

Latent Distance Estimation for Random Geometric Graphs *

Ernesto Araya Valdivia

*Laboratoire de Mathématiques d'Orsay (LMO)
Université Paris-Sud
91405 Orsay Cedex
France*

Yohann De Castro

*Ecole des Ponts ParisTech-CERMICS
6 et 8 avenue Blaise Pascal, Cité Descartes
Champs sur Marne, 77455 Marne la Vallée, Cedex 2
France*

Abstract: Random geometric graphs are a popular choice for a latent points generative model for networks. Their definition is based on a sample of n points X_1, X_2, \dots, X_n on the Euclidean sphere \mathbb{S}^{d-1} which represents the latent positions of nodes of the network. The connection probabilities between the nodes are determined by an unknown function (referred to as the “link” function) evaluated at the distance between the latent points. We introduce a spectral estimator of the pairwise distance between latent points and we prove that its rate of convergence is the same as the nonparametric estimation of a function on \mathbb{S}^{d-1} , up to a logarithmic factor. In addition, we provide an efficient spectral algorithm to compute this estimator without any knowledge on the nonparametric link function. As a byproduct, our method can also consistently estimate the dimension d of the latent space.

MSC 2010 subject classifications: Primary 68Q32; secondary 60F99, 68T01.

Keywords and phrases: Graphon model, Random Geometric Graph, Latent distances estimation, Latent position graph, Spectral methods.

1. Introduction

Random geometric graph (RGG) models have received attention lately as alternative to some simpler yet unrealistic models as the ubiquitous Erdős-Rényi model [11]. They are generative latent point models for graphs, where it is assumed that each node has associated a latent point in a metric space (usually the Euclidean unit sphere or the unit cube in \mathbb{R}^d) and the connection probability between two nodes depends on the position of their associated latent points. In many cases, the connection probability depends only on the distance between the latent points and it is determined by a one-dimensional “link” function.

Because of its geometric structure, this model is appealing for applications in wireless networks modeling [17], social networks [15] and biological networks [13], to name a few. In many of these real-world networks, the probability that a tie exists between two agents (nodes) depends on the similarity of their profiles. In other words, the connection probability depends on some notion of distance between the position of the agents in a metric space, which in the social network literature has been called the *social space*.

*Accepted to NeurIPS 2019

In the classical RGG model, as introduced by Gilbert in [12], we consider n independent and identically distributed latent points $\{X_i\}_{i=1}^n$ in \mathbb{R}^d and we construct the graph with vertex set $V = \{1, 2, \dots, n\}$, where the node i and j are connected if and only if the Euclidean distance $\|X_i - X_j\|_d$ is smaller than a certain predefined threshold τ . The classic reference on the classical RGG model, from the probabilistic point-view, is the monograph [24]. Another good reference is the survey paper [29]. In that case, the “link” function, which we have not yet formally defined, is the *threshold* function $\mathbb{1}_{t \leq \tau}(t)$. That is, the connection probability between two points is one or zero depending if their distance is smaller or larger than τ . In that case, all the randomness lies in the fact that we are sampling the latent points with a certain distribution. We choose to maintain the name of random geometric graphs for more general “link” functions.

We are interested in the problem of recovering the pairwise distances between the latent points $\{X_i\}_{i=1}^n$ for geometric graphs on the sphere \mathbb{S}^{d-1} given a single observation of the network. We limit ourselves to the case when the network is a simple graph. Furthermore, we will assume that the dimension d is fixed and that the “link” function is not known. This problem and some related ones have been studied for different versions of the model and under a different set of hypothesis, see for example the recent work [1] and the references therein. In that work the authors propose a method for estimating the latent distances based on the graph theoretic distance between two nodes (that is the length of the shortest path that starts in one node and finishes on the other). Independently, in [10] the authors develop a similar approach which has slightly less recovery error, but for a less general model. In both cases, the authors consider the cube in \mathbb{R}^d (or the whole \mathbb{R}^d) but not the sphere. Our strategy is similar to the one developed in [25], where they considered the latent point estimation problem in the case of *random dot product graphs*, which is a more restricted model compared to the one considered here. However, they considered more general Euclidean spaces and latent points distributions other than the uniform. Similar ideas have been used in the context of vertex classification for latent position graphs [26].

We will use the notion of graphon function to formalize the concept of “link” function. Graphons are central objects to the theory of dense graph limits. They were introduced by Lovász and Szegedy in [22] and further developed in a series of papers, see [4],[5]. Formally, they are symmetric kernels that take values in $[0, 1]$, thus they will act as the “link” function for the latent points. The spectrum of the graphon is defined as the spectrum of an associated integral operator, as in [21, Chap.7]. In this paper, they will play the role of limit models for the adjacency matrix of a graph, when the size goes to infinity. This is justified in light of the work of Koltchinskii and Giné [19] and Koltchinskii [19]. In particular, the adjacency matrix of the observed graph can be thought of as a finite perturbed version of this operator, combining results from [19] and [2].

We will focus on the case of dense graphs on the sphere \mathbb{S}^{d-1} where the connection probability depends only on the geodesic distance between two nodes. This allows us to use the harmonic analysis on the sphere to have a nice characterization of the graphon spectrum, which has a very particular structure. More specifically, the following two key elements hold: first, the basis of eigenfunctions is fixed (do not depend on the particular graphon considered) and equal to the well-known spherical harmonic polynomials. Second, the multiplicity of each eigenvalue is determined by a sequence of integers that depends only on the dimension d of the sphere and is given by a known formula and the associated eigenspaces are composed of spherical harmonics of the same polynomial degree.

