Phase transition of the maximal overlap of two independent random geometric graphs
Course project report for 6.S896

Hang Du*

December 13, 2023

Abstract

In this note, we consider the maximal overlap among vertex bijections between two independent d-dimensional random geometric graphs, each with n vertices and an average degree of n^θ, where $\theta \in (0,1)$. It is known that the maximal overlap, compared to the size of a single graph, undergoes a phase transition from $1-o(1)$ to $o(1)$ as d increases. we provide upper and lower bounds for the transition threshold that are tight up to $\log\log n$ factors. Additionally, we conjecture that the transition occurs at $d \asymp \log n/\log\log n$.

1 Introduction

The *random geometric graph* is a probabilistic model for graphs with inherent geometric characteristics. We fix two integers, n and d, to represent the size and dimension of the graph, respectively. Let ρ denote the uniform measure on the d-dimensional sphere $\mathbb{S}^d$. The random geometric graph is defined as follows.

Definition 1.1 (Random geometric graph with connecting threshold τ)**.** *Let $u_1, \dots, u_n \overset{i.i.d.}{\sim} \rho$. We construct a graph G on the set $[n]$ by forming edges between pairs (i,j) if and only if $u_i \cdot u_j \ge \tau$. The law of G is denoted by $\mathcal{G}(n,d,\tau)$.*

In this note, we will fix a constant $\theta \in (0,1)$ and define $\tau_* = \tau(n,d,\theta)$ such that a graph $G \sim \mathcal{G}(n,d,\tau_*)$ has an average degree of n^θ. Additionally, we will use $\mathcal{G}$ as shorthand for $\mathcal{G}(n,d,\tau_*)$. We make the following heuristic observations:

- When d is relatively small compared to n, $G \sim \mathcal{G}$ essentially acts as a discretization of $\mathbb{S}^d$ (see, for example, [5] for the case where $d = O(1)$). Therefore, in the low-dimensional regime, G exhibits strong geometric rigidity, and there is essentially "no randomness" in $\mathcal{G}$.
- When d is large enough relative to n, it can be shown that $G \sim \mathcal{G}$ becomes indistinguishable from an Erdős–Rényi graph on n vertices with an average degree of n^θ (refer to [2] for the case when $d \gg n^3$, and see also [1, 6] for some subsequent improvements). Thus, in sharp contrast to the low-dimensional regime, G is "purely random", devoid of any observable geometric properties.

These observations suggest that there must be a certain geometry-randomness phase transition in the random geometric graph as the dimension increases. This phenomenon has been one of the central

*Department of Mathematics, MIT, hangdu@mit.edu.

topics in the study of high-dimensional random geometric graphs over the past decade. While significant progress has been made in determining the regime where randomness predominates, as evidenced by [2, 1, 6], to the best of the author's knowledge, there are few results on the threshold at which geometric rigidity begins to diminish and randomness starts playing a substantial role. In this note, we take a step towards answering this latter question by studying the maximal overlap of two instances independently sampled from $\mathcal{G}$. We begin with the following definition.

Definition 1.2. *For two graphs G_1, G_2 on $[n]$ and a permutation $\pi \in \mathrm{S}_n$, define*

$$\mathrm{O}(G_1, G_2, \pi) := \sum_{1 \le i < j \le n} G^1_{i,j} G^2_{\pi(i),\pi(j)} \tag{1.1}$$

(where $(G^1_{i,j})$ and $(G^2_{i,j})$ are the adjacency matrices for G_1 and G_2, respectively), and we further denote that

$$\Lambda(G_1, G_2) := \frac{\max_{\pi \in \mathrm{S}_n} \mathrm{O}(G_1, G_2, \pi)}{|E(G_1)| \wedge |E(G_2)|} \,. \tag{1.2}$$

It is clear from the definition that $\Lambda(G_1, G_2) \le 1$. Additionally, if $\Lambda(G_1, G_2)$ is close to 1 and the cardinalities of $E(G_1)$ and $E(G_2)$ are approximately the same, then G_1 and G_2 appear very similar. One way to assess the geometric rigidity of a random geometric graph sampled from $\mathcal{G}$ is to examine the behavior of $\Lambda(G_1, G_2)$ for two independent graphs $G_1, G_2 \sim \mathcal{G}$. In the low-dimensional case when $d = O(1)$, it can be inferred from [5, Lemma 2.1] that $\Lambda(G_1, G_2)$ is typically $1 - o(1)$. Conversely, in the high-dimensional case when G_1, G_2 are indistinguishable from a pair of independent Erdős–Rényi graphs, a straightforward union bound implies that $\Lambda(G_1, G_2) = o(1)$ with high probability (see also [3, 4] for a detailed study of the Erdős–Rényi graph case). Heuristically, the threshold at which randomness becomes more prominent more or less corresponds to the point where $\Lambda(G_1, G_2)$ starts to deviate from 1, prompting us to investigate the transition of $\Lambda(G_1, G_2)$ as d increases.

In this note, we provide sufficient conditions on d in terms of n such that with high probability, $\Lambda(G_1, G_2) = 1 - o(1)$ or $o(1)$. In particular, our results determine the transition threshold of the dimension d for $\Lambda(G_1, G_2)$ transitioning from $1 - o(1)$ to $o(1)$, accurate up to a $\log \log n$ factor.

Theorem 1.3. *Let $d_0 = \frac{\log n}{\log \log n}$. Then, for any constant $\theta \in (0, 1)$, it holds that for some constant $\lambda = \lambda(\theta) > 0$,*

$$d \le \lambda d_0 \Rightarrow \mathbb{P}[\Lambda(G_1, G_2) = 1 - o(1)] = 1 - o(1) \,, \tag{1.3}$$

and

$$d \gg \log n \Rightarrow \mathbb{P}[\Lambda(G_1, G_2) = o(1)] = 1 - o(1) \,, \tag{1.4}$$

where the probability $\mathbb{P}$ is taken over $(G_1, G_2) \sim \mathcal{G}^{\otimes 2}$.

We expect our lower bound (1.3) to be tight. More precisely, we formulate the following conjecture:

Conjecture 1.4. *With the same notations, it holds that for some constant $\lambda' = \lambda'(\theta) > 0$,*

$$d \ge \lambda' d_0 \Rightarrow \mathbb{P}[\Lambda(G_1, G_2) = o(1)] = 1 - o(1) \,. \tag{1.5}$$

The note is organized as follows: in Section 2, we introduce several basic properties and estimates for random geometric graphs. In Section 3 and Section 4, we prove the lower bound (1.3) and the upper bound (1.4), respectively. We conclude this note with some further discussions on potential approaches towards addressing the main conjecture.

2 Preliminaries on random geometric graphs

In this section, we present some basic properties of random geometric graphs, which will be useful in later proofs. We begin with a handy estimation on the dot-product of vectors uniformly chosen on $\mathbb{S}^d$.

