphs.

2. More interestingly, there is a probabilistic ansatz which points to this being true. Consider a permutation π chosen uniformly at random from $\mathcal{S}_n$; what does it action on the set $\{0,1\}^{\binom{n}{2}}$ of possible edges over $[\![1, n]\!]$ look like? The number of fixed points of such a permutation follows a law which, for our purposes, can be approximated by $\binom{N}{2}$ with N following a Poisson law of parameter λ; the ansatz is that we may hope to approximate the action of the permutation on the rest of the graph by a uniformly random permutation of all of the remaining edges, conditional on having no fixed points. If this were the case, one could show that the cutoff is indeed at the expected point $\lambda s(1 - (1-s)^{p-1}) = 1$.

4 Future directions for research

The field of graph alignment is still quite young and there are a lot of questions which remain unanswered. Here are a few which are likely both within reach and interesting/useful.

- Can we get a full understanding of the limits of alignment of many Erdős–Rényi graphs? Obviously, our current result for partial alignment is one-sided, and it would be nice to be able to generalise this to a complete cutoff result.
- Can we get a better understanding of the cutoff for polynomial-time graph alignment? This is well-understood in the Erdős–Rényi case, but unclear in the situation of Gaussian graph alignment - and it is not known how adding more graphs to the problem would affect poly-time feasibility.
- Can we generalise our methods and results to other graph models (say, preferential attachment or geometric) which better represent real data in a number of situations? Some work has been done here (see [16], [6]) but more remains to be done.
- Can we obtain theoretical understandings of the algorithms which are used in real-world situations? These are often mathematically messier and may rely on some extra information (or even sometimes black-box neural networks), but obtaining some form of theoretical guarantees would still be nice. Ideally, this could also allow us to improve upon the performance of these algorithms.

5 Bonus

During the work mentioned above, the author was led to prove the following result.

Theorem 5.1. *Let X be a (connected) weighted graph over $[\![1,n]\!]$, with positive weights $(x_{ij})_{1\leq i<j\leq n}$. Define the norms*

$$\|X\|_2 = \sqrt{\frac{1}{\binom{n}{2}} \sum_{1\leq i<j\leq n} x_{ij}^2} \ , \ \|X\|_\infty = \max_{1\leq i<j\leq n} x_{ij}. \tag{5.1}$$

Then, there exists a spanning tree of X with total weight at least

$$(p-1)\frac{2\|X\|_2}{\|X\|_2 + \|X\|_\infty}. \tag{5.2}$$

A proof of this result can be found in [15, Appendix L]; it involves induction and a lot of computations. The author is still searching for a nice explanation of why this is true; any insights are welcome.

References

[1] Taha Ameen and Bruce Hajek. *Exact Random Graph Matching with Multiple Graphs*. 2024. arXiv: 2405.12293 [cs.DS]. URL: https://arxiv.org/abs/2405.12293.

[2] Daniel Cullina and Negar Kiyavash. "Exact alignment recovery for correlated Erdos Renyi graphs". In: *ArXiv* abs/1711.06783 (2017). URL: https://api.semanticscholar.org/CorpusID:126055588.

[3] Daniel Cullina et al. "Partial Recovery of Erdős-Rényi Graph Alignment via k-Core Alignment". In: *ACM SIGMETRICS Performance Evaluation Review* 48 (July 2020), pp. 99–100. DOI: 10.1145/3410048.3410106.

[4] Jian Ding and Hang Du. "Detection Threshold for Correlated Erdős-Rényi Graphs via Densest Subgraph". In: *IEEE Transactions on Information Theory* 69.8 (2023), pp. 5289–5298. DOI: 10.1109/TIT.2023.3265009.

[5] Jian Ding, Hang Du, and Zhangsong Li. *Low-Degree Hardness of Detection for Correlated Erdős-Rényi Graphs*. 2023. arXiv: 2311.15931 [cs.DS]. URL: https://arxiv.org/abs/2311.15931.

[6] Mathieu Even et al. *Aligning Embeddings and Geometric Random Graphs: Informational Results and Computational Approaches for the Procrustes-Wasserstein Problem*. 2024. arXiv: 2405.14532 [stat.ML]. URL: https://arxiv.org/abs/2405.14532.

[7] Xiequan Fan, Ion Grama, and Quansheng Liu. "Sharp large deviation results for sums of independent random variables". In: *Science China Mathematics* 58 (Sept. 2015), pp. 1939–1958. DOI: 10.1007/s11425-015-5049-6.

[8] Zhou Fan et al. "Spectral Graph Matching and Regularized Quadratic Relaxations II: Erdős-Rényi Graphs and Universality". In: *Found. Comput. Math.* 23.5 (June 2022), pp. 1567–1617. ISSN: 1615-3375. DOI: 10.1007/s10208-022-09575-7. URL: https://doi.org/10.1007/s10208-022-09575-7.

[9] Zhou Fan et al. "Spectral graph matching and regularized quadratic relaxations: algorithm and theory". In: *Proceedings of the 37th International Conference on Machine Learning*. ICML'20. JMLR.org, 2020.

[10] Luca Ganassali. "Sharp threshold for alignment of graph databases with Gaussian weights". In: *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*. Ed. by Joan Bruna, Jan Hesthaven, and Lenka Zdeborova. Vol. 145. Proceedings of Machine Learning Research. PMLR, 16–19 Aug 2022, pp. 314–335. URL: https://proceedings.mlr.press/v145/ganassali22a.html.

[11] Luca Ganassali, Laurent Massoulie, and Marc Lelarge. "Impossibility of Partial Recovery in the Graph Alignment Problem". In: *Proceedings of Thirty Fourth Conference on Learning Theory*. Ed. by Mikhail Belkin and Samory Kpotufe. Vol. 134. Proceedings of Machine Learning Research. PMLR, 15–19 Aug 2021, pp. 2080–2102. URL: https://proceedings.mlr.press/v134/ganassali21a.html.

[12] Luca Ganassali, Laurent Massoulié, and Guilhem Semerjian. "Statistical limits of correlation detection in trees". In: *The Annals of Applied Probability* 34.4 (2024), pp. 3701–3734. DOI: 10.1214/23-AAP2048. URL: https://doi.org/10.1214/23-AAP2048.

[13] Cheng Mao et al. “Random Graph Matching at Otter’s Threshold via Counting Chandeliers”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (2022). URL: https://api.semanticscholar.org/CorpusID:252531751.

[14] Sushil Mahavir Varma, Irène Waldspurger, and Laurent Massoulié. *Graph Alignment via Birkhoff Relaxation*. 2025. arXiv: 2503.05323 [stat.ML]. URL: https://arxiv.org/abs/2503.05323.

[15] Louis Vassaux and Laurent Massoulié. *The feasibility of multi-graph alignment: a Bayesian approach*. 2025. arXiv: 2502.17142 [math.ST]. URL: https://arxiv.org/abs/2502.17142.

[16] Haoyu Wang et al. “Random Graph Matching in Geometric Models: the Case of Complete Graphs”. In: *Proceedings of Thirty Fifth Conference on Learning Theory*. Ed. by Po-Ling Loh and Maxim Raginsky. Vol. 178. Proceedings of Machine Learning Research. PMLR, Feb. 2022, pp. 3441–3488. URL: https://proceedings.mlr.press/v178/wang22a.html.