The graphon eigenspace is composed only with linear eigenfunctions (harmonic polynomials

of degree one) will play an important role in the latent distances matrix recovery as all the information we need to reconstruct the distances matrix is contained in those eigenfunctions. We will prove that it is possible to approximately recover this information from the observed adjacency matrix of the graph under regularity conditions (of the Sobolev type) on the graphon and assuming an eigenvalue gap condition (similar hypotheses are made in [7] in the context of matrix estimation and in [20] in the context of manifold learning). We do this by proving that a suitable projection of the adjacency matrix, onto a space generated by exactly d of its eigenvectors, approximates well the latent distances matrix considering the mean squared error in the Frobenius norm. We give nonasymptotic bound for this quantity obtaining the same rate as the nonparametric rate of estimation of a function on the sphere \mathbb{S}^{d-1} . Our approach includes the adaptation of some perturbation theorems for matrix projections from the orthogonal to a “nearly” orthogonal case, which combined with concentration inequalities for the spectrum gives a probabilistic finite sample bound, which is novel to the best of our knowledge. Our method share some similarities with the celebrated UVST method, introduced by Chatterjee in [7], but in that case we obtain an estimator of the probability matrix described in Section 2.2 and not of the population Gram matrix as our method. We develop an efficient algorithm, which we call Harmonic EigenCluster(HEiC) to reconstruct the latent positions from the data and illustrate its usefulness with synthetic data.

2. Preliminaries

2.1. Notation

We will consider \mathbb{R}^d with the Euclidean norm $\|\cdot\|$ and the Euclidean scalar product $\langle \cdot, \cdot \rangle$. We define the sphere $\mathbb{S}^{d-1} := \{x \in \mathbb{R}^d : \|x\| = 1\}$. For a set $A \subset \mathbb{R}^d$ its diameter $\text{diam}(A) := \sup_{x,y \in A} |x - y|$ and if $B \subset \mathbb{R}^d$ the distance between A and B is $\text{dist}(A, B) := \inf_{x \in A, y \in B} |x - y|$. We will use $\|\cdot\|_F$ the Frobenius norm for matrices and $\|\cdot\|_{op}$ for the operator norm. The identity matrix in $\mathbb{R}^{d \times d}$ will be Id_d . If X is a real valued random variable and $\alpha \in (0, 1)$, $X \leq_\alpha C$ means that $\mathbb{P}(X \leq C) \geq 1 - \alpha$.

2.2. Generative model

We describe the generative model for networks which is a generalization of the classical random geometric graph model introduced by Gilbert in [12]. We base our definition on the W -random graph model described in [21, Sec. 10.1]. The central objects will be graphon functions on the sphere, which are symmetric measurable functions of the form $W : \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow [0, 1]$. Throughout this paper, we consider the measurable space $(\mathbb{S}^{d-1}, \sigma)$, where σ is the uniform measure on the sphere. On $\mathbb{S}^{d-1} \times \mathbb{S}^{d-1}$ we consider the product measure $\sigma \times \sigma$.

Now we describe how to generate a simple graph with n nodes from a graphon function W and a sample of points on the sphere $\{X_i\}_{i=1}^n$, known as the latent points.

First, we sample n points $\{X_i\}_{i=1}^n$ independently on the sphere \mathbb{S}^{d-1} , according to the uniform measure σ . These are the so-called latent points. Second, we construct the matrix of distances between these points, called the *Gram matrix* \mathcal{G}^* (we will often call it population Gram matrix) defined by

$$\mathcal{G}_{ij}^* := \langle X_i, X_j \rangle$$

and the so-called *probability matrix*

$$\Theta_{ij} = \rho_n W(X_i, X_j)$$

which is also a $n \times n$ matrix. The function W gives the precise meaning for the “link” function, because it determines the connection probability between X_i and X_j . The introduction of the scale parameter $0 < \rho_n \leq 1$ allow us to control the edge density of the sampled graph given a function W , see [18] for instance. The case $\rho_n = 1$ corresponds to the dense case (the parameter Θ_{ij} do not depend on n) and when $\rho_n \rightarrow 0$ the graph will be sparser. Our main results will hold in the regime $\rho_n = \Omega(\frac{\log n}{n})$, which we call *relatively sparse*. Most of the time we will work with the normalized version of the probability matrix $T_n := \frac{1}{\rho_n} \Theta$. If there exists a function $f : [-1, 1] \rightarrow [0, 1]$ such that $W(x, y) = f(\langle x, y \rangle)$ for all $x, y \in \mathbb{S}^{d-1}$ we will say that W is a geometric graphon.

Finally, we define the random adjacency matrix \hat{T}_n , which is a $n \times n$ symmetric random matrix that has independent entries (except for the symmetry constraint $\hat{T}_n = \hat{T}_n^T$), conditional on the probability matrix, with laws

$$n(\hat{T}_n)_{ij} \sim \mathcal{B}(\Theta_{ij})$$

where $\mathcal{B}(m)$ is the Bernoulli distribution with mean parameter m . Since the probability matrix contains the mean parameters for the Bernoulli distributions that define the random adjacency matrix it has been also called the *parameter matrix* [7]. Observe that the classical RGG model on the sphere is a particular case of the described W -random graph model when $W(x, y) = \mathbb{1}_{\langle x, y \rangle \leq \tau}$. In that case, since the entries of the probability matrix only have values in $\{0, 1\}$, the adjacency matrix and the probability matrix are equal. Depending on the context, we use \hat{T}_n for the random matrix as described above or for an instance of this random matrix, that is for the adjacency matrix of the observed graph. This will be clear from the context.

Thus the generative model can be seen as a two step sampling procedure where first the latent points are generated (which determine the Gram matrix and the probability matrix) and conditional on those points we generate the adjacency matrix.