Proposition 2.1. *For any fixed $u \in \mathbb{S}^d$ and $v \sim \rho$, the dot product $u \cdot v$ has a probability density*

$$\psi(x) = \frac{\Gamma\left(\frac{d+1}{2}\right)}{\Gamma\left(\frac{d}{2}\right)\sqrt{\pi}}(1-x^2)^{\frac{d-2}{2}} \triangleq C_d(1-x^2)^{\frac{d-2}{2}},\ x \in [-1,1]\,. \tag{2.1}$$

Furthermore, for any $0 \le r \le 2$, it holds that

$$\frac{C_d}{2ed}\left(r^2 - \frac{r^4}{4}\right)^{\frac{d}{2}} \le \mathbb{P}[\mathrm{d}(u,v) \le r] \le \frac{C_d}{2}r^2\left(r^2 - \frac{r^4}{2}\right)^{\frac{d-2}{2}}\,. \tag{2.2}$$

Proof. (2.1) follows from standard calculations (see, e.g., [1, Lemma 5.1]), and given this, we have (note that $\mathrm{d}(u,v) \le r \iff u \cdot v \ge 1 - \frac{r^2}{2}$),

$$\mathbb{P}[\mathrm{d}(u,v) \le r] = \int_{1-\frac{r^2}{2}}^{1} C_d(1-x^2)^{\frac{d-2}{2}}\,\mathrm{d}x \le \frac{C_d r^2}{2}\left(1 - \left(1-\frac{r^2}{2}\right)^2\right)^{\frac{d-2}{2}} = \frac{C_d r^2}{2}\left(r^2 - \frac{r^4}{4}\right)^{\frac{d-2}{2}},$$

which gives the upper bound in (2.2). On the other hand, let $r' = \frac{r^2 - \frac{r^4}{4}}{d}$, then

$$\mathbb{P}[\mathrm{d}(u,v) \le r] \ge \int_{1-\frac{r^2}{2}}^{1-\frac{r^2}{2}+r'} C_d(1-x^2)^{\frac{d-2}{2}}\,\mathrm{d}x \ge \frac{C_d\left(r^2 - \frac{r^4}{4}\right)}{2d}\left(1-\frac{1}{d}\right)^{\frac{d-2}{2}}\left(r^2 - \frac{r^4}{4}\right)^{\frac{d-2}{2}},$$

which is bounded below by $\frac{C_d}{2ed}(r^2 - \frac{r^4}{4})^{\frac{d}{2}}$. This verifies (2.2) and completes the proof. □

The next lemma considers the degree of a vertex in the random geometric graph.

Lemma 2.2. *For any $\tau \in [-1,1]$ and $v \in [n]$, let $G \sim \mathcal{G}(n,d,\tau)$. Then the degree d_v of v in G follows the binomial distribution $\mathbf{B}(n,p)$, where $p = \mathbb{P}[\mathrm{d}(u,v) \le \sqrt{2-2\tau}]$. Furthermore, for any $K > 0$, it holds that*

$$\mathbb{P}[|d_v - np| \ge K] \le 2\exp\left(-\frac{K^2}{2(np+K)}\right)\,. \tag{2.3}$$

Proof. The first claim follows from the fact that the vectors $u_1, \dots, u_n$ which generate G are independent. Meanwhile, (2.3) is a standard estimate on the tail probability of binomial variables. □

3 The lower bound

In this section, we assume that $d \le \frac{\theta d_0}{10}$ and prove $\mathbb{P}[\Lambda(G_1,G_2) = 1 - o(1)] = 1 - o(1)$. Note that this proves (1.3) by taking $\lambda = \frac{\theta}{10}$. We begin with a lemma that relates the connecting radius to the size of the neighborhood in the random geometric graph.

Lemma 3.1. *For any constant $\theta' \in (0,1)$, let $r_{\theta'}$ be such that $\mathbb{P}_{u,v\sim\rho^{\otimes 2}}[\mathrm{d}(u,v) \le r_{\theta'}] = n^{-1+\theta'}$. Then it holds that $r_{\theta'} = n^{\frac{-1+\theta'+o(1)}{d}}$.*

Proof. By the definition of $r_{\theta'}$ and the estimations given in Proposition 2.1 (together with the fact that $C_d/2 = n^{o(1)}$ and $C_d/2ed = n^{o(1)}$ when $d \ll d_0$), we have

$$n^{\frac{-2+2\theta'+o(1)}{d}} \le r_{\theta'}^2 - \frac{r_{\theta'}^4}{4} \le r_{\theta'}^2 \le n^{-\frac{2-2\theta'+o(1)}{d}},$$

which implies $r_{\theta'} = n^{\frac{-1+\theta'+o(1)}{d}}$, as desired. □

The core of the proof of (1.3) lies in the following proposition, which is largely inspired by [5].

Proposition 3.2. *For $d \le \frac{\theta d_0}{10}$, let $u_1, \dots, u_n, v_1, \dots, v_n \overset{i.i.d.}{\sim} \rho$. Then, with high probability, there exists a permutation $\pi \in \mathrm{S}_n$ such that for all but $o(n)$ many $i \in [n]$, it holds that*

$$\mathrm{d}(u_i, v_{\pi(i)}) \le n^{-\frac{-1+\theta/3+o(1)}{d}}. \tag{3.1}$$

With Proposition 3.2 in hand, we may now proceed to finish the proof of (1.3) as follows.

Proof of the lower bound (1.3). Recall the definition of $r_{\theta'}$, where $\theta' \in (0,1)$, as in Lemma 3.1. Let $u_1, \dots, u_n \overset{\text{i.i.d.}}{\sim} \rho$, and let G_1 be the random geometric graph generated by forming edges between pairs (i,j) which satisfy $\mathrm{d}(u_i, u_j) \le r_\theta$. We note that $G_1 \sim \mathcal{G}$ by the definition of τ_* and $\mathcal{G} = \mathcal{G}(n, d, \tau_*)$. Moreover, let G_1' be the graph on $[n]$ constructed by forming edges between the pairs (i,j) such that $\mathrm{d}(u_i, u_j) \le r_\theta - r_{\theta/2}$. It is clear that G_1' is a subgraph of G_1. In addition, denoting $p_* = \mathbb{P}[\mathrm{d}(u,v) \le r_\theta] = n^{-1+\theta}$ and $p_{**} = \mathbb{P}[\mathrm{d}(u,v) \le r_\theta - r_{\theta/2}]$ (where $u, v \sim \rho^{\otimes 2}$) as the edge densities of G_1 and G_1', respectively. We then note that

$$\begin{aligned} p_* - p_{**} &= \mathbb{P}[r_\theta - r_{\theta/2} < \mathrm{d}(u,v) \le r_\theta] \\ &\overset{(2.1)}{=} \int_{1-r_\theta^2/2}^{1-(r_\theta - r_{\theta/2})^2/2} C(1-x^2)^{\frac{d-2}{2}} \,\mathrm{d}x \le C_d r_\theta r_{\theta/2} (r_\theta^2 - r_\theta^4/4)^{\frac{d-2}{2}} \\ &\overset{(2.2)}{\le} 2ed \cdot \frac{r_{\theta/2}}{r_\theta} \cdot \mathbb{P}[\mathrm{d}(u,v) \le r_\theta] \overset{\text{Lemma 3.1}}{\le} dn^{-\frac{\theta/2+o(1)}{d}} p_*, \end{aligned}$$

which is $o(p_*)$ since $dn^{-\theta/2d} \ll 1$ by our assumption that $d \le \frac{\theta d_0}{10}$. This suggests that $p_{**} = (1-o(1))p$. Therefore, both G_1 and G_1' have average degrees of $(1+o(1))np_*$.

Now let $v_1, \dots, v_n \overset{\text{i.i.d.}}{\sim} \rho$, and let $G_2 \sim \mathcal{G}$ be generated from $v_1, \dots, v_n$. Define $\mathcal{S}_1$ as the event that $u_1, \dots, u_n, v_1, \dots, v_n$ satisfies the condition in Proposition 3.2, and define $\mathcal{S}_2$ as the event that each vertex in G_1, G_1' and G_2 has degree $(1+o(1))np_*$. From Proposition 3.2 and Lemma 2.2, together with a union bound, $\mathbb{P}[\mathcal{S}_1 \cap \mathcal{S}_2] = 1 - o(1)$.

