

Testing for high-dimensional geometry in random graphs

Sébastien Bubeck^{*} Jian Ding[†] Ronen Eldan[‡] Miklós Z. Rácz[§]

November 20, 2014

Abstract

We study the problem of detecting the presence of an underlying high-dimensional geometric structure in a random graph. Under the null hypothesis, the observed graph is a realization of an Erdős-Rényi random graph $G(n, p)$. Under the alternative, the graph is generated from the $G(n, p, d)$ model, where each vertex corresponds to a latent independent random vector uniformly distributed on the sphere \mathbb{S}^{d-1} , and two vertices are connected if the corresponding latent vectors are close enough. In the dense regime (i.e., p is a constant), we propose a near-optimal and computationally efficient testing procedure based on a new quantity which we call signed triangles. The proof of the detection lower bound is based on a new bound on the total variation distance between a Wishart matrix and an appropriately normalized GOE matrix. In the sparse regime, we make a conjecture for the optimal detection boundary. We conclude the paper with some preliminary steps on the problem of estimating the dimension in $G(n, p, d)$.

1 Introduction

Extracting information from large graphs has become an important statistical problem, as network data is now common in various fields. Whether one talks about social networks, gene networks, or (biological) neural networks, in all cases there is a rapidly growing industry dedicated to extracting knowledge from these graphs. A particularly important task in this spirit is to learn a useful *representation* of the vertices, that is, a mapping from the vertices to some metric space (usually \mathbb{R}^d). In this paper, we take a step back and study the hypothesis testing problem that underlies these investigations: given a large graph, one wants to tell if the observed connections result from a latent geometrical structure in the vertices, or if they are purely random.

The null hypothesis that we consider is that the observed graph G on n vertices has been generated by the standard Erdős-Rényi random graph $G(n, p)$ [16], where each edge appears independently with probability p :

$$H_0 : G \sim G(n, p).$$

On the other hand, for the alternative, we consider the simplest model of a random geometric graph—recall that a geometric graph is such that each vertex is labeled with a point in some metric space, and an edge is present between two vertices if the distance between the corresponding labels is smaller than some prespecified threshold. We focus on the case where the underlying metric space is the Euclidean sphere $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\|_2 = 1\}$, and the latent labels are i.i.d. uniform

^{*}Microsoft Research and Princeton University; sebubeck@microsoft.com.

[†]University of Chicago; jianding@galton.uchicago.edu.

[‡]University of Washington; roneneldan@gmail.com.

[§]University of California, Berkeley; racz@stat.berkeley.edu.

random vectors in \mathbb{S}^{d-1} . We denote this model by $G(n, p, d)$, where p is the probability of an edge between two vertices (p determines the threshold distance for connection), and thus

$$H_1 : G \sim G(n, p, d).$$

Slightly more formally, $G(n, p, d)$ is defined as follows. Let X_1, \dots, X_n be independent random vectors, uniformly distributed on \mathbb{S}^{d-1} . In $G(n, p, d)$, distinct vertices $i \in [n]$ and $j \in [n]$ are connected by an edge if and only if $\langle X_i, X_j \rangle \geq t_{p,d}$, where the threshold value $t_{p,d} \in [-1, 1]$ is such that $\mathbb{P}(\langle X_1, X_2 \rangle \geq t_{p,d}) = p$.

In contrast to most previous works (see Section 1.1), the focus of this paper is on the *high-dimensional* situation, where d can scale as a function of n . This point of view is in line with recent advances in all areas of applied mathematics, and in particular statistics and learning theory, where high-dimensional feature spaces are becoming the new norm. We also consider both the dense regime, where p is a constant independent of n , and the sparse regime, where $p = c/n$ for some constant c .

A natural test to uncover geometric structure is to count the number of triangles in G . Indeed, in a purely random scenario, vertex u being connected to both v and w says nothing about whether v and w are connected. On the other hand, in a geometric setting this would imply that v and w are close to each other due to the triangle inequality, thus increasing the probability of a connection between them. This, in turn, implies that the expected number of triangles is larger in the geometric setting, given the same edge density. More broadly, the number of triangles is a basic statistic in the analysis of networks, as it gives an indication of the dependencies in the networks, see, e.g., [18] and the references therein.

One of the main results of this paper is that in the dense regime, there are much more powerful statistics than triangles to distinguish between $G(n, p)$ and $G(n, p, d)$. In particular, we propose a near-optimal statistic based on a new quantity which we call *signed triangles*. To put it succinctly, if A denotes the adjacency matrix of a graph G , then the total number of triangles is equal to $\text{Tr}(A^3)$, while the total “number” of signed triangles is defined as $\text{Tr}((A - p)^3)$. The key insight motivating this definition is that the variance of signed triangles is dramatically smaller than the variance of triangles, due to the cancellations introduced by the centering of the adjacency matrix. While being elementary, this idea of using cancellations of various patterns seems to be new in the network analysis literature. It could have applications beyond the problem studied in this paper, in particular because it is quite different from other popular statistics, such as those coming from the method of moments [9] (where patterns are always counted positively), or those with the flavor of the clustering coefficient [27] (where ratios of counts are considered). Indeed, the key aspect of signed triangles is that different patterns are reweighted with positive and negative weights: triangles and induced single edges are counted positively, while induced wedges and independent sets on three vertices are counted negatively.

1.1 Related work

In the social networks literature the latent metric space is referred to as the *social space*. Starting with [19], there have been numerous works on estimating positions in the social space for individuals in a social network. Various models are considered in this literature, and interestingly some of them are very closely related to $G(n, p, d)$, see [19, Section 2.2]. Most papers focus on various approximations to the maximum likelihood estimator, and some theoretical results have been obtained with other methods such as counting the number of common neighbors, see [24, 2]. In the

present paper we settle for a less ambitious goal, as we are not learning the representation but simply testing for the presence of a meaningful representation. On the other hand, we obtain much more precise results, such as an almost tight characterization of couples (n, d) for which testing between $G(n, p)$ and $G(n, p, d)$ is possible in the dense regime. Most importantly, this provides, to the best of our knowledge, the first result for the high-dimensional setting, as all previous works on social space inference were in the low-dimensional regime (i.e., d is fixed and $n \rightarrow \infty$).

In probability theory, models such as $G(n, p, d)$ have been studied for a long time in the low-dimensional regime, see, e.g., [23]. The high-dimensional setting was first investigated recently in [12]. In this paper it was observed that with n fixed and $d \rightarrow \infty$, $G(n, p, d)$ converges in total variation to $G(n, p)$. In other words, if the dimension d is very large compared to n , then one cannot distinguish between $G(n, p)$ and $G(n, p, d)$. Our main result is that, in fact, in the dense regime, there is a phase transition when d is of order n^3 , with the total variation distance going from 1 to 0. In spite of previous works, the high-dimensional $G(n, p, d)$ remains mysterious in many ways. Essentially, in the dense regime, the only graph parameter which is well understood is the clique number, due to the results of [12, 5], while in the sparse case basically nothing is known. One of the technical contributions of the present paper is to compute rather precisely the probability of a triangle in the sparse case.

This paper can be seen as part of a series of recent papers studying hypothesis testing in random graph models [6, 26, 11]. For instance, Arias-Castro and Verzelen consider the problem of testing for *community structure*, while we test for *geometric structure*. More precisely, their null is identical to ours, that is $G \sim G(n, p)$, while for the alternative they consider the model $G(n, p, k, q)$ which differs from $G(n, p)$ by having a random subset of k vertices for which edges appear between them with probability q . This problem is closely related to community detection in stochastic block models, a problem which has recently attracted a lot of attention, see, e.g., [22, 1] and the references therein. Interestingly, when testing for community structure, one of the main obstacles is computational rather than information-theoretic. For example, it is obvious that one can distinguish between $G(n, 1/2)$ and $G(n, 1/2, k, 1)$ as long as $k \gg \log(n)$, though when $k = o(\sqrt{n})$, no polynomial test is known for this problem (referred to as the *planted clique problem*), see, e.g., [3, 8]. On the contrary, in the context of testing for geometrical structure, we show that polynomial time methods are near-optimal, at least for the dense regime, since one can efficiently compute signed triangles. This is perhaps surprising, as the worst-case version of our problem, namely recognizing if a graph can be realized as a geometric graph, is known to be NP-hard [10]. Finally, we also note that our new signed triangles statistic is closely related to the tensor introduced in [17] for the planted clique problem. While [17] computes the spectral norm of this tensor, here we simply sum its entries (in the case $p = 1/2$).

1.2 Contributions and content of the paper

The main objective of the paper is to identify the boundary of testability between $G(n, p)$ and $G(n, p, d)$, in both the dense and the sparse regimes. To put it differently, we are interested in studying the total variation distance between these two models, denoted by $\text{TV}(G(n, p), G(n, p, d))$. Recall that in [12] it was proved that, if n is fixed and $d \rightarrow \infty$, then

$$\text{TV}(G(n, p), G(n, p, d)) \rightarrow 0.^1$$

Given this result, our question of interest becomes the following: how large can d be, as a function of n , so that $\text{TV}(G(n, p), G(n, p, d))$ remains bounded away from 0 (or even becomes close to 1)?

¹More precisely, they show that this convergence happens when d is exponential in n .

The core of our contribution to this question is summarized by the following theorem.

Theorem 1 (a) *Let $p \in (0, 1)$ be fixed, and assume that $d/n^3 \rightarrow 0$. Then*

$$\text{TV}(G(n, p), G(n, p, d)) \rightarrow 1. \quad (1)$$

(b) *Let $c > 0$ be fixed, and assume that $d/\log^3(n) \rightarrow 0$. Then*

$$\text{TV}\left(G\left(n, \frac{c}{n}\right), G\left(n, \frac{c}{n}, d\right)\right) \rightarrow 1. \quad (2)$$

(c) *Furthermore, if $d/n^3 \rightarrow \infty$, then*

$$\sup_{p \in [0, 1]} \text{TV}(G(n, p), G(n, p, d)) \rightarrow 0. \quad (3)$$

The results of Theorem 1(a) and 1(c) tightly characterize the dense regime. We conjecture that the result of Theorem 1(b) is tight for the sparse regime:

Conjecture 1 *Let $c > 0$ be fixed, and assume that $d/\log^3(n) \rightarrow \infty$. Then*

$$\text{TV}\left(G\left(n, \frac{c}{n}\right), G\left(n, \frac{c}{n}, d\right)\right) \rightarrow 0.$$

In the following we outline the methods we use to prove Theorem 1.

We first discuss the proof of part (a) of Theorem 1, where our main methodological contribution lies. As pointed out above, a natural test to consider is the number of triangles:

$$T(G) := \sum_{\{i, j, k\} \in \binom{[n]}{3}} A_{i, j} A_{i, k} A_{j, k}, \quad (4)$$

where A denotes the adjacency matrix of the graph G with vertex set $[n]$, and so $A_{i, j} = 1$ if vertices i and j are connected by an edge, and $A_{i, j} = 0$ otherwise. An elementary calculation shows that $\mathbb{E} T(G(n, p)) = \binom{n}{3} p^3$, while $\text{Var}(T(G(n, p)))$ is of order n^4 . A key calculation of the paper is to show that $\mathbb{E} T(G(n, p, d)) \geq \binom{n}{3} p^3 (1 + C_p/\sqrt{d})$ for some constant C_p that depends only on p , see Lemma 1. Intuitively, this shows that triangles have some power as long as

$$|\mathbb{E} T(G(n, p)) - \mathbb{E} T(G(n, p, d))| \gg \sqrt{\text{Var}(T(G(n, p)))},$$

which is equivalent to $d \ll n^2$. One of the main contributions of the paper is to introduce a statistic that is asymptotically powerful as long as $d \ll n^3$, which we refer to as the number of *signed triangles*:

$$\tau(G) := \sum_{\{i, j, k\} \in \binom{[n]}{3}} (A_{i, j} - p)(A_{i, k} - p)(A_{j, k} - p). \quad (5)$$

The key point of signed triangles is the reduction of variance. Namely, $\text{Var}(\tau(G(n, p)))$ is of order n^3 , instead of n^4 for triangles. We study τ in detail in Section 3, where we prove the following result (which implies Theorem 1(a)).

Theorem 2 *Let $p \in (0, 1)$ be fixed, and assume that $d/n^3 \rightarrow 0$. Then*

$$\text{TV}(\tau(G(n, p)), \tau(G(n, p, d))) \rightarrow 1.$$

In the sparse regime we analyze the triangle statistic and prove the following theorem, which implies part (b) of Theorem 1.

