

Sharp threshold for alignment of graph databases with Gaussian weights

Luca Ganassali

LUCA.GANASSALI@INRIA.FR

INRIA, ENS, PSL Research University, Paris, France.

Editors: Joan Bruna, Jan S Hesthaven, Lenka Zdeborova

Abstract

We study the fundamental limits for reconstruction in weighted graph (or matrix) database alignment. We consider a model of two graphs where π^* is a planted uniform permutation and all pairs of edge weights $(A_{i,j}, B_{\pi^*(i),\pi^*(j)})_{1 \leq i < j \leq n}$ are i.i.d. pairs of Gaussian variables with zero mean, unit variance and correlation parameter $\rho \in [0, 1]$. We prove that there is a sharp threshold for exact recovery of π^* : if $n\rho^2 \geq (4 + \varepsilon) \log n + \omega(1)$ for some $\varepsilon > 0$, there is an estimator $\hat{\pi}$ – namely the MAP estimator – based on the observation of databases A, B that achieves exact reconstruction with high probability. Conversely, if $n\rho^2 \leq 4 \log n - \log \log n - \omega(1)$, then any estimator $\hat{\pi}$ verifies $\hat{\pi} = \pi$ with probability $o(1)$.

This result shows that the information-theoretic threshold for exact recovery is the same as the one obtained for detection in a recent work by [Wu et al. \(2020\)](#): in other words, for Gaussian weighted graph alignment, the problem of reconstruction is not more difficult than that of detection. Though the reconstruction task was already well understood for vector-shaped database alignment (that is taking signal of the form $(u_i, v_{\pi^*(i)})_{1 \leq i \leq n}$ where $(u_i, v_{\pi^*(i)})$ are i.i.d. pairs in $\mathbb{R}^{d_u} \times \mathbb{R}^{d_v}$), its formulation for graph (or matrix) databases brings a drastically different problem for which the hard phase is conjectured to be wide.

The proofs build upon the analysis of the MAP estimator and the second moment method, together with the study of the correlation structure of energies of permutations.

Keywords: high-dimensional statistics, matrix alignment, information-theoretic limits, Gaussian matrices

Introduction

Aligning databases We address the following problem: suppose that we have two databases consisting in weighted graphs represented by their adjacency matrices A and B . For simplicity, assume that the two graphs have same size and that each individual appears in both graphs. For a given individual, its attached signal consists in weighted edges with all other users. Across databases, edges that correspond to pairs of matched individuals are correlated. We consider the following question: *if the graphs are shown unlabeled (that is, if users are anonymized), is it possible to recover the corresponding matching between databases by aligning them at the sight of their correlation structure?*

Intuitively, when the matrices are correlated enough, one can learn the true matching between individuals present in the databases. In this study we investigate the precise conditions on correlation under which exact reconstruction (or perfect de-anonymization) is feasible with high probability.

De-anonymization problems aroused great interest when [Narayanan and Shmatikov \(2008\)](#) were able to de-anonymize an unlabeled dataset of film ratings (namely, the Netflix prize dataset) with the observation of a publicly available database (namely the Internet Movie Database), using correlations between the ratings. Since then, they have been studied in recent literature, in several

versions and reformulations. The range of applications has been widened to quantifying privacy issues related to databases (Dwork (2008)) or social networks (Narayanan and Shmatikov (2009)).

Widespread attention was given on the *graph alignment problem*, focusing on more geometrical databases (Cullina and Kiyavash (2017); Cullina et al. (2018a); Ding et al. (2018); Fan et al. (2019a,b); Ganassali and Massoulié (2020)). Lots of other natural applications can be mentioned, such as *pattern recognition* in image processing (Berg et al. (2005); Cour et al. (2007)), *aligning protein interaction networks* in computational biology (Singh et al. (2008)) or performing *ontology alignment* in natural language processing (Haghighi et al. (2005)).

Vector-shaped and graph-shaped databases From the theoretical point of view, fundamental limits for the deanonymisation problem are now well understood when data only consists in vectors u, v of given sizes n (Cullina et al. (2018b); Emre Dai et al. (2019)), that is when each user has its own signal, regardless of its connections with others. In this setting, the problem can be phrased in terms of a *Linear Assignment Problem (LAP)*:

$$\arg \max_{\Pi} \langle \Pi u, v \rangle, \quad (1)$$

where the maximum runs over all permutation matrices of size n . Even if greedy optimization is easily seen to be exponential-time, LAP can be solved efficiently in $O(n^3)$ steps using the classical Hungarian algorithm (Kuhn (1955)).

Another related problem is that of linear regression with an unknown permutation, studied in Pananjady et al. (2016): this time, one observes $y = \Pi^* A x^* + w$, where $x^* \in \mathbb{R}^d$ is an unknown vector, Π^* is an unknown $n \times n$ permutation matrix, and $w \in \mathbb{R}^n$ is additive Gaussian noise. Here again, the permutation Π^* applies only on the left side of A , which corresponds to row permutation.

On the other hand, when the databases are graphs, the problem is different and can be phrased this time in terms of a *Quadratic Assignment Problem (QAP)*:

$$\arg \max_{\Pi} \langle A, \Pi B \Pi^T \rangle. \quad (2)$$

A significant difference with the previous vector-shaped setting is that this problem is known to be NP-hard in the worst case, as well as some of its approximations (Makarychev et al. (2014); Pardalos et al. (1994)). In the case where the signal lies in the graph structure itself – that is, when $(A_{i,j}, B_{\pi^*(i),\pi^*(j)})_{1 \leq i < j \leq n}$ are correlated pairs of Bernoulli variables – recent work (Cullina and Kiyavash (2017); Cullina et al. (2018a)) showed that there exists a sharp threshold for exact recovery, where the signal-to-noise ratio can be expressed in the correlated Erdős-Rényi model in terms of the size n of both graphs, the marginal edge probability p and the correlation parameter s between edges of the two graphs. Indeed, they established that exact (resp. almost exact) reconstruction is feasible with high probability if and only if $nps \geq \log n + \omega(1)$ (resp. $nps \geq \omega(1)$). When the signal is sparser, e.g. $np = \Theta(1)$, the problem of *partial graph alignment* (that is, recovering only a positive fraction of vertices) has been recently explored algorithmically (Ganassali and Massoulié (2020)) and theoretically (Hall and Massoulié (2020)).

Model of Gaussian Wigner matrices This paper focuses on the case where signal lies in weights on edges between all pairs of nodes. In order to rigorously analyze the fundamental limits of our reconstruction problem, we will work in a probabilistic setting. The *correlated Gaussian Wigner model* is first introduced by Ding et al. (2018) as a standard model for random graph alignment, and

has been further investigated for its own sake in recent work (Fan et al. (2019a); Ganassali et al. (2019); Wu et al. (2020)).

Assume that the weighted adjacency matrices A and B of the two graphs G and G' are symmetric, and sampled as follows: first draw the planted permutation π^* uniformly at random in \mathcal{S}_n . Then all pairs of edge weights $(A_{i,j}, B_{\pi^*(i),\pi^*(j)})_{1 \leq i < j \leq n}$ are i.i.d. couples of normal variables with zero mean, unit variance and correlation parameter $\rho \in [0, 1]$. Since all Gaussian variables are independent from π^* , matrix B can also be drawn from A as follows:

$$B = \rho \cdot \Pi^{*T} A \Pi^* + \sqrt{1 - \rho^2} \cdot H, \quad (3)$$

where H is an independent copy of A , and Π^* is the $n \times n$ matrix representation of permutation π^* , that is $\Pi^*_{i,j} = \mathbf{1}_{j=\pi^*(i)}$.




Figure 1: A sample from model (3) with $n = 5$. For representation, edges are colored according to their weights, and the underlying alignment is $u \mapsto u'$ for $u \in \{1, 2, 3, 4, 5\}$.

Detection problem A most recent paper (Wu et al. (2020)) studies fundamental limits for detection, both in correlated Gaussian weighted and correlated Erdős-Rényi graphs. This time, the problem is as follows: *given A, B , are we able to distinguish between model (3) and a null model, where the two graphs are just independent Gaussian weighted graphs?* Intuitively, this problem is less demanding than that of exact alignment, since the task is to detect – wherever in the graph – the presence of a hidden planted alignment. Under the same model (3), Y. Wu, J. Xu and S. Yu showed that detection is feasible with high probability if $n\rho^2 \geq 4 \log n$, whereas it is impossible if $n\rho^2 \leq (4 - \varepsilon) \log n$ for some $\varepsilon > 0$. Their study builds on an analysis of the likelihood ratio – as often in detection problems. The contribution of this paper is to show that this sharp detection threshold is also that of exact reconstruction. Interestingly, for Gaussian weighted graph alignment, the problem of reconstruction is in fact not more difficult than that of detection.