It is worth noting that graphons can be, without loss of generality, defined in $[0, 1]^2$. The previous affirmation means that for any graphon there exists a graphon in $[0, 1]^2$ that generates the same distribution on graphs for any given number of nodes. However, in many cases the $[0, 1]^2$ representation can be less revealing than other representations using a different underlying space. This is illustrated in the case of the *prefix attachment* model in [21, example 11.41].

In the sequel we use the notation $\lambda_0, \lambda_1, \dots, \lambda_{n-1}$ for the eigenvalues of the normalized probability matrix T_n . Similarly, we denote by $\hat{\lambda}_0, \hat{\lambda}_1, \dots, \hat{\lambda}_{n-1}$ the eigenvalues of the matrix \hat{T}_n . We recall that T_n (resp. \hat{T}_n) and $\frac{1}{\rho_n} T_n$ (resp. $\frac{1}{\rho_n} \hat{T}_n$) have the same set of eigenvectors. We will denote by v_j for $1 \leq j \leq n$ the eigenvector of T_n associated to λ_j , which is also the eigenvector of $\frac{1}{\rho_n} T_n$ associated to $\frac{1}{\rho_n} \lambda_j$. Similarly, we denote by \hat{v}_j to the eigenvector associated to the eigenvalue $\rho_n \hat{\lambda}_j$ of \hat{T}_n .

Our main result is that we can recover the Gram matrix using the eigenvectors of \hat{T}_n as follows

Theorem 1 (Informal statement). *There exists a constant $c_1 > 0$ that depends only on the dimension d such that the following is true. Given a graphon W on the sphere such that $W(x, y) = f(\langle x, y \rangle)$ with $f : [-1, 1] \rightarrow [0, 1]$ unknown, which satisfies an eigenvalue gap condition and has Sobolev regularity s , there exists a subset of the eigenvectors of \hat{T}_n , such that $\hat{\mathcal{G}} := \frac{1}{c_1} \hat{V} \hat{V}^T$ converges to the population Gram matrix $\mathcal{G}^* := \frac{1}{n} (\langle X_i, X_j \rangle)_{i,j}$ at rate $n^{\frac{-s}{2s+d-1}}$ (up to a log factor). This estimate $\hat{V} \hat{V}^T$ can be found in linear time given the spectral decomposition of \hat{T}_n .*

We will say that a geometric graphon $W(x, y) = f(\langle x, y \rangle)$ on \mathbb{S}^{d-1} has regularity s if f belongs to the Weighted Sobolev space $Z_\gamma^s([-1, 1])$ with weight function $w_\gamma(t) = (1-t)^{\gamma-\frac{1}{2}}$, as defined in [23]. In order to make the statement of 1 rigorous, we need to precise the eigenvalue gap condition and define the graphon eigensystem.

2.3. Geometric graphon eigensystem

Here we gather some asymptotic and concentration properties for the eigenvalues and eigenfunctions of the matrices \hat{T}_n, T_n and the operator T_W , which allows us to recover the Gram matrix from data. The key fact is that the eigenvalues (resp. eigenvectors) of the matrix $\frac{1}{\rho_n} \hat{T}_n$ and $\frac{1}{\rho_n} T_n$ converge to the eigenvalues (resp. sampled eigenfunctions) of the integral operator $T_W : L^2(\mathbb{S}^{d-1}) \rightarrow L^2(\mathbb{S}^{d-1})$

$$T_W g(x) = \int_{\mathbb{S}^{d-1}} g(y) W(x, y) d\sigma(y)$$

which is compact [14, Sec.6, example 1] and self-adjoint (which follows directly from the symmetry of W). Then by a classic theorem in functional analysis [14, Sec.6, Thm. 1.8] its spectrum is a discrete set $\{\lambda_k^*\}_{k \in \mathbb{N}} \subset \mathbb{R}$ and its only accumulation point is zero. In consequence, we can see the spectra of \hat{T}_n, T_n and T_W (which we denote $\lambda(\hat{T}_n), \lambda(T_n)$ and $\lambda(T_W)$ resp.) as elements of the space \mathcal{C}_0 of infinite sequences that converge to 0 (where we complete the finite sequences with zeros). It is worth noting that in the case of geometric graphons with regularity s (in the Sobolev sense defined above) the rate of convergence of $\lambda(T_W)$ is determined by the regularity parameter s . We have the following:

- The spectrum of $\lambda(\frac{1}{\rho_n} T_n)$ converges to $\lambda(T_W)$ (almost surely) in the δ_2 metric, defined as follows

$$\delta_2(x, y) = \inf_{p \in \mathcal{P}} \sqrt{\sum_{i \in \mathbb{N}} (x_i - y_{p(i)})^2}$$

where \mathcal{P} is the set of all permutations of the non-negative integers. This is proved in [19].

- Matrices \hat{T}_n approach to matrix T_n in operator norm as n gets larger. Applying [2, Cor.3.3] to the centered matrix $Y = \hat{T}_n - T_n$ we get

$$\mathbb{E}(\|\hat{T}_n - T_n\|_{op}) \lesssim \frac{D_0}{n} + \frac{D_0^* \sqrt{\log n}}{n} \quad (1)$$

where \lesssim denotes inequality up to constant factors, $D_0 = \max_{0 \leq i \leq n} \sum_{j=1}^n Y_{ij}(1 - Y_{ij})$ and $D_0^* = \max_{i,j} |Y_{ij}|$. We clearly have that $D_0 = \mathcal{O}(n\rho_n)$ and $D_0^* \leq 1$, which implies that

$$\mathbb{E}\|\hat{T}_n - T_n\|_{op} \lesssim \max \left\{ \frac{\rho_n}{\sqrt{n}}, \frac{\sqrt{\log n}}{n} \right\}$$

We see that this inequality do not improve if ρ_n is smaller than in the relatively sparse case, that is $\rho_n = \Omega(\frac{\log n}{n})$. A similar bound can be obtained for the Frobenius norm replacing \hat{T}_n with \hat{T}_n^{uvst} the UVST estimator defined in [7]. For our main results, Proposition 3 and Theorem 4 the operator norm bound will suffice.