Theorem 3 *Let $c > 0$. There exists a constant $C > 0$ such that, if $d/\log^3(n) \rightarrow 0$, then*

$$\text{TV} \left(T \left(G \left(n, \frac{c}{n} \right) \right), T \left(G \left(n, \frac{c}{n}, d \right) \right) \right) \rightarrow 1.$$

We believe that the upper bound $\log^3(n)$ cannot be significantly improved, as stated above in Conjecture 1. The main reason for this conjecture is that, when $d \gg \log^3(n)$, $G(n, c/n)$ and $G(n, c/n, d)$ seem to be locally equivalent (in particular, they both have the same Poisson number of triangles). Thus the only way to distinguish between them would be to find an emergent global property which is significantly different under the two models, but this seems unlikely to exist.

The final result of Theorem 1, part (c), complements part (a) by giving a matching lower bound. This bound is also valid for the sparse regime, though in this case we believe that it is not tight. The main idea behind the proof of this result is to view the random geometric graph $G(n, p, d)$ as a function of an $n \times n$ Wishart matrix with d degrees of freedom, denoted by $W(n, d)$, and to view $G(n, p)$ as a function of an $n \times n$ GOE random matrix, denoted by $M(n)$. Theorem 1(c) then follows from the following result on random matrices, stating that, if d/n^3 is very large, then the Wishart matrix has approximately the same law as an appropriately centered and scaled GOE random matrix.

Theorem 4 *Let I_n denote the $n \times n$ identity matrix. If $d/n^3 \rightarrow \infty$, then*

$$\text{TV} \left(W(n, d), \sqrt{d}M(n) + dI_n \right) \rightarrow 0.$$

The random ensembles $W(n, d)$ and $M(n)$ are defined more precisely in Section 5, where the above theorem is also proved.

Finally, in Section 6 we touch upon the following question: for which values of n and d can the dimension d be reconstructed by observing a sample of $G(n, p, d)$? We give the following bound for $p = 1/2$, which can be considered as a proof of concept.

Theorem 5 *There exists a universal constant $C > 0$, such that for all integers n and $d_1 < d_2$, one has*

$$\text{TV} (G(n, 1/2, d_1), G(n, 1/2, d_2)) \geq 1 - C \left(\frac{d_1}{n} \right)^2.$$

This bound is tight, as demonstrated by an older result of the third named author, proven in [15]. This result states that when $d \gg n$, the graphs $G(n, 1/2, d)$ and $G(n, 1/2, d + 1)$ are indistinguishable.

Theorem 6 *There exists a universal constant $C > 0$ such that for all integers $n < d$,*

$$\text{TV} (G(n, 1/2, d), G(n, 1/2, d + 1)) \leq C \sqrt{\left(\frac{d + 1}{d - n} \right)^2 - 1}.$$

2 Estimates for the number of triangles in a geometric graph

The point of this section is to give a lower bound for the expected number of triangles in the random geometric graph $G(n, p, d)$, using elementary estimates. To this end, let X_1, X_2 , and X_3 be independent uniformly distributed points in \mathbb{S}^{d-1} . Consider the event

$$E = \left\{ \langle X_1, X_2 \rangle \geq t_{p,d}, \langle X_1, X_3 \rangle \geq t_{p,d}, \langle X_2, X_3 \rangle \geq t_{p,d} \right\}$$

that the corresponding vertices form a triangle; the expected number of triangles in $G(n, p, d)$ is thus $\mathbb{E} T(G(n, p, d)) = \binom{n}{3} \mathbb{P}(E)$. Our main result of this section is the following.

Lemma 1 *There exists a universal constant $C > 0$ such that whenever $p < \frac{1}{4}$ we have that*

$$\mathbb{P}(E) \geq C^{-1} p^3 \frac{(\log \frac{1}{p})^{3/2}}{\sqrt{d}}. \quad (6)$$

Moreover, for every fixed $0 < p < 1$, there exists a constant $C_p > 0$ such that for all $d \geq C_p^{-1}$,

$$\mathbb{P}(E) \geq p^3 \left(1 + \frac{C_p}{\sqrt{d}} \right). \quad (7)$$

Before we move on to the proof, we need some preliminary results, for which we need some additional notation. Denote by $\varphi(x) = (1/\sqrt{2\pi}) e^{-x^2/2}$ the standard Gaussian density and by $\bar{\Phi}(x) = \int_x^\infty \varphi(z) dz$ the associated complementary cumulative distribution function (note that $\bar{\Phi}$ is decreasing). Moreover, define

$$f_d(x) = \frac{\Gamma(d/2)}{\Gamma((d-1)/2)\sqrt{\pi}} (1-x^2)^{(d-3)/2}, \quad x \in [-1, 1]. \quad (8)$$

A standard calculation gives that f_d is the density of a one-dimensional marginal of a uniform random point on \mathbb{S}^{d-1} (see [25, Section 2]).

Lemma 2 *There exists a universal constant $C > 0$ such that for all $0 < p \leq 1/2$ we have that*

$$\left(C^{-1} \left(\frac{1}{2} - p \right) \sqrt{\frac{\log(1/p)}{d}} \right) \wedge (1/2) \leq t_{p,d} \leq C \sqrt{\frac{\log(1/p)}{d}}. \quad (9)$$

and

$$f_{d-1}(t_{p,d}) \geq C^{-1} d p t_{p,d}. \quad (10)$$

Proof We may assume that $d \geq 4$; the statement is easily checked when $d \leq 3$. We begin by recalling the well-known inequality

$$\Gamma(x)\sqrt{x}/100 \leq \Gamma(x+1/2) \leq 2\sqrt{x}\Gamma(x), \quad \forall x > 1. \quad (11)$$

Therefore, we have for $0 < \theta < 1$ that

$$\begin{aligned} \int_\theta^1 f_d(x) dx &\leq \frac{2\sqrt{d}}{\sqrt{2\pi}} \int_\theta^1 (1-x^2)^{(d-3)/2} dx = \frac{2\sqrt{d}}{\sqrt{2\pi}} \int_0^{1-\theta} e^{\frac{d-3}{2} \log(1-(\theta+x)^2)} dx \\ &\leq \frac{2\sqrt{d}}{\sqrt{2\pi}} (1-\theta^2)^{(d-3)/2} \int_0^{1-\theta} e^{\frac{d-3}{2} (\log(1-(\theta+x)^2) - \log(1-\theta^2))} dx \\ &\leq \frac{2\sqrt{d}}{\sqrt{2\pi}} (1-\theta^2)^{(d-3)/2} \int_0^1 e^{-\frac{d-3}{2} (\theta x + x^2)} dx \\ &\leq 8 \left(1 \wedge \frac{1}{\theta\sqrt{d}} \right) (1-\theta^2)^{(d-3)/2}. \end{aligned} \quad (12)$$

Taking $\theta = t_{p,d}$ thus gives $p \leq 8(1 - t_{p,d}^2)^{(d-3)/2}$ and so $\log(p/8) \leq -\frac{d-3}{2}t_{p,d}^2$, which implies the upper bound in (9). Moreover, equation (12) also gives

$$p \leq \frac{8}{t_{p,d}\sqrt{d}}(1 - t_{p,d}^2)^{(d-3)/2} \stackrel{(11)}{\leq} \frac{8C}{dt_{p,d}}f_d(t_{p,d})$$

for a universal constant $C > 0$, which proves (10).

It remains to prove the lower bound in (9). First assume $10^{-6} \leq p \leq 1/2$. For all $0 \leq x \leq 1$ we see that $f_d(x) \leq 2\sqrt{d}$, and so for all $0 \leq \theta \leq 1$ we have that $\int_0^\theta f_d(x)dx \leq 2\theta\sqrt{d}$. Since $\int_0^{t_{p,d}} f_d(x)dx = 1/2 - p$, we get that $t_{p,d} \geq \frac{1/2-p}{2\sqrt{d}}$ as required.

Next, we consider the case for $0 < p \leq 10^{-6}$. In this case it suffices to assume that d is larger than an absolute constant so that the left hand side of (9) is at most $1/2$. By (11), we see that for $0 < \theta \leq 1/2$ we have that

$$\begin{aligned} \int_\theta^1 f_d(x)dx &\geq \frac{\sqrt{d}}{400} \int_\theta^1 (1 - x^2)^{(d-3)/2} dx \\ &\geq \frac{\sqrt{d}}{400} \int_\theta^{\theta + \frac{1}{\sqrt{d}} \wedge \frac{1}{\theta d}} e^{-x^2(d-3)} dx \geq 10^{-4} \left(1 \wedge \frac{1}{\theta\sqrt{d}}\right) (1 - \theta^2)^{(d-3)/2}. \end{aligned} \quad (13)$$

Since $\int_{t_{p,d}}^1 f_d(x)dx = p$, the preceding inequality yields that $t_{p,d} \geq (c\sqrt{\log(1/p)/d}) \wedge (1/2)$ for some absolute constant $c > 0$. This completes the proof. \blacksquare

We are finally ready to prove the main estimate of this section.

Proof of Lemma 1 First, assume that $d \leq \frac{1}{4} \log(1/p)$. Denote by $d(\cdot, \cdot)$ the geodesic distance on \mathbb{S}^{d-1} and set $g(\theta) := \mathbb{P}(d(X_1, X_2) < \theta)$. A well known calculation shows that

$$g(\theta) = \frac{(d-1)\pi^{(d-1)/2}}{\Gamma(\frac{d+1}{2})} \int_0^\theta \sin(x)^{d-2} dx.$$

By the concavity of $\sin(x)$ on the interval $[0, \pi]$, we have that

$$g(\theta/2) \geq 2^{-d}g(\theta), \quad \forall \theta < \pi. \quad (14)$$

Next observe that one has, by the triangle inequality,

$$\begin{aligned} E &= \{d(X_1, X_2) < \arccos(t_{p,d}), d(X_1, X_3) < \arccos(t_{p,d}), d(X_2, X_3) < \arccos(t_{p,d})\} \\ &\supseteq \{d(X_1, X_2) < \frac{1}{2} \arccos(t_{p,d}), d(X_1, X_3) < \frac{1}{2} \arccos(t_{p,d})\}. \end{aligned}$$

Since $d(X_1, X_2)$ and $d(X_1, X_3)$ are independent, we get that

$$\mathbb{P}(E) \geq g\left(\frac{1}{2} \arccos(t_{p,d})\right)^2 \geq g(\arccos(t_{p,d}))^2 2^{-2d} = p^2 2^{-2d} \geq p^2 2^{-\frac{1}{2} \log(1/p)} \geq p^3 \left(1 + c \left(\log \frac{1}{p}\right)^{3/2}\right)$$

for a universal constant $c > 0$, yielding the desired estimate.

From this point on, we may therefore assume that $d \geq \frac{1}{4} \log(1/p)$. We first prove the lemma under the assumption that $p < \frac{1}{2}$. The proof of the case $p \geq \frac{1}{2}$ is similar, as explained below.

Let Z be the first coordinate of a uniform point in \mathbb{S}^{d-1} , hence a random variable whose probability density function is $f_d(\cdot)$. Denote by $E_{i,j}$ the event $\{\langle X_i, X_j \rangle \geq t_{p,d}\}$ and by $E_{i,j}(x)$ the event $\{\langle X_i, X_j \rangle = x\}$. In what follows, we occasionally allow ourselves to condition on the zero probability event $E_{i,j}(x)$. This should be understood as conditioning on $\{\langle X_i, X_j \rangle \in [x - \varepsilon, x + \varepsilon]\}$ and then taking $\varepsilon \rightarrow 0$, which is well-defined thanks to a simple continuity argument.