On the event $\mathcal{S}_1 \cap \mathcal{S}_2$, let $\pi \in \mathrm{S}_n$ and $I \subset [n]$ be the witnesses of $\mathcal{S}_1$, i.e., $|I| = n - o(n)$ and $\mathrm{d}(u_i, v_{\pi(i)}) \le n^{\frac{-1+\theta/3+o(1)}{d}}$ for $i \in I$. Note that for $i, j \in I$ and $(i,j) \in E(G_1')$, it holds that (recall that $r_{\theta'} = n^{\frac{-1+\theta'+o(1)}{d}}$ for $\theta' \in (0,1)$ by Lemma 3.1)

$$\mathrm{d}(v_{\pi(i)}, v_{\pi(j)}) \le \mathrm{d}(u_i, v_{\pi(i)}) + \mathrm{d}(u_i, u_j) + \mathrm{d}(u_j, v_{\pi(j)}) \le r_\theta - r_{\theta/2} + 2n^{\frac{-1+\theta/3+o(1)}{d}} \le r_\theta,$$

and thus $(\pi(i), \pi(j)) \in E(G_2)$. This implies that $O(G_1, G_2, \pi)$ is at least the number of edges in the induced subgraph of G_1' on I. Under the event $\mathcal{S}_2$, the induced subgraph has at least

$$|E(G_1')| - (1+o(1))np_*|[n] \setminus I| \ge (1/2 - o(1))n^{1+\theta} = (1+o(1))|E(G_1)|,$$

which implies that

$$\Lambda(G_1, G_2) \ge \frac{|O(G_1, G_2, \pi)|}{|E(G_1)|} \ge 1 - o(1),$$

and thus $\mathbb{P}[\Lambda(G_1, G_2) \ge 1 - o(1)] \ge \mathbb{P}[\mathcal{S}_1 \cap \mathcal{S}_2] = 1 - o(1)$, as desired. □

Now we turn to the proof of Proposition 3.2. For a set of subsets $A_1, \ldots, A_N \subset [n]$, we say $(a_1, \ldots, a_N)$ is a system of distinct representatives (SDR) for $(A_1, \ldots, A_N)$ if $a_i \in A_i$ for any $1 \le i \le N$ and $a_1, \ldots, a_N$ are distinct. Recall Hall's theorem regarding the existence of an SDR: a collection $(A_1, \ldots, A_N)$ has an SDR if and only if for any $1 \le k \le N$ and any $1 \le i_1 < \cdots < i_k \le N$, it holds that $|A_{i_1} \cup \cdots \cup A_{i_k}| \ge k$.

Proof of Proposition 3.2. Let $u_1, \ldots, u_n, v_1, \ldots, v_n \stackrel{\text{i.i.d.}}{\sim} \rho$ and $N = \lfloor n - n/\log n \rfloor = (1 - o(1))n$. For each $1 \le i \le N$, we define

$$A_i = \{j \in [n] : \mathrm{d}(u_i, v_j) \le r_{\theta/3}\},$$

and for each $1 \le j \le n$, we define $N_j = |\{i \in [N] : j \in A_i\}|$. It is clear that for each $1 \le i \le N$, $|A_i| \sim \mathbf{B}(n, n^{-1+\theta/3})$ and for each $1 \le j \le n$, $N_j \sim \mathbf{B}(N, n^{-1+\theta/3})$. Then it follows from Lemma 2.2, together with a union bound, that with high probability,

$$|A_i| \ge n^{\theta/3} - n^{\theta/6} \log n, \quad \forall 1 \le i \le N,$$

and similarly (note that $N_j \sim \mathbf{B}(N, p)$)

$$N_j \le N n^{-1+\theta/3} + n^{\theta/6} \log n \le n^{\theta/3} - \frac{n^{\theta/3}}{\log n} + n^{\theta/6} \log n, \quad \forall 1 \le j \le n.$$

Therefore, it holds with high probability that

$$\min_{1 \le i \le N} |A_i| \ge \max_{1 \le j \le n} N_j. \tag{3.2}$$

Now we verify Hall's condition under (3.2): for any $1 \le k \le N$ and any $1 \le i_1 < \cdots < i_k \le N$, let $B = A_{i_1} \cup \cdots \cup A_{i_k}$. We have that

$$k \min_{1 \le i \le N} |A_i| \le |\{(i, j) : i \in \{i_1, \ldots, i_k\}, j \in [n], \mathrm{d}(u_i, v_j) \le r_{\theta/3}\}| \le B \max_{1 \le j \le n} N_j \,,$$

and thus $|B| \ge k$ by (3.2). This verifies the Hall's condition and hence $(A_1, \ldots, A_d)$ has a SDR $(j_1, \ldots, j_N)$. Let $\pi \in \mathrm{S}_n$ satisfies $\pi(i) = j_i, 1 \le i \le N$, then by our definition of A_i we see $\mathrm{d}(u_i, v_{\pi(i)}) \le r_{\theta/3}$ for any $1 \le i \le N = (1 - o(1))n$. Combined with Lemma 3.1, this completes the proof. $\square$

4 The upper bound

This section is devoted to proving the upper bound (1.4). We fix some $d \gg \log n$ along with an arbitrary constant $\delta \in (0, 1)$. It suffices to show that

$$\mathbb{P}[\Lambda(G_1, G_2) \ge 2\delta] \le o(1) \,. \tag{4.1}$$

Given that this is true, we can obtain the desired result by letting $\delta \to 0$.

4.1 Random sparsification

The starting point is to attempt to show that for each fixed $\pi \in \mathrm{S}_n$, $\mathbb{P}[\mathrm{O}(G_1, G_2, \pi) \ge \delta n^{1+\theta}] = o(1/n!)$. Once this is established, the desired result follows by applying a union bound over all $\pi \in \mathrm{S}_n$. Unfortunately, the distribution of $\mathrm{O}(G_1, G_2, \pi)$ is far from clear, and even after some simplification, it remains difficult to analyze. This complexity is largely due to the denseness of G_1 and G_2, which results in complicated correlations within $\mathrm{O}(G_1, G_2, \pi)$. However, it turns out that sparse random geometric graphs (where the average degree is $O(1)$) are much more tractable. For instance, many tools for the sparse regime have been developed in [6], and this note draws significant inspiration from this seminal work. In light of this, our first step is to perform a random "sparsification" of G_1.