A remarkable fact in the case of geometric graphons on \mathbb{S}^{d-1} , that is when $W(x, y) = f(\langle x, y \rangle)$, is that the eigenfunctions $\{\phi_k\}_{k \in \mathbb{N}}$ of the integral operator T_W are a fixed set that

do not depend on the particular function f on consideration. This comes from the fact that T_W is a convolution operator on the sphere and its eigenfunctions are the well known *spherical harmonics* of dimension d , which are harmonic polynomials in d variables defined on \mathbb{S}^{d-1} corresponding to the eigenfunctions of the Laplace-Beltrami operator on the sphere. This follows from [8, Thm.1.4.5] and from the Funk-Hecke formula given in [8, Thm.1.2.9]. Let d_k denote the dimension of the k -th spherical harmonic space. It is well known [8, Cor.1.1.4] that $d_0 = 1$, $d_1 = d$ and $d_k = \binom{k+d-1}{k} - \binom{k+d-3}{k-2}$. Another important fact, known as the *addition theorem* [8, Lem.1.2.3 and Thm.1.2.6], is that

$$\sum_{i=d_{k-1}}^{d_k} \phi_j(x)\phi_j(y) = c_k G_k^\gamma(\langle x, y \rangle)$$

where G_k^γ are the Gegenbauer polynomials of degree k with parameter $\gamma = \frac{d-2}{2}$ and $c_k = \frac{2k+d-2}{d-2}$.

The Gegenbauer polynomial of degree one is $G_1^\gamma(t) = 2\gamma t$ (see [8, Appendix B2]), hence we have $G_1^\gamma(\langle X_i, X_j \rangle) = 2\gamma \langle X_i, X_j \rangle$ for every i and j . In consequence, by the addition theorem

$$G_1^\gamma(\langle X_i, X_j \rangle) = \frac{1}{c_1} \sum_{k=1}^d \phi_k(X_i)\phi_k(X_j)$$

where we recall that $d_1 = d$. This implies the following relation for the Gram matrix, observing that $2\gamma c_1 = d$

$$\mathcal{G}^* := \frac{1}{n} (\langle X_i, X_j \rangle)_{i,j} = \frac{1}{2\gamma c_1} \sum_{j=1}^d v_j^* v_j^{*T} = \frac{1}{d} V^* V^{*T} \quad (2)$$

where v_j^* is the \mathbb{R}^n vector with i -th coordinate $\phi_j(X_i)/\sqrt{n}$ and V^* is the matrix with columns v_j^* . In a similar way, we define for any matrix U in $\mathbb{R}^{n \times d}$ with columns u_1, u_2, \dots, u_d , the matrix $\mathcal{G}_U := \frac{1}{d} U U^T$. As part of our main theorem we prove that for n large enough there exists a matrix \hat{V} in $\mathbb{R}^{n \times d}$ where each column is an eigenvector of \hat{T}_n , such that $\hat{\mathcal{G}} := \mathcal{G}_{\hat{V}}$ approximates \mathcal{G}^* well, in the sense that the norm $\|\hat{\mathcal{G}} - \mathcal{G}^*\|_F$ converges to 0 at a rate which is that of the non-parametric estimation of a function on \mathbb{S}^{d-1} .

2.4. Eigenvalue gap condition

In this section we describe one of our main hypotheses on W needed to ensure that the space $\text{span}\{v_1^*, v_2^*, \dots, v_d^*\}$ can be effectively recovered with the vectors $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_d$ using our algorithm. Informally, we assume that the eigenvalue λ_1^* is sufficiently isolated from the rest of the spectrum of T_W . Given a geometric graphon W , we define the *spectral gap* of W relative to the eigenvalue λ_1^* by

$$\text{Gap}_1(W) := \min_{j \notin \{1, \dots, d_1\}} |\lambda_1^* - \lambda_j^*|$$

which quantifies the distance between the eigenvalue λ_1^* and the rest of the spectrum. In particular, we have the following elementary proposition.

Proposition 2. *It holds that $\text{Gap}_1(W) = 0$ if and only if there exists $j \notin \{1, \dots, d_1\}$ such that $\lambda_j^* = \lambda_1^*$ or $\lambda_1^* = 0$.*

Proof. Observe that the unique accumulation point of the spectrum of T_W is zero. The proposition follows from this observation. \square

To recover the population Gram matrix \mathcal{G}^* with our Gram matrix estimator $\hat{\mathcal{G}}$ we require the spectral gap $\Delta^* := \text{Gap}_1(W)$ to be different from 0. This assumption have been made before in the literature, in results that are bases in some versin of the Davis-Kahan $\sin\theta$ theorem (see for instance [7] , [20], [26]). More precisely, our results will hold on the following event

$$\mathcal{E} := \left\{ \delta_2 \left(\lambda \left(\frac{1}{\rho_n} T_n \right), \lambda(T_W) \right) \vee \frac{2^{\frac{9}{2}} \sqrt{d}}{\rho_n \Delta^*} \|T_n - \hat{T}_n\|_{op} \leq \frac{\Delta^*}{4} \right\},$$

for which we prove the following: given an arbitrary α we have that

$$\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\alpha}{2}$$

for n large enough (depending on W and α). The following results are the main results of this paper. Their proofs can be found in the supplementary material.