Note that $\mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(x))$ is an increasing function of x . Therefore,

$$\begin{aligned} & \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}) - \mathbb{P}(E_{1,3}, E_{2,3}) \\ &= \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}) - \frac{1}{2} \mathbb{P}(E_{1,3}, E_{2,3} | \langle X_1, X_2 \rangle \leq 0) - \frac{1}{2} \mathbb{P}(E_{1,3}, E_{2,3} | \langle X_1, X_2 \rangle \geq 0) \\ &\geq \frac{1}{2} \{ \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(0)) \}. \end{aligned} \quad (15)$$

Write $Z_1 = \langle X_3, X_1 \rangle$ and $Z_2 = \left\langle X_3, \frac{\text{Proj}_{X_1^\perp} X_2}{|\text{Proj}_{X_1^\perp} X_2|} \right\rangle$ where $\text{Proj}_{X_1^\perp}$ denotes the orthogonal projection onto the orthogonal complement of X_1 . We have that

$$\begin{aligned} & \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(0)) \\ &= \int_{z_1 \geq t_{p,d}} \{ \mathbb{P}(E_{2,3} | Z_1 = z_1, E_{1,2}(t_{p,d})) - \mathbb{P}(E_{2,3} | Z_1 = z_1, E_{1,2}(0)) \} f_d(z_1) dz_1. \end{aligned} \quad (16)$$

Conditioning on $Z_1 = z_1$, it is easy to verify that Z_2 has the same distribution as $\sqrt{1 - z_1^2} Z'$, where Z' has density f_{d-1} . Therefore

$$\begin{aligned} \mathbb{P}(E_{2,3} | Z_1 = z_1, E_{1,2}(t_{p,d})) &= \mathbb{P}\left(\sqrt{1 - z_1^2} \sqrt{1 - t_{p,d}^2} Z' + t_{p,d} z_1 \geq t_{p,d}\right) \\ &= \mathbb{P}\left(\sqrt{1 - t_{p,d}^2} Z' \geq \sqrt{\frac{1 - z_1}{1 + z_1}} t_{p,d}\right). \end{aligned}$$

Since $\sqrt{(1 - z_1)/(1 + z_1)}$ is a decreasing function of z_1 , the right hand side of the preceding inequality is increasing in z_1 . It is clear that $\mathbb{P}(E_{2,3} | Z_1 = z_1, E_{1,2}(0))$ is a decreasing function of z_1 . Thus the integral in (16) is bounded from below by p times the difference of the conditional probabilities evaluated at $z_1 = t_{p,d}$. Therefore

$$\begin{aligned} & \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(0)) \\ &\geq p \mathbb{P}\left(\frac{t_{p,d}}{1 + t_{p,d}} \leq Z' \leq \frac{t_{p,d}}{\sqrt{1 - t_{p,d}^2}}\right) \geq p \int_{\frac{t_{p,d}}{1 + t_{p,d}}}^{t_{p,d}} f_{d-1}(x) dx. \end{aligned} \quad (17)$$

Since $f_{d-1}(z)/f_d(z) \geq 1/10$ for all $-1 \leq z \leq 1$, using (10) we learn that,

$$f_{d-1}(z) \geq c' d p t_{p,d} \text{ for all } 0 \leq z \leq t_{p,d},$$

where $c' > 0$ is an absolute constant. Plugging the preceding inequality into (17) gives that

$$\mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(0)) \geq c d p^2 t_{p,d}^3 \stackrel{(9)}{\geq} c' \left(\frac{1}{2} - p\right)^3 p^2 \left(\frac{(\log(1/p))^{3/2}}{\sqrt{d}} \wedge d\right),$$

where $c, c' > 0$ are absolute constants and the last inequality follows from (9). By our assumption that $d \geq \frac{1}{4} \log(1/p)$, we get that

$$\mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}, E_{2,3} | E_{1,2}(0)) \geq \frac{1}{16} c' \left(\frac{1}{2} - p\right)^3 p^2 \frac{\log(1/p)^{3/2}}{\sqrt{d}},$$

Plugging the preceding inequality into (15) completes the proof of (6) and (7) for the case $p < 1/2$ (note that $\mathbb{P}(E_{1,3}, E_{2,3}) = p^2$).

It remains to prove the lemma for the case $p \geq 1/2$. Since the case $p = 1/2$ is established in Lemma 5 in Section 6, we may assume that $\frac{1}{2} < p < 1$. This case is treated by considering the event $\tilde{E} = E_{1,2}^C \cap E_{1,3}^C \cap E_{2,3}^C$ and by observing that, since $\mathbb{P}(E_{1,2}) = p$ and $\mathbb{P}(E_{1,2} \cap E_{1,3}) = p^2$, we have by the inclusion-exclusion principle that

$$\mathbb{P}(E) = 1 - 3p + 3p^2 - \mathbb{P}(\tilde{E}). \quad (18)$$

In light of this equation, we need to establish an upper bound on $\mathbb{P}(\tilde{E})$. Our computation in what follows can be considered as a dual of the case for $p < 1/2$. Note that $\mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(x))$ is an increasing function of x . Therefore,

$$\mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C) \leq \frac{1}{2} (\mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(0))).$$

We claim that the proof of the bound (7) (and thus of the entire lemma) is concluded if we manage to show that

$$\mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(0)) \leq -c_p/\sqrt{d}, \quad (19)$$

for a constant $c_p > 0$ which depends only on p . Indeed, using the last two equations we would then have that

$$\mathbb{P}(\tilde{E}) = (1-p)\mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}) \leq -(1-p)c_p/\sqrt{d} + (1-p)^3,$$

which, combined with (18), gives that $\mathbb{P}(E) \geq p^3 + (1-p)c_p/\sqrt{d}$. It thus remains to prove (19).

As above, we write $Z_1 = \langle X_3, X_1 \rangle$ and $Z_2 = \left\langle X_3, \frac{\text{Proj}_{X_1^\perp} X_2}{|\text{Proj}_{X_1^\perp} X_2|} \right\rangle$. We see that

$$\begin{aligned} & \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(0)) \\ &= \int_{z_1 \leq t_{p,d}} (\mathbb{P}(E_{2,3}^C | Z_1 = z_1, E_{1,2}(t_{p,d})) - \mathbb{P}(E_{2,3}^C | Z_1 = z_1, E_{1,2}(0))) f_d(z_1) dz_1. \end{aligned} \quad (20)$$

Conditioning on $Z_1 = z_1$, we see that Z_2 has the same distribution as $\sqrt{1-z_1^2}Z'$, where Z' has density f_{d-1} . Therefore

$$\begin{aligned} \mathbb{P}(E_{2,3}^C | Z_1 = z_1, E_{1,2}(t_{p,d})) &= \mathbb{P}\left(\sqrt{1-z_1^2}\sqrt{1-t_{p,d}^2}Z' + t_{p,d}z_1 \leq t_{p,d}\right) \\ &= \mathbb{P}\left(\sqrt{1-t_{p,d}^2}Z' \leq \sqrt{\frac{1-z_1}{1+z_1}}t_{p,d}\right). \end{aligned}$$

Since $\sqrt{(1-z_1)/(1+z_1)}$ is a decreasing function of z_1 (and $t_{p,d} < 0$ in this case), the right hand side of the preceding inequality is increasing in z_1 . It is clear that $\mathbb{P}(E_{2,3}^C | Z_1 = z_1, E_{1,2}(0))$ is a decreasing function of z_1 . Thus, the right hand side of (20) is bounded from above by p times the integrand evaluated at $z_1 = t_{p,d}$. Therefore,

$$\begin{aligned} \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C | E_{1,2}(0)) &\leq -p\mathbb{P}\left(\frac{t_{p,d}}{1+t_{p,d}} \leq Z' \leq \frac{t_{p,d}}{\sqrt{1-t_{p,d}^2}}\right) \\ &\leq -pt_{p,d} \left(\frac{1}{\sqrt{1-t_{p,d}^2}} - \frac{1}{1+t_{p,d}}\right) f_{d-1}\left(\frac{t_{p,d}}{1+t_{p,d}}\right). \end{aligned}$$

It is easily seen that the above is smaller than $-C_{p,d}$ where $C_{p,d} > 0$ for any fixed d and $\frac{1}{2} < p < 1$. Thus, by choosing the constant c_p to be small enough, we can make sure that equation (19) holds true for any fixed value of d . In other words, we may legitimately assume that $d > \tilde{c}_p$ where \tilde{c}_p depends on p . Also, since $t_{1-p,d} = -t_{p,d}$, equation (9) teaches us that

$$|t_{p,d}| \leq C' \sqrt{\frac{\log \frac{1}{1-p}}{d}}, \quad (21)$$

for a universal constant $C' > 0$. In turn by taking \tilde{c}_p to be large enough, we may assume that $|t_{p,d}| < 1/10$. Plugging this assumption to the last inequality, we deduce that

$$\mathbb{P}(E_{1,3}^C, E_{2,3}^C \mid E_{1,2}(t_{p,d})) - \mathbb{P}(E_{1,3}^C, E_{2,3}^C \mid E_{1,2}(0)) \leq -cpt_{p,d}^2 f_{d-1}(2t_{p,d}).$$

for a universal constant $c > 0$. Using the bounds (11) and (21), we have

$$f_{d-1}(2t_{p,d}) \geq \frac{1}{2} \sqrt{d} (1 - 4t_{p,d}^2)^{(d-3)/2} \geq \frac{1}{2} \sqrt{d} \left(1 - 4(C')^2 \frac{\log \frac{1}{1-p}}{d}\right)^{(d-3)/2} \geq \sqrt{d} c'_p$$

where c'_p depends only on p . Combining the last two inequalities with (9) finally yields (19). \blacksquare

3 Proof of Theorem 2

Recall from the introduction that $T(G)$ and $\tau(G)$ denote the number of triangles and signed triangles, respectively, of a graph G (see (4) and (5)). We denote by A the adjacency matrix of a graph G ; we omit the dependence of A on G , as the graph G will always be obvious from the context. To abbreviate notation, for distinct vertices i, j , and k , we define $\bar{A}_{i,j} := A_{i,j} - \mathbb{E}[A_{i,j}]$ and $\tau_G(i, j, k) := \bar{A}_{i,j} \bar{A}_{i,k} \bar{A}_{j,k}$. The number of signed triangles of a graph G with vertex set $[n]$ can then be written as follows:

$$\tau(G) = \sum_{\{i,j,k\} \subset \binom{[n]}{3}} \tau_G(i, j, k).$$

In order to prove Theorem 2, we need to show that $\tau(G(n, p))$ and $\tau(G(n, p, d))$ behave very differently as long as d is much smaller than n^3 . This is done by establishing that the expectations of $\tau(G(n, p))$ and $\tau(G(n, p, d))$ differ by a quantity which exceeds both of the corresponding standard deviations. We break the proof into four parts accordingly: in Section 3.1 we analyze $\tau(G(n, p))$, in Section 3.2 we give a lower bound on $\mathbb{E} \tau(G(n, p, d))$, in Section 3.3 we give an upper bound on the variance of $\tau(G(n, p, d))$, and finally in Section 3.4 we conclude the proof of Theorem 2.

3.1 Analysis of $\tau(G(n, p))$

The analysis of the statistic τ in the Erdős-Rényi model is straightforward. First, note that for each $\{i, j\} \in \binom{[n]}{2}$ one has $\mathbb{E}[\bar{A}_{i,j}] = 0$. Therefore by the independence of edges it is clear that $\mathbb{E} \tau_{G(n,p)}(i, j, k) = 0$ for all $\{i, j, k\} \subset \binom{[n]}{3}$, and so $\mathbb{E} \tau(G(n, p)) = 0$.

In order to compute the variance, we again use the independence of edges to see that

$$\mathbb{E}[\tau_{G(n,p)}(1, 2, 3) \tau_{G(n,p)}(1, 2, 4)] = \mathbb{E}[\bar{A}_{1,2}^2 \bar{A}_{1,3} \bar{A}_{2,3} \bar{A}_{1,4} \bar{A}_{2,4}] = \mathbb{E}[\bar{A}_{1,2}^2 \bar{A}_{1,3} \bar{A}_{2,3} \bar{A}_{1,4}] \mathbb{E}[\bar{A}_{2,4}] = 0.$$

Moreover, a simple calculation reveals that $\mathbb{E}[\bar{A}_{i,j}^2] = p(1-p)$ for every $\{i, j\} \in \binom{[n]}{2}$. We conclude that for two triplets of vertices with indices i, j, k and i', j', k' we have that

$$\mathbb{E}[\tau_{G(n,p)}(i, j, k) \tau_{G(n,p)}(i', j', k')] = \begin{cases} 0 & \text{if } \{i, j, k\} \neq \{i', j', k'\}, \\ p^3(1-p)^3 & \text{if } \{i, j, k\} = \{i', j', k'\}. \end{cases}$$

Thus the variance of $\tau(G(n, p))$ is

$$\text{Var}(\tau(G(n, p))) = \sum_{\{i, j, k\}} \sum_{\{i', j', k'\}} \mathbb{E}[\tau_{G(n,p)}(i, j, k) \tau_{G(n,p)}(i', j', k')] = \binom{n}{3} p^3 (1-p)^3.$$

We have therefore established that

$$\mathbb{E} \tau(G(n, p)) = 0 \quad \text{and} \quad \text{Var}(\tau(G(n, p))) \leq \binom{n}{3}. \quad (22)$$

3.2 A lower bound for $\mathbb{E} \tau(G(n, p, d))$

Our main goal here is to prove the following.