After this paper was completed, the author was made aware of recent and independent work conducted by Wu et al. (2021), which also obtains – among other things – the results of this paper, albeit with different proof techniques.

Main results

In the sequel, we work with the correlated Gaussian Wigner model described in (3), and establish the precise (sharp) threshold for exact recovery of π^* in this model.

Theorem 1 (Achievability part) *If for n large enough*

$$\rho^2 \geq \frac{(4 + \varepsilon) \log n}{n} \quad (4)$$

for some $\varepsilon > 0$, then there is an estimator (namely, the MAP estimator) $\hat{\pi}$ of π given A, B such that $\hat{\pi} = \pi^$ with probability $1 - o(1)$.*

Theorem 2 (Converse part) *Conversely, if*

$$\rho^2 \leq \frac{4 \log n - \log \log n - \omega(1)}{n} \quad (5)$$

then any estimator $\hat{\pi}$ of π given A, B verifies $\hat{\pi} = \pi^$ with probability $o(1)$.*

Computational limits of exact recovery For the correlated Gaussian Wigner model (3), several algorithms have been studied, usually as a first step in order to analyze further graph alignment algorithms. The state-of-the-art polynomial-time algorithms are either based on *degree profiles* (Ding et al. (2018)), or on a spectral method (Fan et al. (2019a)). In both cases, these methods require the noise parameter $\sqrt{1 - \rho^2}$ to be $O(\log^{-1} n)$. Ganassali et al. (2019) study a simpler algorithm with lower computational complexity ($O(n^2)$ versus $O(n^3)$), requiring $\sqrt{1 - \rho^2}$ to be $O(n^{-7/6})$. In any case, ρ needs to tend to 1, and the regimes in which these methods work well are far from the fundamental limits established in this paper. The present paper thus corroborates the idea that matrix alignment may be computationally hard even in the feasibility regime. In other words, the hard phase can be conjectured to be really wide for this reconstruction problem. Proving a result of that form however remains a very thorny question.

Paper organization

We first define our notations at the beginning of Section 1, and then establish a control on correlations between energies of permutations, using Hanson-Wright inequality. The achievability result is proved in Section 2: after showing that the classical first moment method fails, we take advantage of the correlation structure established before to handle the sharp bound. Then, second moment method is applied in Section 3 to show that lots of small perturbations of the true underlying permutation have lower energies, establishing the converse bound. Finally, some additional proofs are deferred to Appendix A. The proof techniques are not far from those used by Emre Dai et al. (2019), the main novelty being the use of correlation of energies, which is essential to both achievability and impossibility result.

1. Preliminaries

1.1. Definitions and notations

For any positive integer n , let $[n] = \{1, 2, \dots, n\}$. For two positive sequences $\{u_n\}$ and $\{v_n\}$, denote $u_n = O(v_n)$ if there exists $C > 0$ such that $u_n \leq Cv_n$ for all n . We will also write $u_n = o(v_n)$ (resp. $u_n = \omega(v_n)$) if $u_n/v_n \rightarrow 0$ (resp. $v_n/u_n \rightarrow 0$). All limits considered are taken when $n \rightarrow \infty$.

Linear algebra We work with the canonical euclidean norm $\|\cdot\|$ on \mathbb{R}^n , and $\langle \cdot, \cdot \rangle$ the canonical inner product on \mathbb{R}^n or $\mathbb{R}^{n \times n}$. For any $n \times n$ matrix M with real entries, its *Frobenius norm* $\|M\|_F$ and its *operator norm* $\|M\|_{\text{op}}$ are defined as follows:

$$\|M\|_F := \left(\sum_{1 \leq i, j \leq n} A_{i,j}^2 \right)^{1/2} \quad \text{and} \quad \|M\|_{\text{op}} := \sup_{X \in \mathbb{R}^n \setminus \{0\}} \frac{\|MX\|}{\|X\|}.$$

Note that for any normal matrix (that is, if $M^T M = M M^T$), then $\|M\|_{\text{op}}$ equals $\rho(M)$, the spectral radius of M .

Probability When working with model (3), we will denote by \mathbb{P}_A (resp. \mathbb{E}_A) the conditional probability (resp. the conditional expectation) with respect to the random matrix A . Throughout the paper, $\mathcal{N}(\mu, v)$ denotes a Gaussian variable (resp. vector) with mean μ and variance (resp. covariance matrix) v . Such a Gaussian variable (resp. vector) is called *standard* if $\mu = 0$ and $v = 1$ (resp. v is the identity matrix). We say that an event \mathcal{A}_n happens *with high probability (w.h.p)* if $\mathbb{P}(\mathcal{A}_n) \rightarrow 1$ when $n \rightarrow \infty$.

Permutations We denote by \mathcal{S}_m the set of permutations of $[m]$. To any permutation $\sigma \in \mathcal{S}_m$, we can associate its $m \times m$ matrix representation Σ defined by $\Sigma_{i,j} = \mathbf{1}_{j=\sigma(i)}$. Define \mathcal{F}_σ the set of *fixed points* of σ :

$$\mathcal{F}_\sigma := \{i \in [m], \sigma(i) = i\}, \quad (6)$$

and denote $f_\sigma := \#\mathcal{F}_\sigma$. Similarly, we define the set of *unfixed points* of σ :

$$\mathcal{D}_\sigma := [m] \setminus \mathcal{F}_\sigma = \{i \in [m], \sigma(i) \neq i\}, \quad (7)$$

and we denote $d_\sigma := \#\mathcal{D}_\sigma$. For any $d \in \{0, \dots, m\}$ we define $\mathcal{S}_{m,d}$ the set of permutations of \mathcal{S}_m with exactly d unfixed points. Note that $\#\mathcal{S}_{m,1} = 0$ and that we have the inequality

$$\#\mathcal{S}_{m,d} = \binom{m}{m-d} \#\{\sigma \in \mathcal{S}_d, F_\sigma = 0\} \leq \binom{m}{m-d} d! \leq m^d. \quad (8)$$

Similarity between two permutations $\sigma, \sigma' \in \mathcal{S}_n$ is measured by their *overlap*:

$$\text{ov}(\sigma, \sigma') := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\sigma(i)=\sigma'(i)} = \frac{1}{n} f_{\sigma^{-1} \circ \sigma'}.$$

Observe that on a graph of size n , each permutation σ of the vertices $[n]$ has a natural extension to a canonical permutation on edges $\sigma^{\text{E}} : \binom{[n]}{2} \rightarrow \binom{[n]}{2}$ defined as follows:

$$\sigma^{\text{E}} : e = \{i, j\} \mapsto \sigma^{\text{E}}(e) = \{\sigma(i), \sigma(j)\}.$$

Note that the mapping $\sigma \mapsto \sigma^{\text{E}}$ is one-to-one as soon as $n \geq 3$, since for all $i \in [n]$ and $j \neq j' \in [n] \setminus \{i\}$, edges $\sigma^{\text{E}}(\{i, j\})$ and $\sigma^{\text{E}}(\{i, j'\})$ have only one node in common, which is $\sigma(i)$. We will use the notation $\mathcal{F}_\sigma^{\text{E}} = \mathcal{F}_{\sigma^{\text{E}}}$ (resp. $\mathcal{D}_\sigma^{\text{E}} = \mathcal{D}_{\sigma^{\text{E}}}$) the set of *fixed edges* (resp. *unfixed edges*) of σ . Similarly we denote $f_\sigma^{\text{E}} = f_{\sigma^{\text{E}}}$ and $d_\sigma^{\text{E}} := d_{\sigma^{\text{E}}}$, for brevity.

Note that d_σ^{E} and d_σ are closely tied, since for all $\sigma \in \mathcal{S}_n$, we have the inequality

$$d_\sigma \left(n - \frac{d_\sigma}{2} \right) \leq d_\sigma^{\text{E}} \leq d_\sigma \left(n - \frac{d_\sigma - 1}{2} \right). \quad (9)$$

Indeed, observe that

- (i) the number of fixed edges is at least the number of pairs of fixed points, and
- (ii) the number of fixed edges is exactly the number of pairs of fixed points plus the number of pairs $(i, j), i < j$ that are exchanged by σ (that is, the number of transpositions), this number being at most $d_\sigma/2$.

These remarks give that

$$\binom{n - d_\sigma}{2} \leq \binom{n}{2} - d_\sigma^E \leq \binom{n - d_\sigma}{2} + \frac{d_\sigma}{2},$$

which directly implies (9).

Remark 1.1 Note that inequality (9) gives the almost sure equivalents $d_\sigma^E \sim d_\sigma n$ when $d_\sigma = o(n)$, and $d_\sigma^E \sim \frac{1}{2}\alpha(2 - \alpha)n^2$ when $d_\sigma = \alpha n$. In any case, $d_\sigma^E \in [\frac{1}{2}d_\sigma n, d_\sigma n]$.