Proposition 3. *On the event \mathcal{E} , there exists one and only one set Λ_1 , consisting of d eigenvalues of \hat{T}_n , whose diameter is smaller that $\rho_n \Delta^*/2$ and whose distance to the rest of the spectrum of \hat{T}_n is at least $\rho_n \Delta^*/2$. Furthermore, on the event \mathcal{E} , our algorithm (Algorithm 1) returns the matrix $\hat{\mathcal{G}} = (1/c_1) \hat{V} \hat{V}^T$, where \hat{V} has by columns the eigenvectors corresponding to the eigenvalues on Λ_1 .*

Theorem 4. *Let W be a regular geometric graphon on \mathbb{S}^{d-1} , with regularity parameter s , such that $\Delta^* > 0$. Then there exists a set of eigenvectors $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_d$ of \hat{T}_n such that*

$$\|\mathcal{G}^* - \hat{\mathcal{G}}\|_F = O(n^{-\frac{s}{2s+d-1}})$$

where $\hat{\mathcal{G}} = \mathcal{G}_{\hat{V}}$ and \hat{V} is the matrix with columns $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_d$. Moreover, this rate is the minimax rate of non-parametric estimation of a regression function f with Sobolev regularity s in dimension $d - 1$.

The condition $\Delta^* > 0$ allow us to use Davis-Kahan type results for matrix perturbation to prove Theorem 4. With this and concentration for the spectrum we are able to control with high probability the terms $\|\hat{\mathcal{G}} - \mathcal{G}\|_F$ and $\|\mathcal{G} - \mathcal{G}^*\|_F$.

3. Algorithms

The Harmonic EigenCluster algorithm(HEiC) (see Algorithm 1 below) receives the observed adjacency matrix \hat{T}_n and the sphere dimension as its inputs to reconstruct the eigenspace associated to the eigenvalue λ_1^* . In order to do so, the algorithm selects d vectors in the set $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_n$, whose linear span is close to the span of the vectors $v_1^*, v_2^*, \dots, v_d^*$ defined in Section 2.3. The main idea is to find a subset of $\{\hat{\lambda}_0, \hat{\lambda}_2, \dots, \hat{\lambda}_{n-1}\}$, which we call Λ_1 , consisting on d_1 elements (recall that $d_1 = d$) and where all its elements are close to λ_1^* . This can be done assuming that the event \mathcal{E} defined above holds (which occurs with high probability). Once we have the set Λ_1 , we return the span of the eigenvectors associated to the eigenvalues in Λ_1 .

For a given set of indices i_1, \dots, i_d we define

$$\text{Gap}_1(\hat{T}_n; i_1, \dots, i_d) := \min_{i \notin \{i_1, \dots, i_d\}} \max_{j \in \{i_1, \dots, i_d\}} |\hat{\lambda}_j - \hat{\lambda}_i|$$

Algorithm 1: Harmonic EigenCluster(HEiC) algorithm

Input: (\hat{T}_n, d) adjacency matrix and sphere dimension
 $\Lambda^{\text{sort}} = \{\hat{\lambda}_1^{\text{sort}}, \dots, \hat{\lambda}_{n-1}^{\text{sort}}\} \leftarrow$ eigenvalues of \hat{T}_n sorted in decreasing order
 $\Lambda_1 \leftarrow \{\Lambda_1^{\text{sort}}, \dots, \Lambda_{1+d}^{\text{sort}}\}$: where Λ_i^{sort} is the i -th element in Λ^{sort}
Initialize $i = 2$, $\text{gap} = \text{Gap}_1(\hat{T}_n; 1, 2, \dots, d)$
while $i \leq n - d$ **do**
 if $\text{Gap}_1(\hat{T}_n; i, i+1, \dots, i+d) > \text{gap}$ **then**
 $\Lambda_1 \leftarrow \{\Lambda_i^{\text{sort}}, \dots, \Lambda_{i+d}^{\text{sort}}\}$
 end if
 $i = i + 1$
end while
Return: Λ_1, gap

and

$$\text{Gap}_1(\hat{T}_n) := \max_{\{i_1, \dots, i_d\} \in \mathcal{S}_d^n} \text{Gap}_1(\hat{T}_n; i_1, \dots, i_d)$$

where \mathcal{S}_d^n contains all the subsets of $\{1, \dots, n-1\}$ of size d . This definition parallels that of $\text{Gap}_1(W)$ for the graphon. Observe any set of indices in \mathcal{S}_d^n will not include 0. Otherwise stated, we can leave $\hat{\lambda}_0^{\text{sort}}$ out of this definition and it will not be candidate to be in Λ_1 . In the supplementary material we prove that the largest eigenvalue of the adjacency matrix will be close to the eigenvalue λ_0^* and in consequence can not be close enough to λ_1^* to be in the set Λ_1 , given the definition of the event \mathcal{E} .

To compute $\text{Gap}_1(\hat{T}_n)$ we consider the set of eigenvalues $\hat{\lambda}_j$ ordered in decreasing order. We use the notation $\hat{\lambda}_j^{\text{sort}}$ to emphasize this fact. We define the right and left differences on the sorted set by

$$\begin{aligned} \text{left}(i) &= |\hat{\lambda}_i^{\text{sort}} - \hat{\lambda}_{i-1}^{\text{sort}}| \\ \text{right}(i) &= \text{left}(i+1) \end{aligned}$$

where $\text{left}(\cdot)$ is defined for $1 \leq i \leq n$ and $\text{right}(\cdot)$ is defined for $0 \leq i \leq n-1$. With these definition, we have the following lemma, which we prove in the supplementary material.