Lemma 3 *For every $0 < p \leq 1/2$, there exists a constant $C_p > 0$ (depending only on p) such that for all n and d we have*

$$\mathbb{E} \tau(G(n, p, d)) \geq \binom{n}{3} \frac{C_p}{\sqrt{d}}. \quad (23)$$

Proof It suffices to estimate $\mathbb{E} \tau_{G(n,p,d)}(1, 2, 3)$, since $\mathbb{E} \tau(G(n, p, d)) = \binom{[n]}{3} \mathbb{E} \tau_{G(n,p,d)}(1, 2, 3)$. We have that

$$\begin{aligned} \mathbb{E} \tau_{G(n,p,d)}(1, 2, 3) &= \mathbb{E} \bar{A}_{1,2} \bar{A}_{1,3} \bar{A}_{2,3} = \mathbb{E} (A_{1,2} - p) (A_{1,3} - p) (A_{2,3} - p) \\ &= \mathbb{E} A_{1,2} A_{1,3} A_{2,3} - p (\mathbb{E} A_{1,2} A_{1,3} + \mathbb{E} A_{1,2} A_{2,3} + \mathbb{E} A_{1,3} A_{2,3}) \\ &\quad + p^2 (\mathbb{E} A_{1,2} + \mathbb{E} A_{1,3} + \mathbb{E} A_{2,3}) - p^3 \\ &= \mathbb{E} A_{1,2} A_{1,3} A_{2,3} - p^3, \end{aligned}$$

where the last equality follows from the simple facts that $\mathbb{E} A_{i,j} = p$ and $\mathbb{E} A_{i,j} A_{i,k} = p^2$ for all triples $\{i, j, k\} \subset \binom{[n]}{3}$. The bound (7) from Lemma 1 then gives that

$$\mathbb{E} \tau_{G(n,p,d)}(1, 2, 3) \geq \frac{C_p}{\sqrt{d}}$$

for some $C_p > 0$, which concludes the proof. ■

3.3 The variance of $\tau(G(n, p, d))$

The estimation of the variance of $\tau(G(n, p, d))$ requires the following technical lemma.

Lemma 4 *For every $p \in [0, 1]$ we have that*

$$\mathbb{E}[\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4)] \leq \pi^2/d.$$

Proof Define

$$g(x) := \mathbb{E} \left[\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4) \mid \langle X_1, X_2 \rangle = x \right].$$

Note first that by rotational invariance, and by the independence of X_3 and X_4 , we have that

$$\begin{aligned} g(x) &= \mathbb{E} \left[\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4) \mid X_1 = e_1, X_2 = xe_1 + \sqrt{1-x^2}e_2 \right] \\ &= \left(\mathbf{1}_{\{x < t_{p,d}\}} p^2 + \mathbf{1}_{\{x \geq t_{p,d}\}} (1-p)^2 \right) \left(\mathbb{E} \left[\bar{A}_{1,3} \bar{A}_{2,3} \mid X_1 = e_1, X_2 = xe_1 + \sqrt{1-x^2}e_2 \right] \right)^2, \end{aligned}$$

where, as before, $\{e_1, \dots, e_d\}$ denotes the standard basis of \mathbb{R}^d . Define the spherical cap

$$B_x := \left\{ y \in \mathbb{S}^{d-1} : \langle y, xe_1 + \sqrt{1-x^2}e_2 \rangle \geq t_{p,d} \right\}$$

and let $S_x := B_1 \cap B_x$ denote the intersection of two such caps. Let σ denote the normalized Haar measure on the sphere. We have by definition that

$$\begin{aligned} &\mathbb{E} \left[\bar{A}_{1,3} \bar{A}_{2,3} \mid X_1 = e_1, X_2 = xe_1 + \sqrt{1-x^2}e_2 \right] \\ &= (1-p)^2 \sigma(S_x) - p(1-p) (\sigma(B_1 \setminus S_x) + \sigma(B_x \setminus S_x)) + p^2 \sigma(\mathbb{S}^{d-1} \setminus (B_1 \cup B_x)) \\ &= (1-p)^2 \sigma(S_x) - 2p(1-p)(p - \sigma(S_x)) + p^2(1 - 2p + \sigma(S_x)) = \sigma(S_x) - p^2, \end{aligned}$$

and so

$$g(x) = \left(\mathbf{1}_{\{x < t_{p,d}\}} p^2 + \mathbf{1}_{\{x \geq t_{p,d}\}} (1-p)^2 \right) (\sigma(S_x) - p^2)^2 \leq (\sigma(S_x) - p^2)^2.$$

Let X be a random variable with the same law as that of $\langle X_1, X_2 \rangle$; in other words, X has density f_d , see (8) in Section 2. Thus by conditioning on $\langle X_1, X_2 \rangle$ and using the definition of $g(\cdot)$ we have that

$$\mathbb{E} \left[\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4) \right] = \mathbb{E} [g(X)] \leq \mathbb{E} \left[(\sigma(S_X) - p^2)^2 \right]. \quad (24)$$

Next, we claim that

$$\mathbb{E} [\sigma(S_X)] = p^2. \quad (25)$$

To see this, define, using a slight abuse of notation, $B_\theta = \{y \in \mathbb{S}^{d-1} : \langle y, \theta \rangle \geq t_{p,d}\}$ for $\theta \in \mathbb{S}^{d-1}$, and let Y, Z , and W be independent, uniformly distributed points on \mathbb{S}^{d-1} . By rotational invariance, $\sigma(S_X)$ has the same distribution as $\sigma(B_Y \cap B_Z)$, and so

$$\mathbb{E} [\sigma(S_X)] = \mathbb{E} [\sigma(B_Y \cap B_Z)] = \mathbb{E} [\mathbb{P}(W \in B_Y \cap B_Z)] = \mathbb{P}(W \in B_Y \cap B_Z) = (\mathbb{P}(W \in B_Y))^2 = p^2$$

as claimed. Thus by (24) and (25) we have that

$$\mathbb{E} \left[\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4) \right] \leq \mathbb{V}ar [h(X)],$$

where we define $h(x) := \sigma(S_x)$.

Now since $1 = \mathbb{E} [|Y|^2] = d \times \mathbb{E} [X^2]$ by linearity of expectation, we have that

$$\mathbb{E} [X^2] = \frac{1}{d}.$$

We claim that the proof follows by showing that for all $-1 \leq x, y \leq 1$,

$$|h(x) - h(y)| \leq 2 |\arcsin(x) - \arcsin(y)|. \quad (26)$$

Indeed, observe that, since $\arcsin(x)$ is odd and convex on $[0, 1]$, we have that

$$|\arcsin(x)| \leq |x| \arcsin(1) = |x| \frac{\pi}{2}, \quad \forall -1 \leq x \leq 1, \quad (27)$$

and now combining the last four displays we get that

$$\begin{aligned} \mathbb{E} [\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4)] &\leq \text{Var} [h(X)] \leq \mathbb{E} [(h(X) - h(0))^2] \\ &\stackrel{(26)}{\leq} 4 \mathbb{E} [(\arcsin(X))^2] \stackrel{(27)}{\leq} \pi^2 \mathbb{E} [X^2] = \frac{\pi^2}{d}, \end{aligned}$$

which concludes the proof, using (26).

It thus remains to prove inequality (26). To this end, consider the foliation of the sphere

$$\mathbb{S}^{d-1} = \bigcup_{z \in \mathbf{B}^{d-2}} W_z,$$

where \mathbf{B}^{d-2} is the $(d-2)$ -dimensional Euclidean unit ball and $W_z = \{(x, y, z); x^2 + y^2 = 1 - |z|^2\}$, with a corresponding decomposition $\theta \mapsto (x_\theta, y_\theta, z_\theta)$. Let μ denote the measure on \mathbf{B}^{d-2} which is the push-forward of σ under the map $\theta \mapsto z_\theta$, and for $z \in \mathbf{B}^{d-2}$ let σ_z denote the uniform measure on W_z so that

$$\int_{\mathbb{S}^{d-1}} f(\theta) d\sigma(\theta) = \int_{\mathbf{B}^{d-2}} \int_{\mathbb{S}^1} f(x, y, z) d\sigma_z(x, y) d\mu(z)$$

for all measurable $f : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$. Clearly for all $0 < x, y < 1$ we have that

$$\sigma(S_x) - \sigma(S_y) = \int_{\mathbf{B}^{d-2}} (\sigma_z(S_x \cap W_z) - \sigma_z(S_y \cap W_z)) d\mu(z).$$

By the triangle inequality, we have for all $z \in \mathbf{B}^{d-2}$ that

$$|\sigma_z(S_x \cap W_z) - \sigma_z(S_y \cap W_z)| \leq \sigma_z((S_x \Delta S_y) \cap W_z) \leq \sigma_z((B_x \Delta B_y) \cap W_z).$$

But clearly,

$$\sigma_z((B_x \Delta B_y) \cap W_z) \leq 2 |\arcsin(x) - \arcsin(y)| \sigma_z(W_z).$$

Therefore, by the last displays and the triangle inequality, we have that

$$|\sigma(S_x) - \sigma(S_y)| \leq 2 |\arcsin(x) - \arcsin(y)| \sigma(\mathbb{S}^{d-1}) = 2 |\arcsin(x) - \arcsin(y)|,$$

which is exactly (26). ■

Using this lemma we can finally give an upper bound on the variance of $\tau(G(n, p, d))$. By expanding the variance as

$$\text{Var} (\tau(G(n, p, d))) = \sum_{\{i,j,k\}} \sum_{\{i',j',k'\}} \mathbb{E} [\tau_{G(n,p,d)}(i, j, k) \tau_{G(n,p,d)}(i', j', k')] \quad (28)$$

and by using the fact that X_1, \dots, X_n are i.i.d., it is clearly enough to analyze the following terms:

$$\begin{aligned} W_1 &:= \mathbb{E} [\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 3)], \quad W_2 := \mathbb{E} [\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 2, 4)], \\ W_3 &:= \mathbb{E} [\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(1, 4, 5)], \quad W_4 := \mathbb{E} [\tau_{G(n,p,d)}(1, 2, 3) \tau_{G(n,p,d)}(4, 5, 6)]. \end{aligned}$$

By symmetry, each term in the right hand side of (28) corresponds to one of these four ‘‘types’’. Observe first that $\mathbb{E} \bar{A}_{1,3} = 0$, and since $\bar{A}_{1,3}$ is independent of both $\tau_{G(n,p,d)}(1, 4, 5)$ and $\tau_{G(n,p,d)}(4, 5, 6)$, we have that $W_3 = W_4 = 0$. By definition we have $|W_1| \leq 1$, and Lemma 4 states that $W_2 \leq \pi^2/d$. Putting everything together, we finally conclude that

$$\text{Var} (\tau(G(n, p, d))) \leq \binom{n}{3} + \binom{n}{4} \binom{4}{2} \frac{\pi^2}{d} \leq n^3 + \frac{3n^4}{d}. \quad (29)$$

3.4 Concluding the proof

Combining displays (22), (23) and (29) together, we get that

$$\mathbb{E}(\tau(G(n, p))) = 0, \quad \mathbb{E}[\tau(G(n, p, d))] \geq \binom{n}{3} \frac{C_p}{\sqrt{d}}$$

and

$$\max \{ \mathbb{V}\text{ar} [\tau(G(n, p))], \mathbb{V}\text{ar} [\tau(G(n, p, d))] \} \leq n^3 + \frac{3n^4}{d}.$$

Using Chebyshev's inequality this implies that

$$\mathbb{P}(\tau(G(n, p, d)) \leq \frac{1}{2} \mathbb{E}[\tau(G(n, p, d))]) \leq 200 \frac{dn^3 + 3n^4}{C_p^2 n^6},$$

and also

$$\mathbb{P}(\tau(G(n, p)) \geq \frac{1}{2} \mathbb{E}[\tau(G(n, p, d))]) \leq 200 \frac{dn^3 + 3n^4}{C_p^2 n^6}.$$

Putting the two above displays together we thus have that

$$\text{TV}(\tau(G(n, p, d)), \tau(G(n, p))) \geq 1 - \frac{Cd}{n^3} - \frac{C}{n^2}$$

for a constant C that depends only on p . This concludes the proof of Theorem 2.