Definition 4.1. *Fix a constant M such that $M > 100(1-\theta)/\delta^2$. Define $s = Mn^{-\theta}$. Consider $G_1 \sim \mathcal{G}$. We sample a random subgraph $H = H(G_1, s)$ from G_1 by retaining each edge of G_1 in H with probability s, independently. Observe that H possesses an average degree of $M = O(1)$, as determined by our choice of M.*

After sparsifying G_1 to H, we claim that (4.1) simplifies to

$$\mathbb{P}[\Lambda(H, G_1) \geq 2\delta + o(1)] = o(1)\,, \tag{4.2}$$

where $\Lambda(H, G_2) = \frac{\max_{\pi \in \mathrm{S}_n} \mathrm{O}(H,G_2,\pi)}{|E(H)| \wedge |E(G_2)|}$, as defined in Definition 1.2. This is because, with high probability, $|E(G_1)| \wedge |E(G_2)| = (1+o(1))n^{1+\theta}/2$ and $|E(H)| \wedge |E(G_2)| = (1+o(1))Mn/2$, leading to

$$\mathbb{P}[\Lambda(G_1, G_2) \geq 2\delta] \leq \mathbb{P}[\max_{\pi \in \mathrm{S}_n} \mathrm{O}(G_1, G_2, \pi) \geq (1+o(1))n^{1+\theta/2}/2] + o(1)\,,$$

and

$$\mathbb{P}[\Lambda(H, G_2) \geq 2\delta + o(1)] \geq \mathbb{P}[\max_{\pi \in \mathrm{S}_n} \mathrm{O}(H, G_2, \pi) \geq (1+o(1))Mn/2] - o(1)\,.$$

Additionally, for any (G_1, G_2) such that $\max_{\pi \in \mathrm{S}_n} \mathrm{O}(G_1, G_2, \pi) = \mathrm{O}(G_1, G_2, \pi^*) \geq (1+o(1))n^{1+\theta}/2$, we find that

$$\max_{\pi \in \mathrm{S}_n} \mathrm{O}(H, G_2, \pi) \geq \mathrm{O}(H, G_2, \pi^*) \sim \mathbf{B}(\mathrm{O}(G_1, G_2, \pi^*), s)\,,$$

which exceeds $(1+o(1))sn^{1+\theta}/2 = (1+o(1))Mn/2$ with high probability. Therefore,

$$\begin{aligned}
\mathbb{P}[\Lambda(G_1, G_2) \geq 2\delta] &\leq \mathbb{P}[\max_{\pi \in \mathrm{S}_n} \mathrm{O}(G_1, G_2, \pi) \geq (1+o(1))n^{1+\theta}/2] + o(1) \\
&\leq (1+o(1))\mathbb{P}[\max_{\pi \in \mathrm{S}_n} \mathrm{O}(H, G_2, \pi) \geq (1+o(1))Mn/2] + o(1) \\
&\leq (1+o(1))\mathbb{P}[\Lambda(H, G_2) \geq 2\delta + o(1)] + o(1)\,,
\end{aligned}$$

thus verifying the claim.

4.2 Truncations

Now, we proceed to the proof of (4.2). For this purpose, it is necessary to introduce some appropriate truncations on both H and G_2. We begin by specifying the desired properties for H.

Definition 4.2. *We define a vertex $v \in [n]$ as good in H if the following conditions are met:*

- *The 3-neighborhood of v forms a tree structure;*
- *For any vertex u within the 3-neighborhood of v, the degree of u in H does not exceed $2M$.*

A vertex that does not satisfy these conditions is considered bad. Additionally, we denote by $\mathcal{H}$ the event that $|E(H)| = (1+o(1))Mn/2$, and the sum of degrees of bad vertices is at most $\delta Mn/3$.

Lemma 4.3. *For $H = H(s, G_1)$ as in Definition 4.1, the event $\mathcal{H}$ occurs with probability $1 - o(1)$.*

The lemma follows from standard facts about sparse random graphs, and we omit the details here. Next, we consider a slight modification to the random geometric graph model $\mathcal{G}$.

Definition 4.4. *Let $\mu = \rho^{\otimes n}$. For vectors $(u_1, \dots, u_n) \sim \mu$, define $\mathcal{U}$ as the event that $|u_i \cdot u_j| \leq \gamma$ for any $1 \leq i \neq j \leq n$, with $\gamma := (\log n/d)^{1/4}$. Furthermore, denote by $\mathcal{G}'$ the distribution of a graph constructed by forming edges between pairs (i,j) whenever $u_i \cdot u_j \geq \tau_*$, where $(u_1, \dots, u_n)$ is sampled from $\mu[\cdot \mid \mathcal{U}]$.*

We remark that such a truncation will facilitate our analysis of joint probabilities of the form $\mathbb{P}_{v\sim\rho}[v \cdot u_{i_1} \geq \tau_*, \ldots, v \cdot u_{i_K} \geq \tau_*]$ (see Lemma 4.9 for more details). The following lemma shows the close similarity between $\mathcal{G}$ and its modified version $\mathcal{G}'$.

Lemma 4.5. $\mu[\mathcal{U}] = 1 - n^{-\omega(1)}$ *and the total variation distance* $\mathrm{TV}(\mathcal{G}, \mathcal{G}') = 1 - n^{-\omega(1)}$.

Proof. From (2.1), for any distinct i and j, we have

$$\mu[|u_i \cdot u_j| \geq \gamma] \leq 2C_d \left(1 - \sqrt{\frac{\log n}{d}}\right)^{\frac{d-2}{2}} \leq \exp\left(-\Omega\left(\sqrt{d \log n}\right)\right),$$

which is $n^{-\omega(1)}$ as $d \gg \log n$. By applying a union bound, we deduce that $\mu[\mathcal{U}^c] \leq n^2 \cdot n^{-\omega(1)} = n^{-\omega(1)}$. This confirms that $\mu[\mathcal{U}] = 1 - n^{-\omega(1)}$. Consequently, by the data processing inequality, the total variation distance $\mathrm{TV}(\mathcal{G}, \mathcal{G}') \leq \mathrm{TV}(\mu, \mu[\cdot \mid \mathcal{U}]) = n^{-\omega(1)}$, as required. □

It is straightforward to see that

$$\mathbb{P}[\Lambda(H, G_2) \leq 2\delta + o(1)] \leq \mathbb{P}[H \notin \mathcal{H}] + \mathbb{P}[H \in \mathcal{H}, \Lambda(H, G') \leq 2\delta + o(1)] + \mathrm{TV}(\mathcal{G}, \mathcal{G}'),$$

where in the second term G' is sampled from $\mathcal{G}'$. In light of Lemma 4.3 and Lemma 4.5, (4.2) is reduced to proving that for any $H \in \mathcal{H}$, we have $\mathbb{P}[\Lambda(H, G) \leq 2\delta + o(1)] \leq \delta + o(1)$. Here, the probability is taken over the random graph $G \sim \mathcal{G}'$. Furthermore, since $|E(H)| = (1 + o(1))Mn/2$ under $\mathcal{H}$, it remains to show that for any $H \in \mathcal{H}$,

$$\mathbb{P}_{G\sim\mathcal{G}'}[\max_{\pi\in \mathrm{S}_n} \mathrm{O}(H, G, \pi) \leq (1 + o(1))\delta Mn] \leq \delta + o(1). \tag{4.3}$$

4.3 Domination by binomial variable

In this discussion, we consider a fixed graph $H \in \mathcal{H}$ and aim to prove (4.3). Consider any permutation $\pi \in \mathrm{S}_n$ and an index k such that $1 \leq k \leq n$. Let $\mathcal{F}^\pi_{k-1}$ represent the σ-field generated by the set of random variables $\{G_{\pi(i),\pi(j)} : (i, j) \in E(H), i < j < k\}$. The overlap $\mathrm{O}(H, G, \pi)$ can be written as

$$\mathrm{O}(H, G, \pi) = \sum_{(i,j)\in E(H)} G_{\pi(i),\pi(j)} = \sum_{k=1}^{n} O_k(\pi), \tag{4.4}$$

where $O_k(\pi)$ is defined as $O_k(\pi) = \sum_{j<k,(j,k)\in E(H)} G_{\pi(j),\pi(k)}$. It is evident that for $1 \leq j < k$, O_j is measurable with respect to $\mathcal{F}^\pi_{k-1}$. The crucial component for managing $\Lambda(H, G) = \max_{\pi\in \mathrm{S}_n} \mathrm{O}(H, G, \pi)$ involves a stochastic domination relation, as stated in the following proposition:

Proposition 4.6. *For any permutation $\pi \in \mathrm{S}_n$, any good vertex k in H, and any realization of $\mathcal{F}^\pi_{k-1}$, the conditioned random variable O_k given $\mathcal{F}^\pi_{k-1}$, is stochastically dominated by $\frac{\delta M}{3} + 2MI_k$, where $I_k \sim \mathbf{B}(1, n^{-10/\delta})$ denotes a Bernoulli indicator that is independent of $\mathcal{F}^\pi_{k-1}$.*

Proof of (4.3) *assuming Proposition 4.6.* For each fixed $\pi \in \mathrm{S}_n$, it is clear that

$$\sum_{k \text{ is bad in } H} O_k(\pi) \leq \sum_{k \text{ is bad in } H} d_H(k) \leq \frac{\delta Mn}{3}$$

by the condition $H \in \mathcal{H}$. In addition, we conclude from Proposition 4.6 that

$$\sum_{k \text{ is good in } H} O_k(\pi)$$

is stochastically dominated by $\frac{\delta M n}{3} + 2MX$, where $X \sim \mathbf{B}(n, n^{-10/\delta})$ is a binomial variable. From the Chernoff bound for binomial variables, we get that

$$\mathbb{P}[\mathrm{O}(H,G,\pi) \geq (1+o(1))\delta M n] \leq \mathbb{P}[X \geq \delta n/8] \leq \exp\left(-\frac{\delta n}{8} \cdot \left(\log\left(\frac{\delta n/8}{n^{1-\delta/10}}\right) - 1\right)\right),$$

which is upper-bounded by $\exp(-n \log n) \ll 1/n!$. The desired result then follows by taking a union bound over $\pi \in \mathrm{S}_n$. $\square$

The remainder of this section is dedicated to proving Proposition 4.6. Without loss of generality, we assume $\pi = \mathrm{id}$. For simplicity, we denote O_k for $\mathrm{O}_k(\mathrm{id})$ and $\mathcal{F}_{k-1}$ for $\mathcal{F}_{k-1}^{\mathrm{id}}$, where $1 \leq k \leq n$. We consider a fixed $k \in [n]$ which is a good vertex in H, and a specific realization of $\mathcal{F}_{k-1}$. Define the sets

$$\mathcal{C} = \{(i,j) \in E(H) : i < j < k, G_{i,j} = 1\}, \quad \mathcal{D} = \{(i,j) \in E(H) : i < j < k, G_{i,j} = 0\}.$$

Suppose G is generated from the vectors $u_1, \ldots, u_n \in \mathbb{S}^d$. The conditional joint distribution of $u_1, \ldots, u_n$, given $\mathcal{F}_{k-1}$ is then given by

$$\mu_k := \mu[\cdot \mid u_i \cdot u_j \geq \tau_*, \forall (i,j) \in \mathcal{C}, u_i \cdot u_j < \tau_*, \forall (i,j) \in \mathcal{D}, \mathcal{U}], \tag{4.5}$$

where $\mathcal{U}$ is as defined in Definition 4.4. Denote H_k as the induced subgraph of H on the vertex set $[k]$, and let $i_1, \ldots, i_D$ be the neighbors of k in H_k. The distribution of O_k given $\mathcal{F}_{k-1}$ corresponds to the distribution of the sum

$$\sum_{t=1}^{D} \mathbf{1}(u_{i_t} \cdot u_k \geq \tau_*)$$

under μ_k. While μ_k is a measure conditioned on an event with a sparse correlation structure, analyzing the joint probability on all coordinates $(u_1, \ldots, u_n)$ remains difficult. To tackle this, we fix the positions of u_i for vertices i distant from k in H, restricting the correlation to a local neighborhood of k.

Let T denote the 3-neighborhood of k in H_k, which is a tree rooted at k with a maximum degree not exceeding $2M$ under the assumption that k is good in H. For each $1 \leq t \leq D$, define T_t as the subtree of T rooted at i_t, and let $N(T_t)$ be the set of non-leaf vertices in T_t. Additionally, set

$$\mathcal{C}_t = \{(i,j) \in E(T_t) : G_{i,j} = 1\}, \quad \mathcal{D}_t = \{(i,j) \in E(T_t) : G_{i,j} = 0\},$$

and consider a realization of $(u_i)_{i \notin N(T_t)}$ to be compatible with $\mathcal{F}_{k-1}$ if it does not violate the conditioning in (4.5). The joint law on $(u_i)_{i \in N(T_t)}$ of μ_k, conditioned on a compatible realization $(u_i)_{i \notin N(T_t)}$, is then given by

$$\widetilde{\mu}_t := \widetilde{\mu}[\cdot \mid u_i \cdot u_j \geq \tau_*, \forall (i,j) \in \mathcal{C}_t, u_i \cdot u_j < \tau_*, \forall (i,j) \in \mathcal{D}_t, |u_i \cdot u_j| \leq \gamma, \forall i \in N(T_t), j \in [n]], \tag{4.6}$$

where $\widetilde{\mu}$ is the product measure of uniform distributions on $(u_i)_{i \in N(T_t)}$.

For a probability measure ν on $\mathbb{S}^d$ with probability density $\nu(x)$ with respect to ρ, we define

$$F(\nu) = \sup_{x \in \mathbb{S}^d} \mathbb{P}_{y \sim \nu}[x \cdot y \geq \tau_*] = \sup_{x \in \mathbb{S}^d} \int_{y \cdot x \geq \tau_*} \nu(y) \, \mathrm{d}\rho(y). \tag{4.7}$$

We claim that it suffices to show the following control on $F(\widetilde{\nu}_t)$ for $\widetilde{\nu}_t$ be the marginal of u_{i_t} under $\widetilde{\mu}_t$.

Proposition 4.7. *For any realization of $\mathcal{F}_{k-1}$, any $1 \leq t \leq D$, and any compatible specification of $(u_i)_{i \notin N(T_t)}$, let $\widetilde{\mu}_t$ be defined as in (4.6) and let $\widetilde{\nu}_t$ be the marginal of u_{i_t} under $\widetilde{\mu}_t$, then $F(\widetilde{\nu}_t) \leq n^{o(1)} p$.*

Proof of Proposition 4.6 assuming Proposition 4.7. For any distinct $t_1, \dots, t_k \in \{i_1, \dots, i_D\}$, we apply the total probability formula to obtain

$$\mu_k[u_{t_j} \cdot u_k \geq \tau_* \mid u_{t_l} \cdot u_k \geq \tau_*, l < j] = \mathbb{E}_{(u_i)_{i \notin N(T_{t_j})}} \widetilde{\nu}_{t_j}[u_{t_j} \cdot u_k \geq \tau_*] \leq \mathbb{E}_{(u_i)_{i \notin N(T_{t_j})}} F(\widetilde{\nu}_{t_j}),$$

where the expectation over $(u_i)_{i \notin N(T_{t_j})}$ is taken under the conditional measure $\mu_k[\cdot \mid u_{t_l} \cdot u_k \geq \tau_*, l < j]$. This implies that the realizations of $(u_i)_{i \notin N(T_{t_j})}$ are almost surely compatible. By Proposition 4.7, we have $F(\widetilde{\nu}_t) \leq n^{o(1)} p = n^{-1+\theta+o(1)}$ almost surely, and therefore

$$\mu_k[u_{t_j} \cdot u_k \geq \tau_* \mid u_{t_l} \cdot u_k \geq \tau_*, l < j] \leq n^{-1+\theta+o(1)}.$$