1.2. MAP estimation, relative energy of permutations

Since π^* is uniformly chosen, we work in a Bayesian setting: let us evaluate the posterior probability density of π^* given A, B :

$$\begin{aligned} p_{\pi^*|A,B}(\pi|a, b) &\propto p_{\pi^*,A,B}(\pi, a, b) \\ &\propto \exp\left(-\frac{1}{2(1-\rho^2)} \sum_{1 \leq i < j \leq n} (B_{\pi(i),\pi(j)} - \rho A_{i,j})^2\right), \end{aligned}$$

where \propto indicates equality up to some factors that do not depend on σ . Define the *loss function*

$$\mathcal{L}(\pi, A, B) := \sum_{1 \leq i < j \leq n} (B_{\pi(i),\pi(j)} - \rho A_{i,j})^2. \quad (10)$$

This loss function can also be viewed as the *energy* associated with permutation π . Note that the posterior distribution is a Gibbs measure corresponding to this energy \mathcal{L} , with inverse temperature $\beta = \frac{1}{2(1-\rho^2)}$. The MAP (maximum a posteriori) estimator is thus

$$\hat{\pi}_{\text{MAP}} := \arg \max_{\pi} p_{\pi^*|A,B}(\pi|A, B) = \arg \min_{\pi} \mathcal{L}(\pi, A, B), \quad (11)$$

where the minimum is taken over all permutations $\pi \in \mathcal{S}_n$. The above formulation (11) is standard in the literature of graph and matrix alignment and meets the classical QAP formulation (2), since

$$\arg \min_{\pi} \mathcal{L}(\pi, A, B) = \arg \max_{\Pi} \langle A, \Pi B \Pi^T \rangle.$$

Theory from Bayesian optimal estimation guarantees that the best possible estimator for our exact reconstruction problem, in the Bayes risk sense, is $\hat{\pi}_{\text{MAP}}$. Thus, if MAP estimator fails with high probability, then no estimator can succeed. This is why this estimator is often studied in exact reconstruction problems, as already done in previous works (Cullina and Kiyavash (2017); Cullina et al. (2018b); Emre Dai et al. (2019)).

From now on we work conditionally on π^* which can always be assumed to be id without loss of generality. More precisely, we will make the variable change $\sigma = \pi^* \circ \pi^{-1}$; writing B as a function of σ , A and H , (10) becomes

$$\begin{aligned} \mathcal{L}(\sigma, A, H) &= \rho^2 \sum_{1 \leq i < j \leq n} (A_{i,j} - A_{\sigma(i),\sigma(j)})^2 - 2\rho\sqrt{1-\rho^2} \sum_{1 \leq i < j \leq n} H_{i,j} (A_{i,j} - A_{\sigma(i),\sigma(j)}) \\ &\quad + (1 - \rho^2) \sum_{1 \leq i < j \leq n} H_{i,j}^2. \end{aligned}$$

The loss function \mathcal{L} applied to the ground truth $\pi = \pi^*$ – that is $\sigma = \text{id}$ – gives the energy reference $(1 - \rho^2) \sum_{1 \leq i < j \leq n} H_{i,j}^2$. In order to compare any π with π^* – or any σ with id – we further define the *relative energy* of a permutation $\sigma \in \mathcal{S}_n$:

$$\begin{aligned} \delta(\sigma) &:= \mathcal{L}(\sigma, A, H) - \mathcal{L}(\text{id}, A, H) \\ &= \rho^2 \sum_{1 \leq i < j \leq n} (A_{i,j} - A_{\sigma(i),\sigma(j)})^2 - 2\rho\sqrt{1-\rho^2} \sum_{1 \leq i < j \leq n} H_{i,j} (A_{i,j} - A_{\sigma(i),\sigma(j)}). \end{aligned} \quad (12)$$

We next omit in our notations the dependency on A and H of $\delta(\sigma)$.

Remark 1.2 *This relative energy δ , also introduced by Cullina and Kiyavash (2017) for Erdős-Rényi graph alignment, is a measurement of the quality of a proposed alignment: $\delta(\sigma) \leq 0$ means that $\sigma^{-1} \circ \pi^*$ is a better alignment than π^* for A and B in the posterior sense. A crucial set is then*

$$\mathcal{Q} := \{\sigma \in \mathcal{S}_n, \delta(\sigma) \leq 0\}.$$

Points of \mathcal{Q} are alignments on which the posterior distribution puts important weights – at least greater weights than that of the ground truth – or equivalently points of low energy. Note that $\text{id} \in \mathcal{Q}$.

In view of (12), conditionally on A , $\delta(\sigma)$ is as follows:

$$\delta(\sigma) = \rho^2 v_\sigma - 2\rho\sqrt{1-\rho^2} X_\sigma, \quad (13)$$

where

$$v_\sigma := \sum_{1 \leq i < j \leq n} (A_{i,j} - A_{\sigma(i),\sigma(j)})^2,$$

and $X = (X_\sigma)_{\sigma \in \mathcal{S}_n}$ is a Gaussian vector, centered, with covariance given by

$$\text{Cov}(X_\sigma, X_{\sigma'}) = \sum_{1 \leq i < j \leq n} (A_{i,j} - A_{\sigma(i),\sigma(j)}) (A_{i,j} - A_{\sigma'(i),\sigma'(j)}) := c_{\sigma,\sigma'}.$$

Note that for all $\sigma \in \mathcal{S}_n$, $c_{\sigma,\sigma} = v_\sigma$. Elaborating on the correlation structure of these relative energies is the object of the end of this section.

1.3. Control of covariance structure of relative energies

For all $\sigma, \sigma' \in \mathcal{S}_n$, $c_{\sigma, \sigma'}$ can be written as follows

$$c_{\sigma, \sigma'} = \sum_{e \in \binom{[n]}{2}} (A_e - A_{\sigma^E(e)}) (A_e - A_{\sigma'^E(e)})$$

and satisfies

$$\mathbb{E} [c_{\sigma, \sigma'}] = \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E + \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'}^E.$$

In particular,

$$\mathbb{E} [v_\sigma] = d_\sigma^E + d_\sigma^E = 2d_\sigma^E.$$

Random variables $c_{\sigma, \sigma'}$ only depend on the entries of A , which are Gaussian. Moreover, $c_{\sigma, \sigma'}$ being a quadratic form evaluated on a Gaussian vector, it can therefore be controlled using Hanson-Wright inequality:

Lemma 1.1 (Hanson-Wright inequality (Hanson and Wright (1971))) *Let X be a standard Gaussian vector, and M a deterministic matrix. Then there exists a universal constant $c > 0$ such that with probability at least $1 - 2\delta$:*

$$|X^T M X - \text{Tr} M| \leq c \left(\|M\|_F \sqrt{\log(1/\delta)} + \|M\|_{\text{op}} \log(1/\delta) \right). \quad (14)$$

We refer to [Hanson and Wright \(1971\)](#) for a proof. Inequality (14) used in our context leads to the following

Corollary 1.1 *There exists a universal constant $C > 0$ such that with high probability, for every $d \in \{2, \dots, n\}$, for all $\sigma, \sigma' \in \mathcal{S}_{n,d}$,*

$$|c_{\sigma, \sigma'} - \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E - \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'}^E| \leq Cd\sqrt{n \log n}.$$

Proof We first make the following observation: for any $\sigma, \sigma' \in \mathcal{S}_n$,

$$\begin{aligned} c_{\sigma, \sigma'} &= \sum_e (A_e - A_{\sigma^E(e)}) (A_e - A_{\sigma'^E(e)}) \\ &= A^T (I_N - \Sigma)^T (I_N - \Sigma') A, \end{aligned}$$

where $A = (A_e)_e$ is viewed as a standard Gaussian vector of size $N = \binom{n}{2}$, and Σ (resp. Σ') is the $N \times N$ permutation matrix associated with σ^E (resp. σ'^E). Note that

$$\begin{aligned} \text{Tr}((I_N - \Sigma)^T (I_N - \Sigma')) &= N - f_\sigma^E - f_{\sigma'}^E + f_{\sigma^{-1}\circ\sigma'}^E \\ &\stackrel{(a)}{=} \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E + \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'}^E, \end{aligned}$$

where (a) is obtained by noticing that

$$\#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E + \#\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'}^E = d_\sigma^E + d_{\sigma'}^E - \#\mathcal{D}_\sigma^E \cup \mathcal{D}_{\sigma'}^E + f_{\sigma^{-1}\circ\sigma'}^E - \#\mathcal{F}_\sigma^E \cup \mathcal{F}_{\sigma'}^E$$

and that $\#(\mathcal{D}_\sigma^E \cup \mathcal{D}_{\sigma'}^E) + \#(\mathcal{F}_\sigma^E \cup \mathcal{F}_{\sigma'}^E) = N$. For a fixed d and $\sigma, \sigma' \in \mathcal{S}_{n,d}$, one has

$$\begin{aligned} \|(I_N - \Sigma)^T(I_N - \Sigma')\|_F &\leq \|(I_N - \Sigma')\|_F + \|\Sigma^T(I_N - \Sigma')\|_F \\ &= 2\|(I_N - \Sigma')\|_F \\ &\leq 2\sqrt{2d_{\sigma'}^E} \\ &\leq 2\sqrt{2dn}, \end{aligned}$$

where we used (9) in the last step. One also has

$$\begin{aligned} \|(I_N - \Sigma)^T(I_N - \Sigma')\|_{\text{op}} &\leq \rho(I_N - \Sigma) \times \rho(I_N - \Sigma') \\ &\leq 2 \times 2 = 4. \end{aligned}$$