4 Proof of Theorem 3

Fix a constant $c > 0$ and let $p = c/n$. To abbreviate notation, for distinct vertices i, j , and k of a graph G , let $T_G(i, j, k) := A_{i,j}A_{i,k}A_{j,k}$; in other words, $T_G(i, j, k)$ is the indicator that the vertices i, j , and k form a triangle in G . We can then write the number of triangles in a graph G with vertex set $[n]$ as

$$T(G) = \sum_{\{i,j,k\} \subset \binom{[n]}{3}} T_G(i, j, k).$$

First, we have a simple estimate for the expectation of $T(G(n, c/n))$:

$$\mathbb{E}[T(G(n, c/n))] = \binom{n}{3} \left(\frac{c}{n}\right)^3 \leq c^3. \quad (30)$$

We now turn to estimating the first two moments of $G(n, p, d)$. An application of the inequality (6) from Lemma 1 gives that

$$\mathbb{E}[T(G(n, c/n, d))] \geq \kappa c^3 \frac{(\log \frac{n}{c})^{3/2}}{\sqrt{d}} \quad (31)$$

for a universal constant $\kappa > 0$. Note that the right hand side of the inequality (31) goes to infinity when $d/\log^3(n) \rightarrow 0$.

In order to establish an upper bound for the variance of $T(G(n, c/n, d))$, first note that $T_{G(n,p,d)}(i, j, k)$ and $T_{G(n,p,d)}(i', j', k')$ are independent whenever $\{i, j, k\}$ and $\{i', j', k'\}$ do not

share an edge, i.e., whenever $|\{i, j, k\} \cap \{i', j', k'\}| \leq 1$. Furthermore, using the independence of $T_{G(n,p,d)}(1, 2, 3)$ and $A_{1,4}$, we have that

$$\mathbb{E} [T_{G(n,p,d)}(1, 2, 3) T_{G(n,p,d)}(1, 2, 4)] \leq \mathbb{E} [T_{G(n,p,d)}(1, 2, 3) A_{1,4}] = \mathbb{E} [T_{G(n,p,d)}(1, 2, 3)] \times p.$$

Using these facts and expanding $T(G(n, p, d))^2$ as a sum of indicators, we have that

$$\begin{aligned} \mathbb{E} [T(G(n, p, d))^2] &= \sum_{\{i,j,k\}} \sum_{\{i',j',k'\}} \mathbb{E} [T_{G(n,p,d)}(i, j, k) T_{G(n,p,d)}(i', j', k')] \\ &\leq \binom{n}{3}^2 \mathbb{E} [T_{G(n,p,d)}(1, 2, 3) T_{G(n,p,d)}(4, 5, 6)] \\ &\quad + \binom{n}{4} \binom{4}{2} \mathbb{E} [T_{G(n,p,d)}(1, 2, 3) T_{G(n,p,d)}(1, 2, 4)] + \binom{n}{3} \mathbb{E} [T_{G(n,p,d)}(1, 2, 3)^2] \\ &\leq (\mathbb{E} [T(G(n, p, d))])^2 + \binom{n}{3} (2np + 1) \mathbb{E} [T_{G(n,p,d)}(1, 2, 3)], \end{aligned}$$

and so

$$\text{Var} [T(G(n, c/n, d))] \leq (2c + 1) \mathbb{E} [T(G(n, c/n, d))]. \quad (32)$$

Now using this estimate, together with Chebyshev's inequality, gives that

$$\mathbb{P}(T(G(n, \frac{c}{n}, d)) \leq \frac{1}{2} \mathbb{E} [T(G(n, \frac{c}{n}, d))]) \leq \frac{4 \text{Var} (T(G(n, \frac{c}{n}, d)))}{(\mathbb{E} [T(G(n, \frac{c}{n}, d))])^2} \leq \frac{8c + 4}{\mathbb{E} [T(G(n, \frac{c}{n}, d))]}.$$

On the other hand, Markov's inequality and the estimate (30) together give that

$$\mathbb{P}(T(G(n, \frac{c}{n})) \geq \frac{1}{2} \mathbb{E} [T(G(n, \frac{c}{n}, d))]) \leq \frac{2c^3}{\mathbb{E} [T(G(n, \frac{c}{n}, d))]}.$$

Finally, note that both of these upper bounds go to 0 when $d/\log^3(n) \rightarrow 0$, as can be seen from the estimate (31). This concludes the proof of Theorem 3.

5 Proof of the lower bound

Recall that if Y_1 is a standard normal random variable in \mathbb{R}^d , then $Y_1/\|Y_1\|$ is uniformly distributed on the sphere \mathbb{S}^{d-1} . Consequently we can view $G(n, p, d)$ as a function of an appropriate Wishart matrix, as follows. Let Y be an $n \times d$ matrix where the entries are i.i.d. standard normal random variables, and let $W \equiv W(n, d) = YY^T$ be the corresponding $n \times n$ Wishart matrix. Note that $W_{ii} = \langle Y_i, Y_i \rangle = \|Y_i\|^2$ and so $\langle Y_i/\|Y_i\|, Y_j/\|Y_j\| \rangle = W_{ij}/\sqrt{W_{ii}W_{jj}}$. Thus the $n \times n$ matrix A defined as

$$A_{i,j} = \begin{cases} 1 & \text{if } W_{ij}/\sqrt{W_{ii}W_{jj}} \geq t_{p,d} \text{ and } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

has the same law as the adjacency matrix of $G(n, p, d)$. Denote the map that takes W to A by $H_{p,d}$, i.e., $A = H_{p,d}(W)$.

In a similar way we can view $G(n, p)$ as a function of an $n \times n$ matrix drawn from the Gaussian Orthogonal Ensemble (GOE). Let $M(n)$ be a symmetric $n \times n$ random matrix where the diagonal entries are i.i.d. normal random variables with mean zero and variance 2, and the entries above the

diagonal are i.i.d. standard normal random variables, with the entries on and above the diagonal all independent. Then the $n \times n$ matrix B defined as

$$B_{i,j} = \begin{cases} 1 & \text{if } M(n)_{i,j} \geq \bar{\Phi}^{-1}(p) \text{ and } i \neq j \\ 0 & \text{otherwise} \end{cases}$$

has the same law as the adjacency matrix of $G(n, p)$. Note that B only depends on the off-diagonal elements of $M(n)$, so in the definition of B we can replace $M(n)$ with $M(n, d) := \sqrt{d}M(n) + dI_n$, where I_n is the $n \times n$ identity matrix, provided we also replace $\bar{\Phi}^{-1}(p)$ with $\sqrt{d} \times \bar{\Phi}^{-1}(p)$. Denote the map that takes $M(n, d)$ to B by $K_{p,d}$, i.e., $B = K_{p,d}(M(n, d))$. The maps $H_{p,d}$ and $K_{p,d}$ are different, but very similar, as we shall quantify later.

Using the triangle inequality, we can conclude from the previous two paragraphs that for any $p \in [0, 1]$ we have that

$$\begin{aligned} \text{TV}(G(n, p, d), G(n, p)) &= \text{TV}(H_{p,d}(W(n, d)), K_{p,d}(M(n, d))) \\ &\leq \text{TV}(H_{p,d}(W(n, d)), H_{p,d}(M(n, d))) + \text{TV}(H_{p,d}(M(n, d)), K_{p,d}(M(n, d))) \\ &\leq \text{TV}(W(n, d), M(n, d)) + \text{TV}(H_{p,d}(M(n, d)), K_{p,d}(M(n, d))). \end{aligned}$$

The second term in the line above is a smaller order error term which we deal with in Appendix A, and thus Theorem 1(c) follows from the following result, which is a more precise restatement of Theorem 4 from the Introduction.

Theorem 7 *Define the random matrix ensembles $W(n, d)$ and $M(n, d)$ as above. If $d/n^3 \rightarrow \infty$, then*

$$\text{TV}(W(n, d), M(n, d)) \rightarrow 0. \quad (33)$$

After proving this result, we learned of the work of Jiang and Li [20], who also study the question of when can Wishart matrices be approximated by a GOE random matrix. They prove a result analogous to Theorem 7 for the joint eigenvalue distributions, from which Theorem 7 follows by orthogonal invariance. Nonetheless, we present here our proof, for the following two reasons: (1) to keep the paper self-contained, and (2) because our proof is shorter and simpler than the one presented in [20].

Theorem 7 thus states that as $d/n^3 \rightarrow \infty$, all statistics of the random matrix ensembles $W(n, d)$ and $M(n, d)$ have asymptotically the same distribution. We note that in the random matrix literature there has been lots of work showing that particular statistics (e.g., the empirical spectral distribution, or the largest eigenvalue) of these two ensembles have asymptotically the same distribution, even when $d/n \rightarrow c \in [0, \infty]$, see, e.g., [7, 21, 13, 14].

Proof We show (33) by comparing the densities of the two random matrix ensembles. Let $\mathcal{P} \subset \mathbb{R}^{n^2}$ denote the cone of positive semidefinite matrices. It is well known (see, e.g., [28]) that when $d \geq n$, $W(n, d)$ has the following density with respect to the Lebesgue measure on \mathcal{P} :

$$f_{n,d}(A) := \frac{(\det(A))^{\frac{1}{2}(d-n-1)} \exp(-\frac{1}{2}\text{Tr}(A))}{2^{\frac{1}{2}dn} \pi^{\frac{1}{4}n(n-1)} \prod_{i=1}^n \Gamma(\frac{1}{2}(d+1-i))},$$

where $\text{Tr}(A)$ denotes the trace of the matrix A . It is also known that the density of a GOE random matrix with respect to the Lebesgue measure on \mathbb{R}^{n^2} is $A \mapsto (2\pi)^{-\frac{1}{4}n(n+1)} 2^{-\frac{n}{2}} \exp(-\frac{1}{4}\text{Tr}(A^2))$

and so the density of $M(n, d)$ with respect to the Lebesgue measure on \mathbb{R}^{n^2} is

$$g_{n,d}(A) := \frac{\exp\left(-\frac{1}{4d}\text{Tr}\left((A - dI_n)^2\right)\right)}{(2\pi d)^{\frac{1}{4}n(n+1)} 2^{\frac{n}{2}}}.$$

Denote the measure given by this density by $\mu_{n,d}$, let λ denote the Lebesgue measure on \mathbb{R}^{n^2} and write $A \succeq 0$ if A is positive semidefinite. We can then write

$$\begin{aligned} \text{TV}(W(n, d), M(n, d)) &= \frac{1}{2} \int_{\mathbb{R}^{n^2}} |f_{n,d}(A) \mathbf{1}_{\{A \succeq 0\}} - g_{n,d}(A)| d\lambda(A) \\ &= \frac{1}{2} \int_{\mathbb{R}^{n^2}} \left| \frac{f_{n,d}(A) \mathbf{1}_{\{A \succeq 0\}}}{g_{n,d}(A)} - 1 \right| d\mu_{n,d}(A). \end{aligned}$$

Since

$$\int_{\mathbb{R}^{n^2}} \left(\frac{f_{n,d}(A) \mathbf{1}_{\{A \succeq 0\}}}{g_{n,d}(A)} - 1 \right) d\mu_{n,d}(A) = 0,$$

in order to show that $\text{TV}(W(n, d), M(n, d)) = o(1)$, it suffices to show that

$$\frac{f_{n,d}(A) \mathbf{1}_{\{A \succeq 0\}}}{g_{n,d}(A)} = 1 + o(1)$$

as $d/n^3 \rightarrow \infty$ with probability $1 - o(1)$ according to the measure $\mu_{n,d}$.

It is known (see, e.g., [4]) that, with probability $1 - o(1)$, all the eigenvalues of $M(n)$ are in the interval $[-3\sqrt{n}, 3\sqrt{n}]$ and so all the eigenvalues of $M(n, d)$ are in the interval $[d - 3\sqrt{dn}, d + 3\sqrt{dn}]$. Since $d/n^3 \rightarrow \infty$, we thus have $\mathbb{P}(M(n, d) \succeq 0) = 1 - o(1)$, and so we may restrict our attention to \mathcal{P} . Define $\alpha_{n,d}(A) := \log(f_{n,d}(A)/g_{n,d}(A))$. It remains then to show that $\alpha_{n,d}(A) = o(1)$ as $d/n^3 \rightarrow \infty$ with probability $1 - o(1)$ according to the measure $\mu_{n,d}$.