Utilizing the multiplicative rule, we conclude that

$$\mu_k[u_{t_j} \cdot u_k \geq \tau_*, 1 \leq j \leq k] \leq (n^{-1+\theta+o(1)})^k.$$

A union bound then yields

$$\widetilde{\mu}_k[O_k \geq \delta M/3] \leq \binom{2M}{\delta M/3} (n^{-1+\theta+o(1)})^{\delta M/3} \leq n^{-10/\delta},$$

where the second inequality is valid since M is selected to be a constant greater than $100(1-\theta)/\delta^2$. Moreover, it trivially holds that $O_k \leq 2M$, leading to the conclusion that $O_k \leq \delta M/3 + 2M I_k$ for some $I_k \sim \mathbf{B}(1, n^{-10/\delta})$. □

4.4 Compute the marginal via belief propagation

We now turn our attention to Proposition 4.7. We begin by fixing a realization of $\mathcal{F}_{k-1}$ and a $t \in \{1, \dots, D\}$, along with a compatible realization $(u_i)_{i \notin N(T_t)}$. For simplicity, let us define

$$\mathcal{A}_t = \{u_i \cdot u_j \geq \tau_*, \forall (i,j) \in \mathcal{C}_t\}, \ \mathcal{B}_t = \{u_i \cdot u_j < \tau_*, \forall (i,j) \in \mathcal{D}_t\}, \ \mathcal{U}_t = \{u_i \cdot u_j \geq -\gamma, \forall i \in N(T_t), j \in [n]\}.$$

Heuristically, the primary influence from the conditioning arises from $\mathcal{A}_t$, as $\mathcal{B}_t$ and $\mathcal{U}_t$ are typical events under $\widetilde{\mu}$. This intuition will be formalized in Lemma 4.10. We now focus on the measure $\widehat{\mu}_t[\cdot] = \widetilde{\mu}[\cdot \mid \mathcal{A}_t]$, which constitutes the main part of our analysis. Let $\widehat{\nu}_t$ be its marginal on u_{i_t}, and our goal is to establish an appropriate control on $F(\widehat{\nu}_t)$.

Consider $F_t \subset T_t$, the subgraph generated by edges $(i,j) \in \mathcal{C}_t$. It is clear that F_t is a forest, and we denote by $T'_t \subset F_t$ the component containing i_t. Crucially, $\widehat{\mu}_t$ can be viewed as the uniform measure on the solution space of a constraint satisfaction problem (CSP) on $(\mathbb{S}^d)^{N(T_t)}$, defined by the constraints $u_i \cdot u_j \geq \tau_*, \forall (i,j) \in E(F_t)$. The factor graph of this CSP is F_t, allowing $\widehat{\mu}_t$ to decompose into products of marginals over the components of F_t. Consequently, $\widehat{\nu}_t$ is simply the marginal of the uniform measure over the solution space of the CSP defined by $u_i \cdot u_j \geq \tau_*, \forall (i,j) \in E(T'_t)$. Moreover, as T'_t is a tree, $\widehat{\nu}_t$ can be explicitly computed using the belief propagation algorithm. Given that T'_t has a maximum depth of 2, we can derive the explicit expression for belief propagation readily.

For clarity, we redefine our notations as follows: let the root of T'_t be denoted as $\emptyset$, with $\emptyset$ having children $1, \dots, D$. Each child i has its own children $(i,1), \dots, (i, D_i)$ for $1 \leq i \leq D$. Under this context, the vectors $u_{i,j} \in \mathbb{S}^d, 1 \leq i \leq D, 1 \leq j \leq D_i$ are fixed, ensuring that $|u_{i,j} \cdot u_{i',j'}| \leq \gamma$ for all distinct pairs (i,j) and (i',j'), while $u_\emptyset, u_i, 1 \leq i \leq D$ are flexible. Under these new notations, our goal can be restated as follows.

Proposition 4.8. *Given that*

$$\widehat{\mu}[\cdot] = \widetilde{\mu}[\cdot \mid u_\emptyset \cdot u_i \geq \tau_*, u_i \cdot u_{(i,j)} \geq \tau_*, 1 \leq i \leq D, 1 \leq j \leq D_i]$$

where $\widetilde{\mu}$ is the product measure of uniform distributions of $u_\emptyset, u_i, 1 \leq i \leq D$, then the probability density $\widehat{\nu}_\emptyset(x)$ of the marginal of $u_\emptyset$ under $\widehat{\mu}$ satisfies that $F(\widehat{\nu}_t) \leq n^{o(1)} p$.

We now run belief propagation algorithm on T'_t to compute $\widehat{\nu}_t$ as follows:

Step 1 For $1 \le i \le D$ and $1 \le j \le D_i$, define $\nu_{(i,j)\to i}(x) = p^{-1}\mathbf{1}(x \cdot u_{i,j} \ge \tau_*), \forall x \in \mathbb{S}^d$.

Step 2 For $1 \le i \le D$, let $\widehat{\nu}_i(x) \propto \prod_{j=1}^{D_i} \nu_{(i,j)\to i}(x), \forall x \in \mathbb{S}^d$, such that $\widetilde{\nu}_i$ is a probability density on S^d.

Step 3 For $1 \le i \le D$, let $\nu_{i\to\emptyset}(x) = p^{-1}\mathbb{P}_{y\sim\widehat{\nu}_i}[x \cdot y \ge \tau_*], \forall x \in \mathbb{S}^d$.

Step 4 Let $\widehat{\nu}_\emptyset(x) \propto \prod_{i=1}^{D} \nu_{i\to\emptyset}(x), \forall x \in \mathbb{S}^d$, such that ν is a probability density on $\mathbb{S}^d$.

Then it follows that the output of Step 4 is just the probability density of the marginal distribution of $u_\emptyset$, see [6, Section 7] for more details.

We now look more carefully into each step of the belief propagation. Step 1 is straightforward. Plugging the expression in Step 1 into Step 2, we see that for each $1 \le i \le D$,

$$\widehat{\nu}_i(x) \propto \mathbf{1}(x \cdot u_{i,j} \ge \tau_*, \forall 1 \le j \le D_i)\,,$$

and so

$$\widehat{\nu}_i(x) = \frac{\mathbf{1}(x \cdot u_{i,j} \ge \tau_*, \forall 1 \le j \le D_i)}{\mathbb{P}_{u\sim\rho}[u \cdot u_{i,j} \ge \tau_*, 1 \le j \le D_i]}\,.$$

Therefore, in Step 3 we have that for each $1 \le i \le D$,

$$\nu_{i\to\emptyset}(x) = \mathbb{P}_{y\sim\widehat{\nu}_i}[x \cdot y \ge \tau_*] = \frac{\mathbb{P}_{u\sim\rho}[x \cdot u \ge \tau_*, u \cdot u_{i,j} \ge \tau_*, \forall 1 \le j \le D_i]}{p \cdot \mathbb{P}_{u\sim\rho}[u \cdot u_{i,j} \ge \tau_*, \forall 1 \le j \le D_i]}\,. \tag{4.8}$$

And finally, we have

$$\widehat{\nu}_\emptyset(x) = \frac{\prod_{i=1}^{D} \nu_{i\to\emptyset}(x)}{\int_{\mathbb{S}^d} \prod_{i=1}^{D} \nu_{i\to\emptyset}(y)\,\mathrm{d}\rho(y)}\,. \tag{4.9}$$

We have the following simple but useful estimate.