Taking $\delta = n^{-(2d+2)}$, Lemma 1.1 gives that with probability at least $1 - 2\delta$,

$$\begin{aligned} |c_{\sigma, \sigma'} - \#(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E) - \#(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'})| &\leq c \left(2\sqrt{2}\sqrt{d(2d+2)}\sqrt{n \log n} + 4(2d+2) \log n \right) \\ &\leq Cd\sqrt{n \log n}, \end{aligned} \quad (15)$$

for some universal constant $C > 0$. The proof is concluded by checking that this inequality holds w.h.p. for all d and $\sigma, \sigma' \in \mathcal{S}_{n,d}$: the probability that at least one pair (σ, σ') contradicts (15) is upper bounded by

$$n \times (\#\mathcal{S}_{n,d})^2 \times 2\delta \leq 2n^{1+2d-2d-2} = o(1).$$

■

In the rest of the paper we define the event

$$\mathcal{A} := \left\{ \forall d \in [n], \forall \sigma, \sigma' \in \mathcal{S}_{n,d}, |c_{\sigma, \sigma'} - \#(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E) - \#(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E \cap \mathcal{F}_{\sigma^{-1}\circ\sigma'})| \leq Cd\sqrt{n \log n} \right\}, \quad (16)$$

which happens with probability $1 - o(1)$ by Corollary 1.1.

2. Achievability result

In this section, we establish Theorem 1.

2.1. Failure of first moment method

For the achievability result, the first strategy is to use the union bound (or first moment method) to show that under condition (4) of Theorem 1,

$$\mathbb{P}(\text{MAP fails}) = \mathbb{P}(\hat{\pi}_{\text{MAP}} \neq \pi) = o(1).$$

As described hereafter, this naive method does not give the correct bound. Indeed, let us evaluate $\mathbb{P}(\delta(\sigma) \leq 0)$ for a given $\sigma \neq \text{id}$. In view of the conditional distribution (13) of $\delta(\sigma)$ we have

$$\begin{aligned} \mathbb{P}(\delta(\sigma) \leq 0) &= \mathbb{E} \left[\mathbb{E}_A [\mathbf{1}_{\delta(\sigma) \leq 0}] \right] = \mathbb{E} \left[\mathbb{P}_A \left(\rho^2 v_\sigma - 2\rho\sqrt{1-\rho^2} X_\sigma \leq 0 \right) \right] \\ &= \mathbb{E} \left[\mathbb{P}_A \left(\rho^2 v_\sigma - 2\rho\sqrt{1-\rho^2} \sqrt{v_\sigma} \cdot \mathcal{N}(0, 1) \leq 0 \right) \right] \\ &= \mathbb{E} \left[\mathbb{P}_A \left(\mathcal{N}(0, 1) \geq \frac{\rho\sqrt{v_\sigma}}{2\sqrt{1-\rho^2}} \right) \right] \leq \mathbb{E} \left[\exp \left(-\frac{\rho^2}{8(1-\rho^2)} v_\sigma \right) \right], \end{aligned}$$

where we used standard Gaussian concentration in the last inequality: $\mathbb{P}(\mathcal{N}(0, 1) \geq t) \leq \exp(-t^2/2)$. Note that on event \mathcal{A} defined in (16) and inequality (9),

$$\forall d \in [n], \forall \sigma \in \mathcal{S}_{n,d}, v_\sigma \geq 2d_\sigma^E - Cd_\sigma \sqrt{n \log n} \geq d_\sigma^E (2 - 2\varepsilon_n),$$

setting $\varepsilon_n = 2C\sqrt{\log n/n}$. Union bound then gives

$$\begin{aligned} \mathbb{P}(\text{MAP fails}) &\leq \mathbb{P}(\exists \sigma \in \mathcal{S}_n \setminus \{\text{id}\}, \delta(\sigma) \leq 0) \\ &\leq o(1) + \sum_{\sigma \in \mathcal{S}_n \setminus \{\text{id}\}} \mathbb{E} \left[\exp \left(-\frac{\rho^2}{8(1-\rho^2)} v_\sigma \right) \mathbf{1}_{\mathcal{A}} \right] \\ &\leq o(1) + \sum_{\sigma \in \mathcal{S}_n \setminus \{\text{id}\}} \exp \left(-\frac{\rho^2}{8(1-\rho^2)} (2 - 2\varepsilon_n) d_\sigma^E \right) \\ &\leq o(1) + \sum_{\sigma \in \mathcal{S}_n \setminus \{\text{id}\}} \exp \left(-\frac{\rho^2}{4} (1 - \varepsilon_n) d_\sigma^E \right), \end{aligned}$$

where we used $1/(1-\rho^2) > 1$ in the last step. Let us now study the last sum, distinguishing the terms according to $d := d_\sigma$:

- As long as $d = o(n)$, by Remark 1.1, the terms behave like $\exp\left(-\frac{\rho^2}{4}(1-\varepsilon_n)dn\right)$. By (8), $\log \#\mathcal{S}_{n,d} \leq d \log n$ so the partial sum is small if $\frac{\rho^2}{4}(1-\varepsilon_n)n - \log n > 0$, which gives the necessary condition $\rho^2 \geq 4\frac{\log n}{n}$.
- However, the situation is different when it comes to large values of d . For instance, let us study the contribution of *derangements* to the sum (that is, σ such that $d_\sigma = n$). Note that these derangements are very numerous (their number is $\sim e^{-1}n!$). Again by Remark 1.1, their contribution is thus of order

$$e^{-1}n! \exp\left(\rho^2(1-\varepsilon_n)n^2/8(1-o(1))\right) = \exp\left((n \log n - \rho^2 n^2/8)(1-o(1))\right),$$

which gives a more restrictive condition: $\rho^2 \geq 8\frac{\log n}{n}$.

As seen here-above, this naive first moment method enables to ensure feasibility of exact reconstruction only in the regime where $\rho^2 \geq 8\frac{\log n}{n}$, which is not the optimal one. This bound is actually quite rough here, because the variables are substantially correlated when d gets large and their contributions make the first moment explode. The next section takes advantage of these correlations in order to get access to the sharp bound.

2.2. Improving the first moment method with correlations.

For all $d \in \{2, \dots, n\}$, define \mathcal{E}_d the event:

$$\mathcal{E}_d := \{\exists \sigma \in \mathcal{S}_{n,d}, \delta(\sigma) \leq 0\}.$$

In this Section we will assume that

$$\rho \geq (2 + \varepsilon) \sqrt{\frac{\log n}{n}},$$

for some $\varepsilon > 0$. Recall that we work on the event \mathcal{A} defined in (16), and that conditionally on entries of matrix A , we can write

$$\delta(\sigma) = \rho^2 v_\sigma - 2\rho\sqrt{1-\rho^2}X_\sigma, \quad (17)$$

where $X = (X_\sigma)_{\sigma \in \mathcal{S}_{n,d}}$ is a Gaussian vector, centered, with covariance given by $\text{Cov}(X_\sigma, X_{\sigma'}) = c_{\sigma,\sigma'}$. Also note that on event \mathcal{A} , for all $d \leq \alpha n$ and $\sigma \in \mathcal{S}_{n,d}$, inequality (9) gives

$$v_\sigma = (1 - o(1))2dn(1 - \alpha/2). \quad (18)$$

In view of (18), as previously done in Section 2.1, naive first moment method may suffice for $d \leq \alpha n$:

$$\begin{aligned} \mathbb{P}\left(\bigcup_{2 \leq d \leq \alpha n} \mathcal{E}_d\right) &\leq o(1) + \sum_{d=2}^{\alpha n} \#\mathcal{S}_{n,d} \times \mathbb{P}\left(\mathcal{N}(0,1) \geq \frac{\rho\sqrt{v_\sigma}}{2\sqrt{1-\rho^2}} \cap \mathcal{A}\right) \\ &\leq o(1) + \sum_{d=2}^{\alpha n} \#\mathcal{S}_{n,d} \times \mathbb{P}\left(\mathcal{N}(0,1) \geq (1 + \varepsilon/2)\sqrt{2d \log n(1 - \alpha/2)}(1 - o(1))\right) \\ &\leq o(1) + \sum_{d=2}^{\alpha n} \exp(d \log n - d \log n(1 + \varepsilon)(1 - \alpha/2) + o(d \log n)), \end{aligned}$$

which is $o(1)$ as soon as $\alpha < \alpha_0 := \frac{2\varepsilon}{1-\varepsilon/2}$. It then remains to control the probabilities $\mathbb{P}(\mathcal{E}_d)$ for $d \geq \alpha_0 n$. As mentioned earlier, we take advantage of the correlation structure in (17). More precisely, we show that all variables X_σ at a given level $d = \alpha n$ have substantial positive covariance when compared to their variance – of order $\alpha(2 - \alpha)n^2$ on \mathcal{A} by (18) – as shown in Figure 2. To do so, we derive an appropriate lower bound for $c_{\sigma,\sigma'}$ for $\sigma, \sigma' \in \mathcal{S}_{n,\alpha n}$. This is the scope of the following Lemma:

Lemma 2.1 *With high probability, there exists a universal constant $C_1 > 0$ such that for any $d = \alpha n$ with fixed $\alpha > 0$ and $\sigma, \sigma' \in \mathcal{S}_{n,\alpha n}$:*

$$\text{Cov}(X_\sigma, X_{\sigma'}) = c_{\sigma,\sigma'} \geq f(\alpha)n^2 - C_1 n^{3/2} \log^{1/2} n,$$

with

$$f(\alpha) := \begin{cases} \alpha^2 & \text{if } \alpha < 1/2 \\ \alpha^2 - \frac{1}{2}(2\alpha - 1)^2 & \text{if } \alpha \geq 1/2 \end{cases} \quad (19)$$

Thus for any $\varepsilon' > 0$, with high probability, for any $d = \alpha n$ with fixed $\alpha > 0$,

$$\max_{\sigma \in \mathcal{S}_{n,\alpha n}} X_\sigma \leq \sqrt{2\alpha(\alpha(2-\alpha) - f(\alpha))} n^{3/2} \log^{1/2} n + (2 + \varepsilon')n \log^{1/2} n.$$

The proof of this Lemma is obtained by working on event \mathcal{A} defined in (16), and establishing a lower bound on $\#(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E)$, which is simply the number of edges that are deranged both by σ^E and σ'^E . It can be found in Appendix A.1.




Figure 2: Plot on $[0, 1]$ of normalized variance $\alpha(2 - \alpha)$, together with the lower bound on the normalized covariance (function f) defined by (19).

Then, since $f(\alpha) \leq \alpha(2 - \alpha)$ with elementary computations, according to Lemma 2.1, there is an event \mathcal{B} of probability $1 - o(1)$ such that

$$\max_{\sigma \in \mathcal{S}_{n,d}} X_\sigma \leq (1 + o(1)) \sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))} n^{3/2} \log^{1/2} n$$

holds for all $d = \alpha n$ with $\alpha > \alpha_0$. Note that on event $\mathcal{A} \cap \mathcal{B}$, for all $d = \alpha n$ and $\sigma \in \mathcal{S}_{n,d}$,

$$\begin{aligned} \rho^{-1} \delta(\sigma) &\geq \rho v_\sigma - 2\sqrt{1 - \rho^2} \max_{\sigma \in \mathcal{S}_{n,d}} X_\sigma \\ &\geq (1 + o(1)) n^{3/2} \log^{1/2} n \left[(2 + \varepsilon)\alpha(2 - \alpha) - 2\sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))} \right] \\ &\geq (1 + o(1)) \times 2 \times \left[\alpha(2 - \alpha) - \sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))} \right] n^{3/2} \log^{1/2} n \geq 0, \end{aligned}$$

for n large enough, since it can be easily checked (see Appendix A.3) that

Lemma 2.2 For every $\alpha \in [0, 1]$,

$$\alpha(2 - \alpha) - \sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))} \geq 0. \quad (20)$$

Previous computations give that

$$\mathbb{P} \left(\bigcup_{d \geq \alpha n} \mathcal{E}_d \right) \leq 1 - \mathbb{P}(\mathcal{A} \cap \mathcal{B}) = o(1),$$

and ends the proof of Theorem 1.

3. Converse bound: second moment method for transpositions

In this section, we prove Theorem 2. As already stated in the introduction, theory from Bayesian optimal estimation guarantees that the best possible estimator for our exact reconstruction problem, in the Bayes risk sense, is $\hat{\pi}_{\text{MAP}}$. We will show that under assumption (5) of Theorem 2, this MAP estimator fails with high probability, which implies that no estimator can succeed.

This converse bound is obtained by a second moment argument, showing that with high probability, there are lots of permutation $\tau \neq \text{id}$ – in fact, transpositions – such that $\delta(\tau)$ is negative, that is, $\tau^{-1} \circ \pi^*$ is a substantially better alignment than π^* , with lowest energy. Let us denote $\mathcal{T}_n \subset \mathcal{S}_n$ the set of all permutations of $[n]$ that are transpositions. For all $\tau \in \mathcal{T}_n$, we have $d_\tau^{\text{E}} = 2(n-2)$. Corollary 1.1 gives that the event

$$\mathcal{C} := \left\{ \forall \tau, \tau' \in \mathcal{T}_n, |c_{\tau, \tau'} - \#(\mathcal{D}_\tau^{\text{E}} \cap \mathcal{D}_{\tau'}^{\text{E}}) - \#(\mathcal{D}_\tau^{\text{E}} \cap \mathcal{D}_{\tau'}^{\text{E}} \cap \mathcal{F}_{\tau \circ \tau'}^{\text{E}})| \leq C \sqrt{n \log n} \right\}$$

happens with probability $1 - o(1)$ for $C > 0$ large enough. In particular, on \mathcal{C} , for $C > 0$ large enough,

$$\forall \tau \in \mathcal{T}_n, |v_\tau - 4n| \leq C \sqrt{n \log n}.$$

In this section we are working under the assumption (5) that we recall here:

$$\rho^2 \leq \frac{4 \log n - \log \log n - \omega(1)}{n}$$

We are about to show the following: under condition (5), with high probability,

$$\# \{ \tau \in \mathcal{T}_n, \delta(\tau) < 0 \} = \omega(1). \quad (21)$$

To do so, we use the classical Paley-Zygmund inequality, proven in Appendix A.4 for self-containment:

Lemma 3.1 (Paley-Zygmund inequality) *Let Y be a real random variable with positive mean and finite variance. Then for all $0 \leq c \leq 1$,*

$$\mathbb{P}(Y \geq c \mathbb{E}[Y]) \geq (1-c)^2 \frac{\mathbb{E}[Y]^2}{\mathbb{E}[Y^2]}.$$

Thus, in the case where $\mathbb{E}[Y^2] \sim \mathbb{E}[Y]^2$, taking $c \rightarrow 0$ implies that $Y \geq o(\mathbb{E}[Y])$ with high probability.

Define

$$X := \sum_{\tau \in \mathcal{T}_n} \mathbf{1}_{\delta(\tau) < 0}. \quad (22)$$

Using a standard coupling argument in (22), one can see that X is decreasing with ρ , thus we can assume without loss of generality that

$$\rho^2 = \frac{4 \log n - \log \log n - a_n}{n}, \quad (23)$$

with a sequence $(a_n)_n$ such that $a_n = \omega(1)$ and $a_n = o(\log \log n)$. We compute the first moment of X , in view of the conditional distribution of $\delta(\tau)$ given in (13):

$$\begin{aligned} \mathbb{E}[X] &\geq \mathbb{E}[X\mathbf{1}_{\mathcal{C}}] = \frac{n(n-1)}{2} \mathbb{E} \left[\mathbb{P}_A \left(\mathcal{N}(0,1) \geq \frac{\rho\sqrt{v_\tau}}{2\sqrt{1-\rho^2}} \cap \mathcal{C} \right) \right] \\ &\geq \frac{n(n-1)}{2} \mathbb{E} \left[(1-o(1)) \mathbb{P}_A \left(\mathcal{N}(0,1) \geq \frac{1}{2} \sqrt{4 \log n - \log \log n - a_n} \sqrt{4 - Cn^{-1/2} \log^{1/2} n} \right) \right] \\ &= \frac{n(n-1)}{2} \mathbb{E} \left[(1-o(1)) \mathbb{P}_A \left(\mathcal{N}(0,1) \geq \sqrt{4 \log n - \log \log n - a_n} - o(1) \right) \right] \\ &\sim \frac{n^2}{4\sqrt{2\pi}\sqrt{\log n}} \exp \left(-2 \log n + \frac{\log \log n}{2} + \frac{a_n}{2} \right) \\ &= \frac{1}{4\sqrt{2\pi}} \exp \left(\frac{a_n}{2} \right) \rightarrow \infty. \end{aligned}$$