Denote the eigenvalues of an $n \times n$ matrix A by $\lambda_1(A) \leq \dots \leq \lambda_n(A)$; when the matrix is obvious from the context, we omit the dependence on A . Recall that $\det(A) = \prod_{i=1}^n \lambda_i$ and $\text{Tr}(A) = \sum_{i=1}^n \lambda_i$. We then have

$$\begin{aligned} \alpha_{n,d}(A) &= \frac{1}{2} \sum_{i=1}^n \left\{ (d - n - 1) \log \lambda_i - \lambda_i + \frac{1}{2d} (\lambda_i - d)^2 \right\} \\ &\quad + \left\{ \frac{n(n+3)}{4} - \frac{dn}{2} \right\} \log 2 + \frac{n}{2} \log \pi + \frac{n(n+1)}{4} \log d - \sum_{i=1}^n \log \Gamma\left(\frac{1}{2}(d+1-i)\right). \end{aligned}$$

By Stirling's formula we know that $\log \Gamma(z) = (z - \frac{1}{2}) \log z - z + \frac{1}{2} \log(2\pi) + O(\frac{1}{z})$ as $z \rightarrow \infty$, so

$$\begin{aligned} \alpha_{n,d}(A) &= \frac{1}{2} \sum_{i=1}^n \left\{ (d - n - 1) \log \lambda_i - \lambda_i + \frac{1}{2d} (\lambda_i - d)^2 \right\} \\ &\quad + \frac{n(n+1)}{4} \log d - \frac{1}{2} \sum_{i=1}^n (d - i) \log(d + 1 - i) + \frac{1}{2} \sum_{i=1}^n (d + 1 - i) + O\left(\frac{n}{d}\right). \end{aligned}$$

Now writing $\log(d+1-i) = \log d + \log(1 - \frac{i-1}{d}) = \log d - \frac{i-1}{d} + O(\frac{i^2}{d^2})$ we get that

$$\begin{aligned} \alpha_{n,d}(A) &= \frac{1}{2} \sum_{i=1}^n \left\{ (d-n-1) \log \lambda_i - \lambda_i + \frac{1}{2d} (\lambda_i - d)^2 \right\} \\ &\quad + \frac{1}{2} \left\{ -nd \log d + nd + n(n+1) \log d \right\} + O\left(\frac{n^3}{d}\right). \end{aligned}$$

Defining $h(x) := \frac{1}{2} \left\{ (d-n-1) \log(x/d) - (x-d) + \frac{1}{2d} (x-d)^2 \right\}$, we have that

$$\alpha_{n,d}(A) = \sum_{i=1}^n h(\lambda_i) + O\left(\frac{n^3}{d}\right). \quad (34)$$

Note that the derivatives of h at d are $h(d) = 0$, $h'(d) = -\frac{n+1}{2d}$, $h''(d) = \frac{n+1}{2d^2}$, and $h^{(3)}(d) = \frac{d-n-1}{d^3}$, and also $h^{(4)}(x) = -\frac{3(d-n-1)}{x^4}$. Approximating h with its third order Taylor polynomial around d we get

$$h(x) = -\frac{n+1}{2d} (x-d) + \frac{n+1}{4d^2} (x-d)^2 + \frac{d-n-1}{6d^3} (x-d)^3 - \frac{d-n-1}{8\xi^4} (x-d)^4,$$

where ξ is some real number between x and d . By (34) we need to show that $\sum_{i=1}^n h(\lambda_i) = o(1)$ with probability tending to 1. Recall that with probability $1 - o(1)$ all eigenvalues of $M(n, d)$ are in the interval $[d - 3\sqrt{dn}, d + 3\sqrt{dn}]$. If $\lambda_i \in [d - 3\sqrt{dn}, d + 3\sqrt{dn}]$ and ξ_i is between λ_i and d for all $i \in [n]$, then it is immediate that

$$\sum_{i=1}^n \left| \frac{n+1}{4d^2} (\lambda_i - d)^2 - \frac{d-n-1}{8\xi_i^4} (\lambda_i - d)^4 \right| \leq \frac{cn^3}{d}$$

for some constant c , and so what remains to show is that

$$\sum_{i=1}^n \left\{ -\frac{n+1}{2d} (\lambda_i - d) + \frac{d-n-1}{6d^3} (\lambda_i - d)^3 \right\} = o(1) \quad (35)$$

as $d/n^3 \rightarrow \infty$ with probability $1 - o(1)$. This follows from known results about the moments of the empirical spectral distribution of Wigner matrices. In particular, [4, Theorem 2.1.31] shows that

$$\frac{1}{(nd)^{1/2}} \sum_{i=1}^n (\lambda_i - d) \quad \text{and} \quad \frac{1}{(nd)^{3/2}} \sum_{i=1}^n (\lambda_i - d)^3$$

both converge weakly to a normal distribution with appropriate variance, from which (35) immediately follows. ■

6 Dimension estimation

In this section we prove Theorem 5. The idea for the proof is very similar to that of Theorem 2, and also uses the statistic $\tau(G)$ counting the “number” of signed triangles, analyzed in Section 3.

However, dimension estimation is a slightly more subtle matter, since here it is necessary to have a bound on the difference of the expected number of triangles between consecutive dimensions, rather than just a lower bound on the expected number of triangles in the random geometric graph $G(n, p, d)$, as in Lemma 1. The next lemma gives a bound on this difference; but note that this lemma only deals with the case $p = 1/2$. We believe that this result should hold true for any fixed $0 < p < 1$, but the proof seems to be much more involved.

Lemma 5 *Let $\{e_i : 1 \leq i \leq d\}$ denote the standard basis in \mathbb{R}^d , and define, for each dimension d ,*

$$h(d) := \mathbb{P}(\langle X_1, e_1 \rangle \geq 0, \langle X_2, e_1 \rangle \geq 0, \langle X_1, X_2 \rangle \geq 0),$$

where X_1 and X_2 are independent uniformly distributed random vectors on \mathbb{S}^{d-1} . Then

$$h(d-1) - h(d) \geq cd^{-3/2} \tag{36}$$

for some universal constant $c > 0$.

Proof Define the event $E := \{\langle X_1, e_1 \rangle \geq 0, \langle X_2, e_1 \rangle \geq 0, \langle X_1, X_2 \rangle \geq 0\}$. We first claim that

$$\mathbb{P}(E \mid \langle X_1, e_1 \rangle = t) = \frac{\pi - \arccos t}{2\pi} \mathbf{1}_{\{t \in [0, 1]\}} \tag{37}$$

for all $t \in [-1, 1]$. To see this, first note that by the rotational invariance of the vectors X_1 and X_2 , and since the event E is invariant under the action of the orthogonal group on the three vectors simultaneously, when conditioning on $\langle X_1, e_1 \rangle = t$ we can assume that $X_1 = te_1 + \sqrt{1-t^2}e_2$. Thus for all $t \in [0, 1]$ we have that

$$\mathbb{P}(E \mid \langle X_1, e_1 \rangle = t) = \mathbb{P}\left(\langle X_2, e_1 \rangle \geq 0, t\langle X_2, e_1 \rangle + \sqrt{1-t^2}\langle X_2, e_2 \rangle \geq 0\right). \tag{38}$$

Since the projection of X_2 on the span of e_1 and e_2 is distributed according to a law which is invariant under rotations of this span, we have that the probability in (38) is proportional to the angle of the wedge defined by $\{x \in \text{span}(e_1, e_2) : \langle x, e_1 \rangle \geq 0, t\langle x, e_1 \rangle + \sqrt{1-t^2}\langle x, e_2 \rangle \geq 0\}$. The angle of this wedge is exactly $\pi - \arccos t$, and thus formula (37) follows.

Let Z_d be a random variable with the same law as that of $\langle X_1, e_1 \rangle$; in other words, Z_d has density f_d , see (8) in Section 2. Defining $Y_d = \arccos(Z_d)$ we therefore have that

$$\mathbb{P}(E) = \mathbb{E}\left[\left(\frac{1}{2} - \frac{Y_d}{2\pi}\right) \mathbf{1}_{\{Y_d \in [0, \pi/2]\}}\right] = \frac{1}{4} - \frac{1}{2\pi} \mathbb{E}[Y_d \mathbf{1}_{\{Y_d \in [0, \pi/2]\}}]$$

Using the change of variables $y = \arccos z$ and the formula for f_d in (8), we have that the density of Y_d is given by

$$g_d(x) = f_d(\cos(x)) \sin(x) = \frac{\Gamma(d/2)}{\Gamma((d-1)/2)\sqrt{\pi}} (\sin(x))^{d-2}, \quad x \in [0, \pi].$$

To abbreviate notation, define $A_d := \frac{\Gamma(d/2)}{\Gamma((d-1)/2)\sqrt{\pi}}$. By integrating we thus get that

$$h(d) = \mathbb{P}(E) = \frac{1}{4} - \frac{A_d}{2\pi} \int_0^{\pi/2} x (\sin(x))^{d-2} dx. \tag{39}$$

Let μ_d be the probability measure on $[0, \pi/2]$ defined by

$$\frac{d\mu_d(x)}{dx} = 2A_d (\sin(x))^{d-2}.$$

We then have that $\frac{d\mu_d}{d\mu_{d-1}}(x) = \frac{A_d}{A_{d-1}} \sin(x)$, and so

$$\int_0^{\pi/2} \sin(x) d\mu_{d-1}(x) = \frac{A_{d-1}}{A_d}. \quad (40)$$

Elementary estimates concerning the Γ function give that

$$\frac{A_d}{A_{d-1}} = \frac{\Gamma(\frac{d}{2}) \Gamma(\frac{d-2}{2})}{(\Gamma(\frac{d-1}{2}))^2} \in \left(1 + \frac{1}{12d}, 1 + \frac{2}{d}\right), \quad \text{and} \quad \frac{A_{d-1}}{A_d} \in \left(1 - \frac{2}{d}, 1 - \frac{1}{12d}\right). \quad (41)$$

In the following we write $\nu \equiv \mu_{d-1}$ to abbreviate notation. We then have that

$$\begin{aligned} h(d-1) - h(d) &= \frac{1}{2\pi} \left(A_d \int_0^{\pi/2} x (\sin(x))^{d-2} dx - A_{d-1} \int_0^{\pi/2} x (\sin(x))^{d-3} dx \right) \\ &= \frac{1}{4\pi} \left(\int x d\mu_d(x) - \int x d\mu_{d-1}(x) \right) \\ &= \frac{1}{4\pi} \int x \left(\frac{A_d}{A_{d-1}} \sin(x) - 1 \right) d\nu(x) \\ &= \frac{1}{4\pi} \frac{A_d}{A_{d-1}} \int x \left(\sin(x) - \frac{A_{d-1}}{A_d} \right) d\nu(x) \\ &= \frac{1}{4\pi} \frac{A_d}{A_{d-1}} \int x \left(\sin(x) - \int \sin(y) d\nu(y) \right) d\nu(x). \end{aligned} \quad (42)$$

Since the function $x \mapsto \sin(x)$ is continuous and monotone on the interval $[0, \pi/2]$, there exists a unique $x_0 \in (0, \pi/2)$ such that

$$\sin(x_0) = \int_0^{\pi/2} \sin(y) d\nu(y). \quad (43)$$

Since the function $x \mapsto \sin(x)$ is concave on the interval $[0, \pi/2]$, we have that $\sin(x) - \sin(x_0) \leq \cos(x_0)(x - x_0)$ for all $x \in [0, \pi/2]$. Thus continuing from (42) and also using that $A_d/A_{d-1} \geq 1$ from (41), we have that

$$\begin{aligned} h(d-1) - h(d) &= \frac{1}{4\pi} \frac{A_d}{A_{d-1}} \int_0^{\pi/2} (x - x_0) (\sin(x) - \sin(x_0)) d\nu(x) \\ &\geq \frac{1}{4\pi} \cos(x_0) \int_0^{x_0} (x - x_0)^2 d\nu(x). \end{aligned} \quad (44)$$

Using (40), (41), and (43), we have that

$$x_0 = \arcsin\left(\frac{A_{d-1}}{A_d}\right) \in \left(\arcsin\left(1 - \frac{2}{d}\right), \arcsin\left(1 - \frac{1}{12d}\right)\right),$$

and so

$$\cos(x_0) \geq \cos\left(\arcsin\left(1 - \frac{1}{12d}\right)\right) = \cos\left(\arccos\left(\sqrt{1 - \left(1 - \frac{1}{12d}\right)^2}\right)\right) = \sqrt{\frac{1}{6d} - \frac{1}{144d^2}} \geq \frac{1}{4\sqrt{d}}.$$

Furthermore, since $x \mapsto \arcsin(x)$ is convex on $[0, 1]$ with derivative $1/\sqrt{1-x^2}$, we have that for all $x \in [0, \arcsin(1 - \frac{3}{d})]$,

$$(x - x_0)^2 \geq (\arcsin(1 - \frac{2}{d}) - \arcsin(1 - \frac{3}{d}))^2 \geq \left(\frac{1}{d} \times \frac{1}{\sqrt{1 - (1 - \frac{3}{d})^2}} \right)^2 = \frac{1}{6d - 9} \geq \frac{1}{6d}.$$