Lemma 5. *On the event \mathcal{E} , the following equality holds*

$$\text{Gap}_1(\hat{T}_n) = \max \left\{ \max_{1 \leq i \leq n-d} \min \{ \text{left}(i), \text{right}(i+d) \}, \text{left}(n-d+1) \right\}$$

The set Λ_1 has the form $\Lambda_1 = \{\hat{\lambda}_{i^*}^{\text{sort}}, \hat{\lambda}_{i^*+1}^{\text{sort}}, \dots, \hat{\lambda}_{i^*+d}^{\text{sort}}\}$ for some $1 \leq i^* \leq n-d$. We have that either

$$i^* = \arg \max_{1 \leq i \leq n-d} \min \{ \text{left}(i), \text{right}(i+d) \}$$

or $i^* = n-d$ depending whether or not one has $\max_{1 \leq i \leq n-d} \min \{ \text{left}(i), \text{right}(i+d) \} > \text{left}(n-d+1)$. The algorithm then constructs the matrix \hat{V} having columns $\{\hat{v}_{i^*}, \hat{v}_{i^*+1}, \dots, \hat{v}_{i^*+d}\}$ and returns $\hat{V}\hat{V}^T$.

It is worth noting that Algorithm 1 time complexity $n^3 + n$, where n^3 comes from the fact that we compute the eigenvalues and eigenvectors of the $n \times n$ matrix \hat{T}_n and the linear term is because we explore the whole set of eigenvalues to find the maximum gap for the size d . In terms of space complexity the algorithm is n^2 because we need to store the matrix \hat{T}_n .

Remark 1. *If we change \hat{T}_n in the input of Algorithm 1 to \hat{T}_n^{usvt} (obtained by the UVST algorithm [7]) we predict that the algorithm will give similar results. This is because discarding some eigenvalues below a prescribed threshold do not have effect on our method. However, as preprocessing step the UVST might help in speeding up the eigenspace detection, but this step is already linear in time. The study of the effect of UVST as preprocessing step is left for future work.*

3.1. Estimation of the dimension d

So far we have focused on the estimation of the population Gram matrix \mathcal{G}^* . We now give an algorithm to find the dimension d , when it is not provided as input. This method receives the matrix \hat{T}_n as input and uses Algorithm 1 as a subroutine to compute a score, which is simply the value of the variable $\text{Gap}_1(\hat{T}_n)$ returned by Algorithm 1. We do this for each d in a set of candidates, which we call \mathcal{D} . This set of candidates will be usually fixed to $\{1, 2, 3, \dots, d_{\max}\}$. Once we have computed the scores, we pick the candidate that have the maximum score.

Given the guarantees provided by Theorem 4, the previously described procedure will find the correct dimension, with high probability (on the event \mathcal{E}), if the true dimension of the graphon is on the candidate set \mathcal{D} . This will happen, in particular, if the assumptions of Theorem 4 are satisfied. We recall that the main hypothesis on the graphon is that the spectral gap $\text{Gap}_1(W)$ should be different from 0.

4. Experiments

We generate synthetic data using different geometric graphons. In the first set of examples, we focus in recovering the Gram matrix when the dimension is provided. In the second set we tried to recover the dimension as well. The Python code of these experiments is provided in the supplementary material.

4.1. Recovering the Gram matrix

We start by considering the graphon $W_1(x, y) = \mathbb{1}_{\langle x, y \rangle \leq 0}$ which defines, through the sampling scheme given in Section 2.2, the same random graph model as the classical RGG model on \mathbb{S}^{d-1} with threshold 0. Thus two sampled points $X_i, X_j \in \mathbb{S}^{d-1}$ will be connected if and only if they lie in the same semisphere.

We consider different values for the sample size n and for each of them we sample 100 Gram matrices in the case $d = 3$ and run the Algorithm 1 for each. We compute each time the mean squared error, defined by

$$MSE_n = \frac{1}{n^2} \|\hat{\mathcal{G}} - \mathcal{G}^*\|_F^2$$

In Figure 1 we put the MSE_n for different values of n , showing how MSE_n decrease in terms of n . For each n , the MSE_n we plot is the mean over the 100 sampled graphs.

4.2. Recovering the dimension d

We conducted a simulation study using graphon W_1 , sampling 1000 point on the sphere of dimension $d = 3$ and we use Algorithm 1 to compute a score and recover d . We consider a

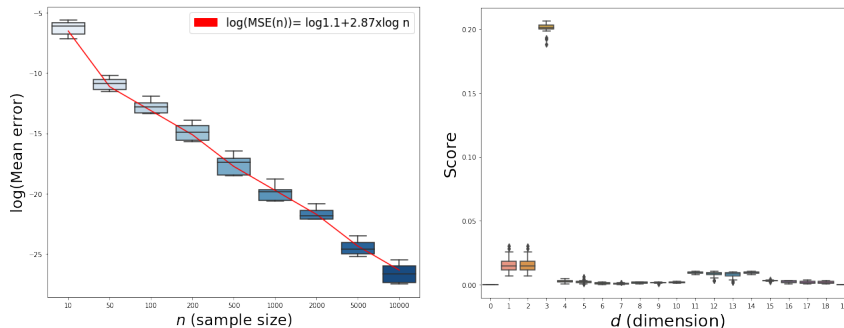


FIGURE 1. In the left we have a boxplot of MSE_n for different values of n . In the right, we plot the score for a set of candidate dimensions $\mathcal{D} = \{1, \dots, 15\}$. Data was sampled with W_1 on \mathbb{S}^{d-1} with $d = 3$.

set of candidates with $d_{max} = 15$. In Figure 1 we provide a boxplot for the score of each candidate repeating the procedure 50 times. We see that for this graphon, the algorithm can each time differentiate the true dimension from the “noise”. We include more experiments in the supplementary material.