Note that (23) is thus precisely the condition ensuring that $\mathbb{E}[X\mathbf{1}_{\mathcal{C}}] \rightarrow \infty$. The second moment argument computation being a little more technical, we encapsulate it into the following Lemma:

Lemma 3.2 (Second moment computation of $X\mathbf{1}_{\mathcal{C}}$) *Let $Y := X\mathbf{1}_{\mathcal{C}}$. Under assumption (23),*

$$\mathbb{E}[Y^2] \leq (1+o(1))\mathbb{E}[Y]^2.$$

Proof [Proof of Lemma 3.2]

We represent a transposition τ by its only 2-cycle $(i j)$ with $i < j$. We then distinguish two cases in couples $\tau = (i j) \neq \tau' = (k \ell) \in \mathcal{T}_n$:

- We write $\tau \cap \tau' = \emptyset$ when τ and τ' have no common point in their 2-cycle: $i \neq k$ and $j \neq \ell$. When $\tau \in \mathcal{T}_n$ is fixed, note that

$$\#\{\tau' \in \mathcal{T}_n, \tau \cap \tau' = \emptyset\} = \frac{(n-2)(n-3)}{2}.$$

- We write $\tau \cap \tau' \neq \emptyset$ when τ and τ' are different but share one common point: for instance $\tau = (3 5)$ and $\tau' = (5 11)$ verify $\tau \cap \tau' \neq \emptyset$. When $\tau \in \mathcal{T}_n$ is fixed, note that

$$\#\{\tau' \in \mathcal{T}_n, \tau \cap \tau' \neq \emptyset\} = 2(n-2).$$

Note that

$$\mathbb{E}[Y^2] = \mathbb{E}[Y] + \sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' = \emptyset} \mathbb{P}(\delta(\tau) < 0, \delta(\tau') < 0, \mathcal{C}) + \sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' \neq \emptyset} \mathbb{P}(\delta(\tau) < 0, \delta(\tau') < 0, \mathcal{C}).$$

We now evaluate these two sums. For this, we will need the following Lemma, which proof is deferred to Appendix A.5.

Lemma 3.3 (Control of deviation probabilities for correlated Gaussians) *Let Z_1, Z_2 be two Gaussian variables with mean 0, variance 1 and correlation $\alpha_n \in [0, 1]$. For any t_n such that $t_n \rightarrow \infty$,*

- (i) *If $\alpha_n t_n \rightarrow 0$, then for n large enough*

$$\mathbb{P}(Z_1 > t_n, Z_2 > t_n) \leq e^{-2t_n^2} + (1+o(1))\mathbb{P}(Z_1 > t_n)\mathbb{P}(Z_2 > t_n). \quad (24)$$

- (ii) *More generally,*

$$\mathbb{P}(Z_1 > t_n, Z_2 > t_n) \leq (1+o(1)) \frac{1+\alpha_n}{\sqrt{2\pi} t_n} \exp\left(-\frac{t_n^2}{1+\alpha_n}\right). \quad (25)$$

First case: $\tau \cap \tau' = \emptyset$. Without loss of generality we can assume that $\tau = (1\ 2)$ and $\tau' = (3\ 4)$. The following diagram shows the simple action of τ and τ' on an interesting (overlapping) subset of edges.

$$\begin{array}{ccc} \{1, 3\} & \xleftrightarrow{\tau} & \{2, 3\} \\ \tau' \updownarrow & & \updownarrow \tau' \\ \{1, 4\} & \xleftrightarrow{\tau} & \{2, 4\} \end{array}$$

We then see that $\sharp(\mathcal{D}_\tau^E \cap \mathcal{D}_{\tau'}^E) + \sharp(\mathcal{D}_\tau^E \cap \mathcal{D}_{\tau'}^E \cap \mathcal{F}_{\tau \circ \tau'}^E) = 4 + 0 = 4$. So, denoting $\alpha_{\tau, \tau'} := \frac{c_{\tau, \tau'}}{\sqrt{v_\tau v_{\tau'}}}$, on \mathcal{C} ,

$$|\alpha_{\tau, \tau'}| \leq \frac{C\sqrt{n \log n} + 4}{4n - C\sqrt{n \log n}} = O\left(\sqrt{\frac{\log n}{n}}\right).$$

In view of the conditional distribution of $\delta(\tau)$ given in (13):

$$\sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' = \emptyset} \mathbb{P}(\delta(\tau) < 0, \delta(\tau') < 0, \mathcal{C}) = (1 - o(1)) \sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' = \emptyset} \mathbb{P}(Z_\tau > t_n, Z_{\tau'} > t_n), \quad (26)$$

with $t_n = \sqrt{4 \log n - \log \log n - a_n}$, where $Z_\tau, Z_{\tau'}$ are two Gaussian variables of mean 0, with correlation coefficient α_n of order $O(\log^{1/2} n^{-1/2})$. Since $\alpha_n t_n \rightarrow 1$, by lemma 3.3 case (i), the sum in (26) is upper bounded by

$$\begin{aligned} & (1 - o(1)) \frac{n(n-1)}{2} \times \frac{(n-2)(n-3)}{2} \times \left[C e^{-2t_n^2} + (1 - o(1)) \mathbb{P}(Z_1 > t_n) \mathbb{P}(Z_2 > t_n) \right] \\ & \leq (1 + o(1)) \mathbb{E}[Y]^2. \end{aligned}$$

Second case: $\tau \cap \tau' \neq \emptyset$. Without loss of generality we can assume that $\tau = (1\ 2)$ and $\tau' = (2\ 3)$. We can immediately deduce that $\sharp(\mathcal{D}_\tau^E \cap \mathcal{D}_{\tau'}^E) + \sharp(\mathcal{D}_\tau^E \cap \mathcal{D}_{\tau'}^E \cap \mathcal{F}_{\tau \circ \tau'}^E) = (n-2) + 0 = n-2$. So, denoting $\alpha_{\tau, \tau'} := \frac{c_{\tau, \tau'}}{\sqrt{v_\tau v_{\tau'}}}$, on \mathcal{C} ,

$$|\alpha_{\tau, \tau'}| \leq \frac{C\sqrt{n \log n} + n - 2}{4n - C\sqrt{n \log n}} \sim \frac{1}{4}.$$

Again, in view of the conditional distribution of $\delta(\tau)$ given in (13):

$$\sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' \neq \emptyset} \mathbb{P}(\delta(\tau) < 0, \delta(\tau') < 0, \mathcal{C}) = (1 - o(1)) \sum_{\tau \in \mathcal{T}_n} \sum_{\tau', \tau \cap \tau' \neq \emptyset} \mathbb{P}(Z_\tau > t_n, Z_{\tau'} > t_n), \quad (27)$$

with $t_n = \sqrt{4 \log n - \log \log n - a_n}$, where $Z_\tau, Z_{\tau'}$ are two Gaussian variables of mean 0, with correlation coefficient $\alpha_n \sim 1/4$. By Lemma 3.3 case (ii), the sum in (27) is upper bounded by

$$\begin{aligned} & (1 - o(1)) \frac{n(n-1)}{2} \times 2(n-2) \times \left[(1 + o(1)) \frac{1 + \alpha_n}{\sqrt{2\pi} t_n} \exp\left(-\frac{t_n^2}{1 + \alpha_n}\right) \right] \\ & \leq C'' n^3 \log^{-1/2}(n) \exp\left(-\frac{16}{5} \log n + o(\log n)\right) = o(1) = o(\mathbb{E}[Y]^2). \end{aligned}$$



Lemma 3.2 together with Payley-Zigmond inequality (Lemma 3.1) implies that $Y \geq o(\mathbb{E}[Y])$ with high probability and thus proves (21) and the converse result of Theorem 2.

Remark 3.1 *We have shown here that under condition (5), there is with high probability a great number of negative relative energy points near the ground truth, none of them being of significant interest to recover exactly our permutation. We may also study this relative energy far from the planted permutation, which would be interesting to address the problem of almost exact (resp. partial) alignment, which consists in finding an estimator $\hat{\pi}$ that coincides with π on at least $n - o(n)$ (resp. some positive fraction of n) points. In the light of our result which shows that exact recovery is not more difficult than detection, we can also conjecture that the same threshold $n\rho^2 / \log n = 4$ is sharp for the tasks of almost exact and partial recovery.*

Acknowledgments

The author would like to thank Léo Miolane, Marc Lelarge, and Laurent Massoulié for helpful discussions. This work was supported by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR19-P3IA-0001 (PRAIRIE 3IA Institute).