Plugging the estimates from the previous three displays back into (44), we get that

$$h(d-1) - h(d) \geq \frac{1}{100\pi d^{3/2}} \nu([0, \arcsin(1 - \frac{3}{d})]). \quad (45)$$

Let Z' be a random variable with density f_{d-1} . Then by the definition of ν we have that

$$\begin{aligned} \nu([0, \arcsin(1 - \frac{3}{d})]) &= 2\mathbb{P}(\arccos(Z') \in [0, \arcsin(1 - \frac{3}{d})]) = 2\mathbb{P}(Z' \geq \cos(\arcsin(1 - \frac{3}{d}))) \\ &= 2\mathbb{P}\left(Z' \geq \cos\left(\arccos\left(\sqrt{1 - (1 - \frac{3}{d})^2}\right)\right)\right) \geq 2\mathbb{P}\left(Z' \geq \frac{\sqrt{6}}{\sqrt{d-1}}\right), \end{aligned}$$

which, by (13), is bounded from below by a universal constant. Together with (45), this concludes the proof. \blacksquare

Proof of Theorem 5 Recall from the proof of Lemma 3 that

$$\mathbb{E}[\tau(G(n, \frac{1}{2}, d))] = \binom{n}{3} (h(d) - (1/2)^3).$$

Thus by Lemma 5 we have for all $d_1 < d_2$ that

$$\begin{aligned} \mathbb{E}[\tau(G(n, \frac{1}{2}, d_1))] - \mathbb{E}[\tau(G(n, \frac{1}{2}, d_2))] &= \binom{n}{3} (h(d_1) - h(d_2)) \geq \binom{n}{3} (h(d_1) - h(d_1 + 1)) \\ &\geq \binom{n}{3} \frac{c}{(d_1 + 1)^{3/2}} \geq \frac{c_1 n^3}{d_1^{3/2}} \end{aligned}$$

for a universal constant $c_1 > 0$. By (29) we have a bound on the variance of these statistics:

$$\max\{\text{Var}(\tau(G(n, \frac{1}{2}, d_1))), \text{Var}(\tau(G(n, \frac{1}{2}, d_2)))\} \leq n^3 + \frac{3n^4}{d_1}.$$

Using Chebyshev's inequality thus gives that

$$\mathbb{P}\left(\tau(G(n, \frac{1}{2}, d_1)) \leq \frac{1}{2} (\mathbb{E}[\tau(G(n, \frac{1}{2}, d_1))] + \mathbb{E}[\tau(G(n, \frac{1}{2}, d_2))])\right) \leq \frac{12n^3 d_1^3 + n^4 d_1^2}{c_1^2 n^6} \leq \frac{24d_1^2}{c_1^2 n^2},$$

where the second inequality holds if $d_1 \leq n$; if $d_1 > n$ then the upper bound on the probability is vacuously true by choosing $c_1 \leq 1$. Similarly we also have that

$$\mathbb{P}\left(\tau(G(n, \frac{1}{2}, d_2)) \geq \frac{1}{2} (\mathbb{E}[\tau(G(n, \frac{1}{2}, d_1))] + \mathbb{E}[\tau(G(n, \frac{1}{2}, d_2))])\right) \leq \frac{12n^3 d_1^3 + n^4 d_1^2}{c_1^2 n^6} \leq \frac{24d_1^2}{c_1^2 n^2}.$$

Putting the two previous displays together concludes the proof, and shows that we can take $C = 48 \max\{1/c_1^2, 1\}$ for the constant in the statement of the theorem. \blacksquare

Acknowledgements

We thank Nouredine El Karoui, Gábor Lugosi, and Johan Ugander for helpful discussions and useful references, and Benedek Valkó for pointing us to ref. [20]. This work was done while J.D. and R.E. were visiting researchers and M.Z.R. was an intern at the Theory Group of Microsoft Research. They thank the Theory Group for their hospitality. J.D. acknowledges support from NSF grant DMS 1313596, and M.Z.R. acknowledges support from NSF grant DMS 1106999.

References

- [1] Emmanuel Abbe, Afonso Bandeira, and Georgina Hall. Exact Recovery in the Stochastic Block Model. *arXiv preprint arXiv:1405.3267*, 2014.
- [2] Ittai Abraham, Shiri Chechik, David Kempe, and Aleksandrs Slivkins. Low-distortion inference of latent similarities from a multiplex social network. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1853–1883. SIAM, 2013.
- [3] Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. In *Proceedings of the Ninth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 594–598. Society for Industrial and Applied Mathematics, 1998.
- [4] Greg W. Anderson, Alice Guionnet, and Ofer Zeitouni. *An Introduction to Random Matrices*. Number 118. Cambridge University Press, 2010.
- [5] Ery Arias-Castro, Sébastien Bubeck, and Gábor Lugosi. Detecting Positive Correlations in a Multivariate Sample. *Bernoulli*, to appear, 2014+.
- [6] Ery Arias-Castro and Nicolas Verzelen. Community Detection in Random Networks. *arXiv preprint arXiv:1302.7099*, 2013.
- [7] Zhidong Bai and Jack W. Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*. Springer, 2nd edition, 2009.
- [8] Quentin Berthet and Philippe Rigollet. Complexity Theoretic Lower Bounds for Sparse Principal Component Detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.
- [9] Peter J. Bickel, Aiyou Chen, and Elizaveta Levina. The method of moments and degree distributions for network models. *The Annals of Statistics*, 39(5):2280–2301, 2011.
- [10] Heinz Breu and David G. Kirkpatrick. Unit disk graph recognition is np-hard. *Computational Geometry*, 9(1):3–24, 1998.
- [11] Sébastien Bubeck, Elchanan Mossel, and Miklós Z. Rácz. On the influence of the seed graph in the preferential attachment model. *arXiv preprint arXiv:1401.4849*, 2014.
- [12] Luc Devroye, András György, Gábor Lugosi, and Frederic Udina. High-dimensional random geometric graphs and their clique number. *Electronic Journal of Probability*, 16:2481–2508, 2011.
- [13] Nouredine El Karoui. On the largest eigenvalue of Wishart matrices with identity covariance when n, p and $p/n \rightarrow \infty$. *arXiv preprint math/0309355*, 2003.

- [14] Nouredine El Karoui. Tracy-Widom limit for the largest eigenvalue of a large class of complex sample covariance matrices. *The Annals of Probability*, pages 663–714, 2007.
- [15] Ronen Eldan. An efficiency upper bound for inverse covariance estimation. *Israel Journal of Mathematics*, to appear, 2014+.
- [16] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 5:17–61, 1960.
- [17] Alan Frieze and Ravi Kannan. A new approach to the planted clique problem. In Ramesh Hariharan, Madhavan Mukund, and V Vinay, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 2 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 187–198, Dagstuhl, Germany, 2008. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [18] Rishi Gupta, Tim Roughgarden, and C. Seshadhri. Decompositions of Triangle-Dense Graphs. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 471–482. ACM, 2014.
- [19] Peter D. Hoff, Adrian E. Raftery, and Mark S. Handcock. Latent Space Approaches to Social Network Analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.
- [20] Tiefeng Jiang and Danning Li. Approximation of Rectangular Beta-Laguerre Ensembles and Large Deviations. *Journal of Theoretical Probability*, to appear, 2014+.
- [21] Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29(2):295–327, 2001.
- [22] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, to appear, 2014+.
- [23] Mathew Penrose. *Random Geometric Graphs*, volume 5 of *Oxford Studies in Probability*. Oxford University Press, 2003.
- [24] Purnamrita Sarkar, Deepayan Chakrabarti, and Andrew W. Moore. Theoretical Justification of Popular Link Prediction Heuristics. In *Conference on Learning Theory*, pages 295–307, 2010.
- [25] Sasha Sodin. Tail-Sensitive Gaussian Asymptotics for Marginals of Concentrated Measures in High Dimension. In Vitali D. Milman and Gideon Schechtman, editors, *Geometric Aspects of Functional Analysis*, volume 1910 of *Lecture Notes in Mathematics*, pages 271–295. Springer Berlin Heidelberg, 2007.
- [26] Nicolas Verzelen and Ery Arias-Castro. Community Detection in Sparse Random Networks. *arXiv preprint arXiv:1308.2955*, 2013.
- [27] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [28] John Wishart. The Generalised Product Moment Distribution in Samples from a Normal Multivariate Population. *Biometrika*, 20A(1/2):32–52, 1928.

A The error term in the proof of the lower bound

Recalling notation from Section 5, we prove here the following lemma, which is necessary in order to conclude the proof of Theorem 1(c).

Lemma 6 *If $d/n^3 \rightarrow \infty$, then $\sup_{p \in [0,1]} \text{TV}(H_{p,d}(M(n,d)), K_{p,d}(M(n,d))) \rightarrow 0$.*

Define $\Psi_d(x) = \int_{x/\sqrt{d}}^1 f_d(y) dy$, and note that by the definition of $t_{p,d}$, $p = \int_{t_{p,d}}^1 f_d(x) dx = \Psi_d(t_{p,d}\sqrt{d})$. During the proof of Lemma 6 we use the following result of Sodin [25].

Lemma 7 *There exist constants $C, C_1, C_2 > 0$, and a sequence $\varepsilon_d \searrow 0$ that satisfies $\varepsilon_d = O(1/d)$, such that the following inequalities hold for all $0 < t < C\sqrt{d}$:*

$$(1 - \varepsilon_d) \bar{\Phi}(t) \exp(-C_1 t^4/d) \leq \Psi_d(t) \leq (1 + \varepsilon_d) \bar{\Phi}(t) \exp(-C_2 t^4/d).$$

We note that in [25, Lemma 1] the fact that ε_d satisfies $\varepsilon_d = O(1/d)$ is not specified, but this can be read from the proof, where this error comes from the error in Stirling's formula.

Proof of Lemma 6 To abbreviate notation, we write $X := H_{p,d}(M(n,d))$, $Y := K_{p,d}(M(n,d))$, and also $M \equiv M(n)$. Define $\widehat{D}_{i,j} := \sqrt{\left(1 + M_{i,i}/\sqrt{d}\right) \left(1 + M_{j,j}/\sqrt{d}\right)}$ and recall that for $1 \leq i < j \leq n$ we have that

$$\begin{aligned} X_{i,j} &= \mathbf{1}_{\{\widehat{D}_{i,j}^{-1} M_{i,j} \geq \sqrt{dt_{p,d}}\}}, \\ Y_{i,j} &= \mathbf{1}_{\{M_{i,j} \geq \bar{\Phi}^{-1}(p)\}}, \end{aligned} \tag{46}$$

and define also

$$Z_{i,j} = \mathbf{1}_{\{M_{i,j} \geq \sqrt{dt_{p,d}}\}}. \tag{47}$$

By the triangle inequality we have that

$$\text{TV}(X, Y) \leq \text{TV}(X, Z) + \text{TV}(Y, Z),$$

and we deal with the second term first. By a union bound, we have that

$$\text{TV}(Y, Z) \leq \mathbb{P}(\exists 1 \leq i < j \leq n : Y_{i,j} \neq Z_{i,j}) \leq n^2 \mathbb{P}(Y_{1,2} \neq Z_{1,2}),$$

and thus it remains to show that $\mathbb{P}(Y_{1,2} \neq Z_{1,2}) = o(n^{-2})$; since $n^3/d \rightarrow 0$, it is enough to show that $\mathbb{P}(Y_{1,2} \neq Z_{1,2}) = O(d^{-2/3})$. Since $M_{1,2}$ is a standard normal random variable, we have that

$$\mathbb{P}(Y_{1,2} \neq Z_{1,2}) = \left| \bar{\Phi}\left(\sqrt{dt_{p,d}}\right) - p \right| = \left| \bar{\Phi}\left(\sqrt{dt_{p,d}}\right) - \Psi_d\left(\sqrt{dt_{p,d}}\right) \right|. \tag{48}$$

For $p = 1/2$ this expression is zero, and noting that $t_{p,d} = -t_{1-p,d}$ by symmetry, it is enough to bound the expression in (48) for $p \in [0, 1/2)$. For $p \in (0, 1/2)$ fixed, note that Lemma 2 states that there exists a constant $C_p < \infty$ such that $0 \leq \sqrt{dt_{p,d}} \leq C_p$. Together with Lemma 7 applied at $t = \sqrt{dt_{p,d}}$, this implies that the expression in (48) is $O(1/d)$.