5. Discussion

Although on this paper we have focused on the sphere as the latent metric space, our main result can be extended to other latent space where the distance is translation invariant, such as compact Lie groups or compact symmetric spaces. In that case, the geometric graphon will be of the form $W(x, y) = f(\cos \rho(x, y))$ where x, y are points in the compact Lie group \mathbb{S} and $\rho(\cdot, \cdot)$ is the metric in this space. We will have

$$f(\cos \rho(x, y)) = f(\cos \rho(x \cdot y^{-1}, e_1)) = \tilde{f}(x \cdot y^{-1})$$

where e_1 is the identity element in \mathbb{S} and $\tilde{f}(x) = f(\rho(x, e_1))$. In consequence $W(x, y) = \tilde{f}(x \cdot y^{-1})$. In addition, there exist an addition theorem in this case (which is central in our recovery result). Similar regularity notions to the one considered in this work also exist. They are related to rate of convergence to zero of the eigenvalues of integral operator associated to the graphon. In [9] the authors give more details on the model of geometric graphon in compact lie groups with focus on the estimation of the graphon function.

References

- [1] E Arias-Castro, A Channarond, B Pelletier, and N Verzelen. On the estimation of latent distances using graph distances. *arXiv:1804.10611*, 2018.
- [2] A Bandeira and R. Van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Annals of Probability*, 44(4):2479–2506, 2016.
- [3] R Bathia. *Matrix Analysis*. Springer Verlag New York, 1997.
- [4] C Borgs, J.T Chayes, L. Lovasz, V.T Sos, and K. Vesztergombi. Convergent sequences of dense graphs i. subgraph frequencies, metric properties and testing. *Adv. Math*, 219(6):1801–1851, 2008.
- [5] C Borgs, J.T Chayes, L. Lovasz, V.T Sos, and K. Vesztergombi. Convergent sequences of dense graphs ii. multiway cuts and statistical physics. *Annals of Mathematics*, 176(1):151–219, 2012.
- [6] M Braun. *Spectral Properties of the Kernel Matrix and their Relation to Kernel Methods in Machine Learning*. PhD thesis, Rheinische Friedrich-Wilhelms-Universität, Bonn, 2005.
- [7] S Chatterjee. Matrix estimation by universal singular value thresholding. *Annals of Statistics*, 43(1):177–214, 2015.
- [8] F. Dai and Y. Xu. *Approximation theory and harmonic Analysis on spheres and balls*. Springer Verlag Monographs in Mathematics, 2013.
- [9] Y. De Castro, C. Lacour, and T.M Pham Ngoc. Adaptive estimation of nonparametric geometric graphs. 2017.
- [10] J Diaz, C McDiarmid, and D Mitsche. Learning random points from geometric graphs or orderings. *arXiv:1804.10611*, 2018.
- [11] P Erdős and A Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci*, 5:17–60, 1960.
- [12] E.N Gilbert. Random plane networks. *J.Soc.Industrial Applied Mathematics*, 9(5):533–543, 1961.
- [13] D.J Higham, M Rasajski, and N Przulj. Fitting a geometric graph to a protein-protein interaction network. *Bioinformatics*, 24(8):1093–1099, 2008.
- [14] F Hirsch and G Lacombe. *Elements of functional analysis*. Springer-Verlag New York, 1999.
- [15] P Hoff, A Raftery, and M Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.
- [16] R Horn and C Johnson. *Matrix Analysis*. Cambridge University Press, 2012.
- [17] X Jia. Wireless networks and random geometric graphs. *Proc. Int. Symp. Parallel Architectures, Algorithms and Networks*, pages 575–579, 2004.
- [18] O Klopp, A Tsybakov, and N Verzelen. Oracle inequalities for network models and sparse graphon estimation. *Annals of Statistics*, 45(1):316–354, 2017.
- [19] V Koltchinskii and E Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, pages 113–167, 2000.
- [20] K Levin and V Lyzinski. Laplacian eigenmaps from sparse, noisy similarity measurements. *IEEE Transactions on Signal Processing*, 65:1998–2003, 2017.
- [21] L. Lovasz. *Large networks and graph limits*. Colloquium Publications (AMS), 2012.
- [22] L. Lovász and B. Szegedy. Limits of dense graph sequences. *J. Combin. Theory. Ser B*, 96(6):197–215, 2006.
- [23] S. Nicaise. Jacobi polynomials, weighted sobolev spaces and approximation results of some singularities. *Math. Nachr.*, 213:117–140, 2000.

- [24] M Penrose. *Random geometric graphs*. Oxford University Press, first edition, 2003.
- [25] D.L Sussman, M Tang, and C.E Priebe. Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 36:48–57, 2014.
- [26] M Tang, D.L Sussman, and C.E Priebe. Universally consistent vertex classification for latent position graphs. *Annals of Statistics*, 41:1406–1430, 2013.
- [27] R Vershynin. How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability*, 25:655–686, 2012.
- [28] R Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018.
- [29] M Walters. Random geometric graphs. *Surveys in Combinatorics*, pages 365–402, 2011.
- [30] Y Yu, T Wang, and R.J Samworth. A useful variant of the davis-kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

Appendix A: Graphon regularity

One way to define the regularity of a geometric graphon on \mathbb{S}^{d-1} is through the notion of weighted Sobolev spaces on the interval $[-1, 1]$. In that context, the regularity is related to the rate at which the eigenvalue sequence $\{\lambda_i^*\}_{i=0}^\infty$ tends to 0. Here we follow [23]. For a function of the form $f(t) = \sum_{k \geq 0} \mu_k c_k G_k^\gamma(t)$, we define the norm

$$\|f\|_{Z_\gamma^s}^2 = \sum_{k=0}^{\infty} d_k |\mu_k|^2 (1 + k(k + 2\gamma + 1))^s$$

We will say that f belongs to weighted Sobolev space Z_γ^s if $\|f\|_{Z_\gamma^s} \leq \infty$. We will refer to s as the regularity parameter. As in the case of classical Sobolev spaces, there is a definition of weighted Sobolev spaces that involves the integrability (with respect to the measure $w_\gamma(t)dt$) of the weak derivatives of a function. Then a function f belongs to Z_γ^s if it has s weak derivatives that are integrable with respect to the weighted L^2 norm in $[-1, 1]$ with weight $w_\gamma(t) = (1 - t)^{\gamma - \frac{1}{2}}$. In [23] the authors prove that both definitions are in fact equivalent.