Lemma 4.9. *For any $K = O(1)$ and $v_1, \dots, v_k \in \mathbb{S}^d$ with $|v_i \cdot v_j| \le o(1), \forall 1 \le i \ne j \le K$, it holds that*

$$\mathbb{P}_{u\sim\rho}[u \cdot v_j \ge \tau_*, \forall 1 \le j \le K] = n^{o(1)} p^K\,.$$

Proof. By rotation symmetry and the condition, we may assume that

$$\begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_K \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & \cdots & 0 \\ o(1) & 1-o(1) & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots & \cdots & 0 \\ o(1) & o(1) & o(1) & \cdots & 1-o(1) & \cdots & 0 \end{pmatrix}$$

First, recall $\gamma = (\log n/d)^{1/4} = o(1)$, from a simple union bound it holds that

$$\mathbb{P}_{u\sim\rho}[u \cdot u_j \ge \tau_*, \forall 1 \le j \le K] = \mathbb{P}_{u\sim\rho}[\tau_* \le u \cdot v_j \le \gamma, \forall 1 \le j \le K] + n^{-\omega(1)}\,.$$

Now let $u = (u_1, \dots, u_{d+1}) \sim \rho$, we note that conditioned on the first m coordinates, $(u_{m+1}, \dots, u_{d+1})$ is just uniformly chosen from the $(d-m)$-dimensional sphere with radius $(1 - u_1^2 - \cdots - u_m^2)^{1/2}$. With the observation in mind, it can be easily deduced from 2.1 that for each $0 \le m \le K-1$, we have

$$\mathbb{P}_{u\sim\rho}[\tau_* \le u \cdot v_{m+1} \le \gamma \mid \tau_* \le u \cdot v_i \le \gamma, i \le m] = n^{o(1)} p\,,$$

and hence the result follows from the multiplicative rule. $\square$

Now we can easily prove Proposition 4.8.

Proof of Proposition 4.8. From our choice of $u_{i,j}, 1 \le i \le D, 1 \le j \le D_i$, we have that $|u_{i,j} \cdot u_{i,j'}| \le \gamma = o(1)$ for all $i, j \ne j'$. Therefore, we have the denominator in (4.8) is $n^{o(1)} p^{D_i+1}$. In addition, we define

$$B = \{x \in \mathbb{S}^d : |x \cdot u_{i,j}| \le \gamma, \forall 1 \le i \le D, 1 \le j \le D_i\},$$

then it is clear that $\rho(B) = 1 - n^{-\omega(1)}$ by a union bound. Be lemma 4.9 we have for $x \in B$, the numerator in (4.8) is also $n^{o(1)} p^{D_i+1}$. Hence for $x \in B$ it holds $\nu_{i \to \emptyset}(x) = n^{o(1)}$. For $x \in \mathbb{S}^d \setminus B$, it trivially holds that $\nu_{i \to \emptyset}(x) \le p^{-1} \le n$. Given this, it follows that

$$F(\widehat{\nu}_\emptyset) = \sup_{x \in \mathbb{S}^d} \int_{y \cdot x \ge \tau_*} \widetilde{\nu}_\emptyset(y) \, \mathrm{d}\rho(y) \le p \cdot n^{o(1)} + n^{-\omega(1)} \cdot n = p n^{o(1)} . \qquad \square$$

Finally, we relate $\widehat{\mu}_t = \widetilde{\mu}[\cdot \mid \mathcal{A}_t]$ with $\widetilde{\mu}_t = \widetilde{\mu}[\cdot \mid \mathcal{A}_t, \mathcal{B}_t, \mathcal{U}_t]$, the probability measure we really care.

Lemma 4.10. *It holds that $\widetilde{\mu}[\mathcal{B}_t, \mathcal{U}_t \mid \mathcal{A}_t] \ge 1 - o(1)$, and hence $\widetilde{\mu}_t[\cdot] \le (1 + o(1))\widehat{\mu}_t[\cdot]$.*

Proof. We only sketch the proof of the first claim. It suffices to show that $\widetilde{\mu}[\mathcal{B}_t^c \mid \mathcal{A}_t]$ and $\widetilde{\mu}[\mathcal{U}_t^c \mid \mathcal{A}_t]$ are both $o(1)$.

For the former one, we show that for any $(i, j) \in \mathcal{D}_t$, $\widetilde{\mu}[u_i \cdot u_j \ge \tau_* \mid \mathcal{A}_t]$. If either of i, j equals to i_t, then the result just follows from $F(\widehat{\nu}_t) \le n^{o(1)} p$ as shown in Proposition 4.8. Otherwise, one of u_i, u_j is fixed, say u_j. It can be shown by analyzing a similar belief propagation that this reduces to some probability of the form

$$\frac{\mathbb{P}[u_i \cdot u_j \ge \tau_*, u_i \cdot u_{j_t} \ge \tau_*, 1 \le t \le K]}{\mathbb{P}[u_i \cdot u_{j_t} \ge \tau_*, \forall 1 \le t \le K]},$$

which can be done via Lemma 4.9.

For the latter one, we note that $\widetilde{\mu}[\mathcal{U}_t^c \mid \mathcal{A}_t] \le \widetilde{\mu}[\mathcal{U}_t^c] / \widetilde{\mu}[\mathcal{A}_t]$, and $\widetilde{\mu}[\mathcal{U}_t^c] \le n^{-\omega(1)}$ by a union bound. It can be shown via Lemma 4.9 that $\widetilde{\mu}[\mathcal{A}_t] \ge n^{O(1)}$, and thus the result follows.

Now given with the first claim, we have

$$\widetilde{\mu}_t[\cdot] = \frac{\widetilde{\mu}[\cdot, \mathcal{A}_t, \mathcal{B}_t, \mathcal{U}_t]}{\widetilde{\mu}[\mathcal{A}_t, \mathcal{B}_t, \mathcal{U}_t]} \le \frac{\widetilde{\mu}[\cdot, \mathcal{A}_t]}{\widetilde{\mu}[\mathcal{A}_t] \widetilde{\mu}[\mathcal{B}_t, \mathcal{U}_t \mid \mathcal{A}_t]} = \frac{\widetilde{\mu}[\cdot \mid \mathcal{A}_t]}{\widetilde{\mu}[\mathcal{B}_t, \mathcal{U}_t \mid \mathcal{A}_t]} = (1 + o(1))\widehat{\mu}[\cdot] . \qquad \square$$

Proof of Proposition 4.7. It follows from the lemma that $F(\widetilde{\nu}_t) \le (1 + o(1)) F(\widehat{\nu}_t)$, which is $n^{o(1)} p$ as shown in Proposition 4.8. This proves Proposition 4.7 and thus finishes the proof of (1.4). $\square$

5 Further discussions

5.1 The power of rotation is futile

In this section, we explore a potential strategy to enhance the lower bound established in Section 3. This approach manage to leverage an underutilized aspect in our previous proofs: the power of rotations on $\mathbb{S}^d$. Specifically, we investigate whether the additional degrees of freedom provided by rotations in $\mathrm{O}(d)$ (the orthogonal group in d dimensions, comprising isometries on $\mathbb{S}^d$) can significantly improve our lower bound. The mind map is once it can be shown that for typical graphs G_1 and G_2 generated from vertices $u_1, \ldots, u_n \in \mathbb{S}^d$ and $v_1, \ldots, v_n \in \mathbb{S}^d$ respectively, there exists a transformation $T \in \mathrm{O}(d)$ such that for a majority of indices $i \in [n]$, the distance $\mathrm{d}(Tu_i, v_{\pi(i)})$ is less than a certain radius r, thereby implying $\Lambda(G_1, G_2) = 1 - o(1)$.