More generally, by Lemma 2 there exists a universal constant $C > 0$, such that if $p \in (n^{-\alpha}, 1/2)$, then $0 \leq \sqrt{dt_{p,d}} \leq C\sqrt{\alpha \log(n)}$, and if $p \in [0, n^{-\alpha}]$, then $\sqrt{dt_{p,d}} \geq C^{-1}\sqrt{\alpha \log(n)}$. When

$p \in (n^{-\alpha}, 1/2)$, Lemma 7 applied at $t = \sqrt{dt_{p,d}}$ then implies that the expression in (48) is $O(\alpha^2 \log^2(n)/d)$, which is $O(d^{-2/3})$ for constant α . When $p \in [0, n^{-\alpha}]$ we then have that

$$\mathbb{P}(Y_{1,2} \neq Z_{1,2}) \leq p + \bar{\Phi}\left(\sqrt{dt_{p,d}}\right) \leq n^{-\alpha} + \bar{\Phi}\left(C^{-1}\sqrt{\alpha \log(n)}\right) \leq n^{-\alpha} + \exp\left(-\frac{\alpha \log(n)}{2C^2}\right).$$

By choosing $\alpha := \max\{3, 6C^2\}$, this expression becomes $O(n^{-3})$. This concludes the proof that

$$\sup_{p \in [0,1]} \mathbb{P}(Y_{1,2} \neq Z_{1,2}) = O(d^{-2/3} + n^{-3}).$$

We now deal with the term $\text{TV}(X, Z)$. For a matrix A , let $D'(A)$ denote the diagonal matrix obtained from A by setting the non-diagonal terms to zero, and let $D(A) := A - D'(A)$. With $I \equiv I_n$ denoting the $n \times n$ identity matrix, define the function

$$f(A) := \left(I_n + \frac{1}{\sqrt{d}}D'(A)\right)^{-1/2} A \left(I_n + \frac{1}{\sqrt{d}}D'(A)\right)^{-1/2}.$$

Note that Z can be obtained from $D(M)$ by thresholding the entries appropriately (see (47)), and X can be obtained from $D(f(M))$ in the same way (see (46)). Consequently we have that

$$\sup_{p \in [0,1]} \text{TV}(X, Z) \leq \text{TV}(D(f(M)), D(M)), \quad (49)$$

and in the remainder of the proof we show that this latter quantity goes to zero as $n^3/d \rightarrow 0$.

Let $\Omega = \mathbb{R}^{(n^2-n)/2}$, which we canonically identify with the space of symmetric $n \times n$ matrices with zero on the diagonal, and let $\Omega' = \mathbb{R}^n$, which we think of as the space of the diagonal entries on an $n \times n$ matrix. By slight abuse of notation, we allow ourselves to think of D and D' as functions from the space of symmetric $n \times n$ matrices to Ω and Ω' respectively. We can thus naturally view the function $D \circ f$ as a mapping from $\Omega \oplus \Omega'$ to Ω .

In order to bound the right hand side of (49), we need an estimate for the density of $D(f(M))$, which we denote by $w(x)$. Note that $w(x)$ is the density of the push-forward under the function $D \circ f$ of the measure whose density is $\gamma(x, y)$, where

$$\gamma(x, y) := \frac{1}{2^{n/2}(2\pi)^{(n^2+n)/2}} \exp\left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2 - \frac{1}{4} \sum_{i=1}^n y_i^2\right)$$

is the Gaussian GOE density on $\Omega \oplus \Omega'$. Eventually, we would like to compare $w(x)$ to the density of $D(M)$, which we denote by $q(x)$. Since $q(x)$ is the push-forward of γ under D , we have that

$$q(x) = \frac{1}{(2\pi)^{(n^2-n)/2}} \exp\left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2\right)$$

which is just the standard Gaussian density on Ω .

Our next goal is thus to provide a formula for $w(x)$, which we do in a rather straightforward manner, by integration over fibers of the map $(D \circ f)^{-1}$. To this end, fix $x \in \Omega$. Observe that, by definition, D' is a bijection from $(D \circ f)^{-1}(x)$ to Ω' . In other words, for $y \in \Omega'$, defining

$$g(x, y) := \left(\frac{1}{\sqrt{d}}y + I_n\right)^{1/2} x \left(\frac{1}{\sqrt{d}}y + I_n\right)^{1/2},$$

it can easily be checked by the definition of f and g that

$$(D \circ f)^{-1}(x) = \{(D(g(x, y)), y); y \in \Omega'\}.$$

Define a mapping $T : \Omega \oplus \Omega' \rightarrow \Omega \oplus \Omega'$ by

$$T(x, y) := (D(g(x, y)), y).$$

Let $\varphi(\cdot)$ be a test function on Ω . Then, substituting $(x, y) = T(x', y')$ and noting that $D(f(T(x', y'))) = x'$, we have that

$$\begin{aligned} & \int_{\Omega} \varphi(x) w(x) dx \\ &= \int_{\Omega \times \Omega'} \varphi(D(f(x, y))) \gamma(x, y) dx dy = \int_{\Omega \times \Omega'} \varphi(D(f(T(x', y')))) \gamma(T(x', y')) J(x', y') dx' dy' \\ &= \int_{\Omega \times \Omega'} \varphi(x') \gamma(T(x', y')) J(x', y') dx' dy' = \int_{\Omega} \varphi(x') \left(\int_{\Omega'} \gamma(T(x', y')) J(x', y') dy' \right) dx' \end{aligned}$$

where $J(\cdot, \cdot)$ is the Jacobian of T . Since the above holds true for any measurable test function, recalling the definition of T we conclude that

$$w(x) = \int_{\Omega'} \gamma(D(g(x, y)), y) J(x, y) dy. \quad (50)$$

Let us now calculate the Jacobian $J(\cdot, \cdot)$. Recall that $T : \Omega \oplus \Omega' \rightarrow \Omega \oplus \Omega'$. Writing ∇T as a block matrix corresponding to this decomposition, we clearly have

$$\nabla T = \begin{pmatrix} \frac{\partial g}{\partial x} & \frac{\partial g}{\partial y} \\ 0 & I_n \end{pmatrix}.$$

Therefore, we have $J(x, y) = \det \left(\frac{\partial g}{\partial x} \right)$. Now by the definition of g we have that

$$\frac{\partial g_{i,j}}{\partial x_{i',j'}} = \mathbf{1}_{\{i=i', j=j'\}} \left(\frac{1}{\sqrt{d}} y_i + 1 \right)^{1/2} \left(\frac{1}{\sqrt{d}} y_j + 1 \right)^{1/2},$$

so we finally get that

$$J(x, y) = \prod_{1 \leq i < j \leq n} \left| \frac{1}{\sqrt{d}} y_i + 1 \right|^{1/2} \left| \frac{1}{\sqrt{d}} y_j + 1 \right|^{1/2} = \prod_{i=1}^n \left| 1 + \frac{y_i}{\sqrt{d}} \right|^{(n-1)/2}.$$

Together with the formula (50), this yields

$$\begin{aligned} w(x) &= \frac{1}{2^{n/2} (2\pi)^{(n^2+n)/2}} \int_{\Omega'} J(x, y) \exp \left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} g(x, y)_{i,j}^2 - \frac{1}{4} \sum_j y_j^2 \right) dy \\ &= \frac{1}{2^{n/2} (2\pi)^{(n^2+n)/2}} \int_{\Omega'} J(x, y) \exp \left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \left(1 + \frac{y_i + y_j}{\sqrt{d}} + \frac{y_i y_j}{d} \right) - \frac{1}{4} \sum_j y_j^2 \right) dy \\ &= \frac{1}{(2\pi)^{(n^2-n)/2}} \mathbb{E} \left[\prod_{i=1}^n \left| 1 + \frac{\Gamma_i}{\sqrt{d/2}} \right|^{(n-1)/2} \exp \left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \left(1 + \frac{\Gamma_i + \Gamma_j}{\sqrt{d/2}} + \frac{2\Gamma_i \Gamma_j}{d} \right) \right) \right], \end{aligned}$$

where $\Gamma = (\Gamma_1, \dots, \Gamma_n)$ is a standard Gaussian random vector. We can now estimate

$$\begin{aligned}
\frac{w(x)}{q(x)} &= \mathbb{E} \left[\prod_{i=1}^n \left| 1 + \frac{\Gamma_i}{\sqrt{d/2}} \right|^{(n-1)/2} \exp \left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \left(\frac{\Gamma_i + \Gamma_j}{\sqrt{d/2}} + \frac{2\Gamma_i \Gamma_j}{d} \right) \right) \right] \\
&\geq \mathbb{E} \left[\prod_{i=1}^n \left| 1 + \frac{\Gamma_i}{\sqrt{d/2}} \right|^{(n-1)/2} \exp \left(-\frac{1}{2} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \left(\frac{\Gamma_i + \Gamma_j}{\sqrt{d/2}} + \frac{2(\Gamma_i^2 + \Gamma_j^2)}{d} \right) \right) \right] \\
&= \prod_{i=1}^n \mathbb{E} \left[\left| 1 + \frac{\Gamma_i}{\sqrt{d/2}} \right|^{(n-1)/2} \exp \left(-\sum_{j \neq i} x_{i,j}^2 \left(\frac{\Gamma_i}{\sqrt{2d}} + \frac{\Gamma_i^2}{d} \right) \right) \right] \\
&\geq \prod_i \mathbb{E} \left[\left(1 + \frac{n}{2\sqrt{2d}} \Gamma_i \right) \left(1 - \frac{\sum_{j \neq i} x_{i,j}^2}{\sqrt{2d}} \Gamma_i - \frac{\sum_{j \neq i} x_{i,j}^2}{d} \Gamma_i^2 \right) \mathbf{1}_{\{|\Gamma_i| \leq \sqrt{d}/2\}} \right] \\
&= \prod_i \mathbb{E} \left[\left(1 + \frac{\frac{n}{2} - \sum_{j \neq i} x_{i,j}^2}{\sqrt{2d}} \Gamma_i - \frac{(n/4 + 1) \sum_{j \neq i} x_{i,j}^2}{d} \Gamma_i^2 - \frac{n \sum_{j \neq i} x_{i,j}^2}{2\sqrt{2d}^{3/2}} \Gamma_i^3 \right) \mathbf{1}_{\{|\Gamma_i| \leq \sqrt{d}/2\}} \right] \\
&= \prod_i \mathbb{E} \left[\left(1 - \frac{(n/4 + 1) \sum_{j \neq i} x_{i,j}^2}{d} \Gamma_i^2 \right) \mathbf{1}_{\{|\Gamma_i| \leq \sqrt{d}/2\}} \right].
\end{aligned}$$

A well-known estimate concerning the Gaussian distribution gives $\mathbb{P}(|\Gamma_1| > \sqrt{d}/2) < 2e^{-d/8}$. Consequently we have for all $1 \leq i \leq n$ that

$$\mathbb{E} \left[\left(1 - \frac{(n/4 + 1) \sum_{j \neq i} x_{i,j}^2}{d} \Gamma_i^2 \right) \mathbf{1}_{\{|\Gamma_i| \leq \sqrt{d}/2\}} \right] \geq 1 - 2e^{-d/8} - \frac{(n/4 + 1) \sum_{j \neq i} x_{i,j}^2}{d},$$

and the two last displays give that

$$\begin{aligned}
\frac{w(x)}{q(x)} &\geq \left(1 - 2ne^{-d/8} - \frac{n/4 + 1}{d} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \right) \mathbf{1}_{\left\{ \frac{(n/4+1) \sum_{j \neq i} x_{i,j}^2}{d} < \frac{1}{2}, \forall 1 \leq i \leq n \right\}} \\
&\geq \left(1 - 2ne^{-d/8} - \frac{n/4 + 1}{d} \sum_{1 \leq i < j \leq n} x_{i,j}^2 \right) \mathbf{1}_{\left\{ x_{i,j}^2 < \frac{d}{n^2}, \forall 1 \leq i, j \leq n \right\}}.
\end{aligned}$$

Finally, let $U = \{U_{i,j}\}_{1 \leq i < j \leq n}$ be a standard Gaussian vector on Ω . The above display gives that

$$\begin{aligned}
\text{TV}(D(f(M)), D(M)) &= \mathbb{E} \left| \frac{w(D(M))}{q(D(M))} - 1 \right| = 2\mathbb{E} \left(1 - \frac{w(U)}{q(U)} \right)_+ \\
&\leq 2ne^{-d/8} + \frac{n/4 + 1}{d} \mathbb{E} \left[\sum_{i,j} U_{i,j}^2 \right] + \mathbb{P} \left(\max_{1 \leq i < j \leq n} U_{i,j}^2 > \frac{d}{n^2} \right) \\
&\leq 2ne^{-d/8} + \frac{2n^3}{d} + 2n^2 e^{-\frac{d}{2n^2}},
\end{aligned}$$

which completes the proof. ■