References

- A. C. Berg, T. L. Berg, and J. Malik. Shape matching and object recognition using low distortion correspondences. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 26–33 vol. 1, 2005. doi: 10.1109/CVPR.2005.320.
- Timothee Cour, Praveen Srinivasan, and Jianbo Shi. Balanced graph matching. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 313–320. MIT Press, 2007. URL <http://papers.nips.cc/paper/2960-balanced-graph-matching.pdf>.
- Daniel Cullina and Negar Kiyavash. Exact alignment recovery for correlated Erdős-Rényi graphs, 2017.
- Daniel Cullina, Negar Kiyavash, Prateek Mittal, and H. Vincent Poor. Partial recovery of Erdős-Rényi graph alignment via k-core alignment. *CoRR*, abs/1809.03553, 2018a. URL <http://arxiv.org/abs/1809.03553>.
- Daniel Cullina, P. Mittal, and N. Kiyavash. Fundamental limits of database alignment. *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 651–655, 2018b.
- Jian Ding, Zongming Ma, Yihong Wu, and Jiaming Xu. Efficient random graph matching via degree profiles. *arXiv e-prints*, art. arXiv:1811.07821, Nov 2018.
- Cynthia Dwork. Differential privacy: A survey of results. In Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li, editors, *Theory and Applications of Models of Computation*, pages 1–19, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. ISBN 978-3-540-79228-4.
- Osman Emre Dai, Daniel Cullina, and Negar Kiyavash. Database Alignment with Gaussian Features. *arXiv e-prints*, art. arXiv:1903.01422, March 2019.
- Zhou Fan, Cheng Mao, Yihong Wu, and Jiaming Xu. Spectral graph matching and regularized quadratic relaxations I: The gaussian model, 2019a.
- Zhou Fan, Cheng Mao, Yihong Wu, and Jiaming Xu. Spectral graph matching and regularized quadratic relaxations II: Erdős-Rényi graphs and universality, 2019b.
- L. Ganassali, M. Lelarge, and L. Massoulié. Spectral alignment of correlated Gaussian random matrices. *arXiv e-prints*, art. arXiv:1912.00231, November 2019.
- Luca Ganassali and Laurent Massoulié. From tree matching to sparse graph alignment. volume 125 of *Proceedings of Machine Learning Research*, pages 1633–1665. PMLR, 09–12 Jul 2020. URL <http://proceedings.mlr.press/v125/ganassali20a.html>.

- Aria D. Haghighi, Andrew Y. Ng, and Christopher D. Manning. Robust textual inference via graph matching. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, HLT '05, pages 387–394, Stroudsburg, PA, USA, 2005. Association for Computational Linguistics. doi: 10.3115/1220575.1220624. URL <https://doi.org/10.3115/1220575.1220624>.
- Georgina Hall and Laurent Massoulié. Partial Recovery in the Graph Alignment Problem. *arXiv e-prints*, art. arXiv:2007.00533, July 2020.
- D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *Ann. Math. Statist.*, 42(3):1079–1083, 06 1971. doi: 10.1214/aoms/1177693335. URL <https://doi.org/10.1214/aoms/1177693335>.
- H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. doi: 10.1002/nav.3800020109. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109>.
- Konstantin Makarychev, Rajsekar Manokaran, and Maxim Sviridenko. Maximum quadratic assignment problem: Reduction from maximum label cover and lp-based approximation algorithm. *CoRR*, abs/1403.7721, 2014. URL <http://arxiv.org/abs/1403.7721>.
- A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125, May 2008. doi: 10.1109/SP.2008.33.
- A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *2009 30th IEEE Symposium on Security and Privacy*, pages 173–187, May 2009. doi: 10.1109/SP.2009.22.
- Ashwin Pananjady, Martin J. Wainwright, and Thomas A. Courtade. Linear Regression with an Unknown Permutation: Statistical and Computational Limits. *arXiv e-prints*, art. arXiv:1608.02902, August 2016.
- Panos Pardalos, Franz Rendl, and Henry Wolkowicz. *The Quadratic Assignment Problem: A Survey and Recent Developments*, pages 1–42. 08 1994. doi: 10.1090/dimacs/016/01.
- Rohit Singh, Jinbo Xu, and Bonnie Berger. Global alignment of multiple protein interaction networks with application to functional orthology detection. *Proceedings of the National Academy of Sciences*, 105(35):12763–12768, 2008. ISSN 0027-8424. doi: 10.1073/pnas.0806627105. URL <https://www.pnas.org/content/105/35/12763>.
- Yihong Wu, Jiaming Xu, and Sophie H. Yu. Testing correlation of unlabeled random graphs. *arXiv e-prints*, art. arXiv:2008.10097, August 2020.
- Yihong Wu, Jiaming Xu, and Sophie H. Yu. Settling the sharp reconstruction thresholds of random graph matching, 2021.

Appendix A. Additional proofs

A.1. Proof of Lemma 2.1: lower bound on correlations of relative energies

Proof Recall that we work under event \mathcal{A} . Fix $\alpha \in (0, 1]$ and take $d = \alpha n$ and $\sigma, \sigma' \in \mathcal{S}_{n,d}$. The proof is obtained by establishing a fine lower bound on $\sharp(\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E)$, which is simply the number of edges that are deranged both by σ^E and σ'^E . In order to establish this lower bound, let us assume that σ and σ' have $\sharp(\mathcal{D}_\sigma \cap \mathcal{D}_{\sigma'}) = \beta n$ common unfixed points, with $\beta \in [0, \alpha]$. We then form edges in $\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E$ in the following way:

- First, by taking all pairs but the pairs made of points in the complement of $\mathcal{D}_\sigma \cap \mathcal{D}_{\sigma'}$ and those made of pairs (i, j) that are transpositions of σ or σ' , we obtain at least $\frac{1}{2}\beta(2 - \beta)n^2 - \alpha n$ edges.
- Then, add new edges made of one extremity in $\mathcal{D}_\sigma \setminus \mathcal{D}_{\sigma'}$ and one in $\mathcal{D}_{\sigma'} \setminus \mathcal{D}_\sigma$. Since \mathcal{D}_σ (resp $\mathcal{D}_{\sigma'}$) is stable by σ (resp. by σ'), all these $(\alpha - \beta)^2 n^2$ edges are in $\mathcal{D}_\sigma^E \cap \mathcal{D}_{\sigma'}^E$.

Finally we formed $g(\alpha, \beta)n^2 - \alpha n$ edges, with

$$g(\alpha, \beta) := \frac{1}{2}\beta^2 + (1 - 2\alpha)\beta + \alpha^2, \quad (28)$$

which is minimal on $[0, \alpha]$ at $\beta = 2\alpha - 1$ if $\alpha \geq 1/2$, or at $\beta = 0$ if $\alpha < 1/2$. In any case, this minimum is $f(\alpha)$. The first inequality is established by applying inequality (16) of event \mathcal{A} .

For the second part, consider a centered vector $Z = (Z_\sigma)_{\sigma \in \mathcal{S}_{n,\alpha n}}$ such that all Z_σ have same variance v_α and $\text{Cov}(Z_\sigma, Z_{\sigma'}) = c_\alpha$ for $\sigma \neq \sigma'$, with v_α, c_α defined as follows:

$$\begin{aligned} v_\alpha &:= \alpha(2 - \alpha)n^2 - C_1 n^{3/2} \log^{1/2} n, \\ c_\alpha &:= f(\alpha)n^2 - C_1 n^{3/2} \log^{1/2} n. \end{aligned}$$

for some $C_1 > 0$ large enough. Note that on event \mathcal{A} , for all $\alpha \in (0, 1]$, all $\sigma, \sigma' \in \mathcal{S}_{n,\alpha n}$,

$$\text{Cov}(Z_\sigma, Z_{\sigma'}) \leq \text{Cov}(X_\sigma, X_{\sigma'}),$$

so one has that for all $t > 0$,

$$\mathbb{P}\left(\max_{\sigma \in \mathcal{S}_{n,\alpha n}} X_\sigma > t \cap \mathcal{A}\right) \leq \mathbb{P}\left(\max_{\sigma \in \mathcal{S}_{n,\alpha n}} Z_\sigma > t\right). \quad (29)$$

We now control the right-hand side of (29) with this classical Lemma, which proof is find hereafter in Appendix A.2:

Lemma A.1 (Maximum of totally correlated Gaussian variables) *Let Z be a centered Gaussian vector of size N , such that all Z_i have same variance v and $\text{Cov}(Z_i, Z_j) = c$ for $i \neq j$. Then*

$$\mathbb{P}\left(\max_{1 \leq i \leq N} Z_i > \sqrt{2(v - c) \log N} + 2\sqrt{v \log \log N}\right) \leq \frac{2}{\log N}. \quad (30)$$

Note that for v_α, c_α previously defined, one has

$$\sqrt{2(v_\alpha - c_\alpha) \log \#\mathcal{S}_{n,\alpha n}} \leq \sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))n^{3/2} \log^{1/2} n}, \quad (31)$$

and for n large enough,

$$2\sqrt{v_\alpha \log \log \#\mathcal{S}_{n,\alpha n}} \leq 2\sqrt{\alpha(2 - \alpha)n} \sqrt{\log n + \log \log n} \leq (2 + \varepsilon')n \log^{1/2} n. \quad (32)$$