However, this section proposes that the inclusion of rotational transformations does not necessarily transcend the conjectured threshold d_0. Suppose we have $d_0 \ll d \ll \log n$. In line with the arguments presented in the proof of (1.3) in Section 3, establishing $\Lambda(G_1, G_2) = 1 - o(1)$ with high probability requires demonstrating that, typically, for $u_1, \dots, u_n$ and $v_1, \dots, v_n$ uniformly and independently selected from $\mathbb{S}^d$, there exists a transformation $T \in \mathrm{O}(d)$ and a permutation $\pi \in \mathrm{S}_n$ such that

$$\mathrm{d}(Tu_i, v_{\pi(i)}) = o\left(\frac{r_\theta}{d}\right), \text{ for } n - o(n) \text{ indices } i \in [n]. \tag{5.1}$$

However, we will demonstrate that this condition typically does not hold, suggesting that our previous approach cannot be extended to dimensions exceeding d_0.

Let ε be such that $\mathbb{P}[\mathrm{d}(u,v) \le \varepsilon] = n^{-2}$, then it follows that $\varepsilon = n^{-(2+o(1))/d}$ by similar arguments as in Lemma 3.1, and so $\varepsilon \gg r_\theta/d$ by our choice of d. Moreover, we take a minimal $\varepsilon/2$-net $\mathcal{N}$ of $\mathrm{O}(d)$ with respect to the operator norm $\|\cdot\|_{\mathrm{op}}$, then from [8, Propsition 6] we have $|\mathcal{N}| \le (c_0\varepsilon)^{-d^2/2}$ for some universal constant $c_0 > 0$. If there exists $T \in \mathrm{O}(d)$ and $\pi \in \mathrm{S}_n$ such that (5.1) holds, then there exists $T_0 \in \mathcal{N}$, such that $\mathrm{d}(T_0 u_i, v_{\pi(i)}) \le \varepsilon$ for at least half of $i \in [n]$. Denote $B_{T_0,i} = \{v \in \mathbb{S}^d, \mathrm{d}(T_0 u_i, v) \le \varepsilon\}$, then the aforementioned event implies that $B_{T_0,i} \cap \{v_1, \dots, v_n\} \neq \emptyset$ for at least half of $i \in [n]$. It can be shown by a union bound that with high probability for $u_1, \dots, u_n$, $\mathrm{d}(u_i, u_j) > 2\varepsilon, \forall i \neq j \in [n]$, and so $B_{T_0,i}, i = 1, \dots, n$ are pairwise disjoint. We fix any choice of $u_1, \dots, u_n$ satisfying this condition, then for each fixed $T_0 \in \mathcal{N}$, the events $\{B_{T_0,i} \cap \{v_1, \dots, v_n\} \neq \emptyset\}$ are negative-correlated with each other, and so we have

$$\begin{aligned}
&\mathbb{P}\big[\exists T \in \mathrm{O}(d), \pi \in \mathrm{S}_n \text{ s.t. } \mathrm{d}(T_0 u_i, v_{\pi(i)}) \le r_\theta/d \text{ for at least half of } i \in [n]\big] \\
\le\ &\sum_{T_0 \in \mathcal{N}} \mathbb{P}\big[B_{T_0,i} \cap \{v_1, \dots, v_n\} \neq \emptyset \text{ for at least half of } i \in [n]\big] \\
\le\ &(c_0\varepsilon)^{-d^2/2} \cdot \binom{n}{n/2} \cdot (2n^{-2})^{n/2} = n^{O(\log n) - \Omega(n)} = o(1)\,,
\end{aligned}$$

where in the last inequality, we first employ a union bound on the half of $i \in [n]$ satisfying $B_{T_0,i} \cap \{v_1, \dots, v_n\} \neq \emptyset$, then use the fact that these events are negative-correlated together with a simple estimate $\mathbb{P}[B_{T_0,i} \cap \{v_1, \dots, v_n\} \neq \emptyset] \le 2n^{-2}$ followed by Poisson approximation. This proves that the additional power of rotations does not essentially improve the previous lower bound.

5.2 Breaking the $\log n$ upper bound

In this final section, we explore potential enhancements to our proof for the upper bound, specifically aiming to achieve a bound down to $\log n$. A crucial element in our analysis in Section 4 is the application of a belief propagation algorithm on a tree with depth 2, as detailed in Proposition 4.8. The justification for limiting the depth to 2 stems from the observation that the connecting threshold $\tau_* = o(1)$ provided with the dimension $d \gg \log n$. In this case, a random walk on $\mathbb{S}^d$ with transition kernel $P(x,y) = p_*^{-1}\mathbf{1}(x \cdot y \ge \tau_*)$ exhibits nice mixing properties within just 2 or 3 steps, which partly explain why restricting to a 3-neighborhood suffices for our purpose.

Conversely, for smaller values of d, where $\tau_* = \Omega(1)$ or even approaches 1, it becomes imperative to consider trees of greater depth for analogous results to Proposition 4.8. Intuitively, the optimal depth D_* should be neither too small, such that given the positions of all descendants up to D_*-generations, the posterior distribution of the root approximates the uniform distribution, nor too large, such that the D_*-neighborhood of a sparse random graph is typically a tree. For the latter lower bound on D_*, as a minimal requirement, it must at least match the mixing time of the corresponding random walk on $\mathbb{S}^d$. This aspect of the random walk mixing time is explored in [7], and it is reassuring to note that the lower bound for D_* remains below the threshold at which the local neighborhood in H begins to exhibit loop-like structures, provided $d \ge d_0$. In this sense, it is promising that our proof scheme can be extended to prove Conjecture 1.4. Nevertheless, handling a large D_* introduces complexity in explicitly writing out each step of the belief propagation algorithm, necessitating innovative approaches for its analysis. We leave this to future works.

Acknowledgments

The author would like to thank Sidhanth Mohanty for introducing the problem, and thank Byron Chin, Nike Sun and Sidhanth Mohanty for helpful discussions.

References

[1] M. Brennan G. Bresler and D. Nagaraj, Phase Transitions for Detecting Latent Geometry in Random Graphs, *Probability Theory and Related Fields*, Vol. 178, 2020, pp. 1215-1289.

[2] S. Bubeck, J. Ding, R. Eldan and M. Rácz, Testing for high-dimensional geometry in random graphs, *Random Structures and Algorithms*, vol. 49, no. 3, 2016, pp. 503-532.

[3] J. Ding, H. Du and S. Gong, A polynomial-time approximation scheme for the maximal overlap of two independent Erdős-Rényi graphs, Preprint, arXiv:2210.07823.

[4] H. Du, S. Gong and R. Huang, The Algorithmic Phase Transition of Random Graph Alignment Problem, Preprint, arXiv:2307.06590.

[5] A. Goel, S. Rai, and B. Krishnamachari, Monotone properties of random geometric graphs have sharp thresholds. *Annals of Applied Probability*, vol. 15, no. 4, 2005, pp. 2535–2552.

[6] S. Liu, S. Mohanty, T. Schramm, and E. Yang, Testing thresholds for high-dimensional sparse random geometric graphs, *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, June 2022, pp. 672–677.

[7] S. Li and T. Schramm, Spectral clustering in the Gaussian mixture block model. Preprint, arXiv:2305.00979.

[8] S. J. Szarek, Nets of Grassmann manifold and orthogonal group, *Proceeding of research workshop on Banach Space Theory*, 1981. pp. 169-186.