Finally, we use equations (29)–(32) to conclude that for n large enough:

$$\begin{aligned} & \mathbb{P} \left(\exists d = \alpha n, \alpha > \alpha_0, \max_{\sigma \in \mathcal{S}_{n,d}} X_\sigma > \sqrt{2\alpha(\alpha(2 - \alpha) - f(\alpha))n^{3/2} \log^{1/2} n} + (2 + \varepsilon')n \log^{1/2} n \right) \\ & \leq 1 - \mathbb{P}(\mathcal{A}) + \sum_{d=\alpha n, \alpha > \alpha_0} \mathbb{P} \left(\max_{\sigma \in \mathcal{S}_{n,\alpha n}} Z_i > \sqrt{2(v_\alpha - c_\alpha) \log \#\mathcal{S}_{n,\alpha n}} + 2\sqrt{v_\alpha \log \log \#\mathcal{S}_{n,\alpha n}} \right) \\ & \leq o(1) + \sum_{d=\alpha n, \alpha > \alpha_0} \frac{2}{\log \#\mathcal{S}_{n,\alpha n}} \leq o(1) + \frac{2n}{\log \#\mathcal{S}_{n,\alpha_0 n}} = o(1) + \frac{2}{\alpha_0 \log n} = o(1), \end{aligned}$$

and Lemma 2.1 is proved. \blacksquare

A.2. Proof of Lemma A.1: maximum of totally correlated Gaussian variables

Proof Let us make a change of variables which preserves the joint distribution:

$$(Z_1, Z_2, \dots, Z_N) = (\sqrt{c} \xi_0 + \sqrt{v - c} \xi_1, \dots, \sqrt{c} \xi_0 + \sqrt{v - c} \xi_N),$$

where ξ_0, \dots, ξ_N are independent standard Gaussian random variables. The maximum thus writes

$$\max_{1 \leq i \leq N} Z_i = \sqrt{c} \xi_0 + \sqrt{v - c} \max_{1 \leq i \leq N} \xi_i$$

Then, with the classical inequality $\mathbb{P}(\mathcal{N}(0, 1) \geq t) \leq e^{-t^2/2}$, then with probability at least $1 - 1/(\log N)$, one has:

$$\sqrt{c} \xi_0 \leq \sqrt{2c \log \log N}, \quad \text{and} \quad \sqrt{v - c} \max_{1 \leq i \leq N} \xi_i \leq \sqrt{2(v - c) \log N \left(1 + \frac{\log \log N}{\log N} \right)},$$

so with probability at least $1 - 2/(\log N)$:

$$\begin{aligned} \max_{1 \leq i \leq N} Z_i & \leq \sqrt{2(v - c) \log N} + \sqrt{2 \log \log N} (\sqrt{c} + \sqrt{v - c}) \\ & \leq \sqrt{2(v - c) \log N} + 2\sqrt{v \log \log N}, \end{aligned}$$

where we used $\sqrt{c} + \sqrt{v - c} \leq \sqrt{2v}$ in the last step. \blacksquare

A.3. Proof of Lemma 2.2**Proof** For $\alpha \in (0, 1]$,

$$\begin{aligned}
(20) \iff \alpha^2(2 - \alpha)^2 &\geq 2\alpha(\alpha(2 - \alpha) - f(\alpha)) \\
&\iff f(\alpha) \geq \alpha^2 - \alpha^3/2.
\end{aligned}$$

The inequality is verified for $\alpha < 1/2$. To conclude the proof of (20), it remains to check that for $1 \geq \alpha \geq 1/2$, $f(\alpha) \geq \alpha^2 - \alpha^3/2$, which is equivalent to

$$\begin{aligned}
\alpha^2 - \frac{1}{2}(2\alpha - 1)^2 \geq \alpha^2 - \alpha^3/2 &\iff \alpha^3 - 4\alpha^2 + 4\alpha - 1 \geq 0 \\
&\iff (\alpha - 1)(\alpha^2 - 3\alpha + 1) \geq 0 \\
&\iff \alpha^2 - 3\alpha + 1 \leq 0 \iff \alpha \geq \frac{3 - \sqrt{5}}{2} \sim 0.382\dots
\end{aligned}$$

■

A.4. Proof of Lemma 3.1: Payley-Zygmund inequality**Proof** Using Cauchy-Schwarz inequality,

$$\begin{aligned}
\mathbb{E}[Y] &= \mathbb{E}[Y\mathbf{1}_{Y < c\mathbb{E}[Y]}] + \mathbb{E}[Y\mathbf{1}_{Y \geq c\mathbb{E}[Y]}] \\
&\leq c\mathbb{E}[Y] + \mathbb{E}[Y^2]^{1/2}\mathbb{P}(Y \geq c\mathbb{E}[Y])^{1/2},
\end{aligned}$$

which gives $\mathbb{E}[Y]^2(1 - c)^2 \leq \mathbb{E}[Y^2]\mathbb{P}(Y \geq c\mathbb{E}[Y])$.

■

A.5. Proof of Lemma 3.3: Control of deviation probabilities for correlated Gaussians**Proof** Let us first make a change of variable which preserves the joint distribution:

$$(Z_1, Z_2) = (Z, \alpha_n Z + \sqrt{1 - \alpha_n^2} Z'),$$

with Z, Z' two independent standard Gaussian variables.

Proof of (i): Note that standard Gaussian concentration gives $\mathbb{P}(Z > 2t_n | Z > t_n) \sim \frac{1}{2}e^{-3t_n^2/2}$. Thus, for n large enough

$$\begin{aligned}
\mathbb{P}(Z_1 > t_n, Z_2 > t_n) &\leq \mathbb{P}(Z > t_n)e^{-3t_n^2/2} + \mathbb{P}(Z > t_n)\mathbb{P}(\alpha_n Z + \sqrt{1 - \alpha_n^2} Z' > t_n, Z \leq 2t_n | Z > t_n) \\
&\leq e^{-2t_n^2} + \mathbb{P}(Z > t_n)\mathbb{P}(Z' > t_n - 2\alpha_n t_n + O(t_n \alpha_n^2)) \\
&\leq e^{-2t_n^2} + \mathbb{P}(Z > t_n)\mathbb{P}(Z' > t_n - o(1)) \\
&\leq e^{-2t_n^2} + (1 + o(1))\mathbb{P}(Z > t_n)\mathbb{P}(Z' > t_n) \\
&= e^{-2t_n^2} + (1 + o(1))\mathbb{P}(Z_1 > t_n)\mathbb{P}(Z_2 > t_n).
\end{aligned}$$

Proof of (ii): For any (s_n) such that $s_n \leq t_n$ for all n , one has

$$\begin{aligned} \mathbb{E} [e^{s_n Z} \mid Z > t_n] &= \frac{1}{\sqrt{2\pi}} \int_{t_n}^{+\infty} e^{s_n z - z^2/2} dz \left(\frac{1}{\sqrt{2\pi}} \int_{t_n}^{+\infty} e^{-z^2/2} dz \right)^{-1} \\ &= e^{s_n^2/2} \int_{t_n - s_n}^{+\infty} e^{-z^2/2} dz \left(\int_{t_n}^{+\infty} e^{-z^2/2} dz \right) \\ &\sim \frac{t_n}{t_n - s_n} \exp \left(s_n^2/2 - (t_n - s_n)^2/2 + t_n^2/2 \right) = \frac{t_n}{t_n - s_n} e^{s_n t_n}. \end{aligned}$$

Using independence of Z, Z' and Chernoff bound, we get, taking s_n such that $\alpha s_n = ut_n$ with $u < 1$, for n large enough,

$$\begin{aligned} \mathbb{P} \left(\alpha Z + \sqrt{1 - \alpha^2} Z' > t_n \mid Z > t_n \right) &\leq (1 + o(1)) \frac{t_n}{t_n - \alpha s_n} \exp \left(\alpha s_n t_n + \frac{1 - \alpha^2}{2} s_n^2 - s_n t_n \right) \\ &\leq (1 + o(1)) \frac{1}{1 - u} \exp \left(\left(u + \frac{u^2(1 - \alpha^2)}{2\alpha^2} - \frac{u}{\alpha} \right) t_n^2 \right) \\ &\stackrel{(a)}{\leq} (1 + o(1))(1 + \alpha) \exp \left(-\frac{1 - \alpha}{1 + \alpha} \cdot \frac{t_n^2}{2} \right) \end{aligned}$$

where we took $u = \frac{\alpha}{1 + \alpha} < 1$ in (a). The proof follows from this last inequality, together with the bound $\mathbb{P}(Z > t_n) \leq \frac{1}{\sqrt{2\pi}t_n} \exp \left(-\frac{t_n^2}{2} \right)$. ■