Appendix B: Geometric Graphons have λ_0^* as the largest eigenvalue

To avoid border issues in Algorithm 1, we use the fact that the eigenvalue μ_0^* associated to the Gegenbauer polynomial $G_0^\gamma(t) = \mathbb{1}(t) := 1$ for $t \in [-1, 1]$ is the largest one, which in the notation of the paper can be written as $\lambda_0^{\text{sort}} = \mu_0^*$. This is true for all geometric graphons.

Lemma 6. *If $W : \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow [0, 1]$ is such that*

$$W(x, y) = f(\langle x, y \rangle)$$

for $f : [-1, 1] \rightarrow [0, 1]$, then

$$d_W(x) := \int_{\mathbb{S}^{d-1}} W(x, y) d\sigma(y)$$

is constant.

Proof. The proof follows from a change of variable. □

The following theorem is an analogous result to a classical theorem of spectral graph theory

Theorem 7. *For a graphon $W : \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow [0, 1]$ we have*

$$\int_{\mathbb{S}^{d-1} \times \mathbb{S}^{d-1}} W(x, y) d\sigma(x) d\sigma(y) \leq \lambda_0^* \leq \max_{x \in \mathbb{S}^{d-1}} d(x)$$

Proof. By Courant-Fisher min-max principle we have

$$\lambda_0^* = \max_{f \in L^2([-1, 1])} \frac{\langle T_W f, f \rangle}{\langle f, f \rangle}$$

In particular, if we take the function $\mathbb{1}(x) := 1$ for $x \in [-1, 1]$ we have

$$\begin{aligned} \lambda_0^* &\geq \frac{\langle T_W \mathbb{1}, \mathbb{1} \rangle}{\langle \mathbb{1}, \mathbb{1} \rangle} \\ &= \frac{\int_{\mathbb{S}^{d-1}} W(x, y) d\sigma(x) d\sigma(y)}{\int_{\mathbb{S}^{d-1}} d\sigma(y)} \\ &= d_W \end{aligned}$$

the last follows from the definition of d_W and the fact that σ is a probability measure on the sphere. On the other hand, if f_0 is an eigenfunction associated with λ_0 we can choose x^* such that $f_0(x^*) \geq f_0(x)$ for $x \in [-1, 1]$. Without loss of generality, assume that $f_0(x^*) \neq 0$. So

$$\begin{aligned}\lambda_0^* &= \frac{T_W f_0(x^*)}{f_0(x^*)} \\ &= \int_{\mathbb{S}^{d-1}} W(x^*, y) \frac{f_0(y)}{f_0(x^*)} d\sigma(y) \\ &\leq \int_{\mathbb{S}^{d-1}} W(x^*, y) d\sigma(y) \\ &= d_W(x^*)\end{aligned}$$

which finish the proof □

Since $G_0^\gamma(t) = \mathbb{1}(t)$ we have by Lemma 6 and Theorem 7 that the $\mu_0^* = \lambda_0^{\text{sort}}$.

Appendix C: Proof of the rate of convergence of the Algorithm

This section is devoted to the proof of the main theorem, Theorem 2.2. In the sequel, the sentence “ n large enough” means that n is bigger than some $n_0 \geq 1$ that may depend on W and α . Recall that the result obtained will hold with probability $1 - \alpha$ with $\alpha > 0$ arbitrarily small. We used through the paper, the notation $X \leq_\alpha C$, where X is random variable and C a constant, to indicate that the inequality holds with probability bigger than $1 - \alpha$.

The aim is to bound $\|\mathcal{G}^* - \hat{\mathcal{G}}\|_F$ and we will split it into two terms as follows

$$\|\mathcal{G}^* - \hat{\mathcal{G}}\|_F \leq \|\mathcal{G}^* - \mathcal{G}\|_F + \|\mathcal{G} - \hat{\mathcal{G}}\|_F$$

where the matrix \mathcal{G} will be defined later (see Proposition 8) using a subset of eigenvectors V of T_n . We will treat these terms separately starting with $\|\mathcal{G} - \hat{\mathcal{G}}\|_F$ in Section C.2 and the other term in Section C.3.

The first step is to control the probability of the following event \mathcal{E}

$$\mathcal{E} := \left\{ \delta_2 \left(\lambda \left(\frac{1}{\rho_n} T_n \right), \lambda(T_W) \right) \vee \frac{2^{\frac{9}{2}} \sqrt{d}}{\rho_n \Delta^*} \|T_n - \hat{T}_n\|_{op} \leq \frac{\Delta^*}{8} \right\},$$

where Δ^* is the spectral gap $\text{Gap}_1(W)$. We will prove in Section C.1 that this event holds with probability $1 - \alpha/2$ when n is large enough. This event ensures that the “noise level” is lower than the spectral gap $\text{Gap}_1(W)$ and it guarantees that our algorithm recovers the right subset of eigenvectors as will see in Proposition 8, Section C.1.

C.1. Event guaranteeing the algorithm convergence

Invoke Theorem 12 with $Y = \hat{T}_n - T_n$, which by definition have independent centered entries (conditional to latent points $\{X_i\}_{i=1}^n$), to obtain

$$\mathbb{P} \left(\|\hat{T}_n - T_n\|_{op} \geq \frac{3\sqrt{2D_0}}{n} + C_0 \frac{\sqrt{\log n/\alpha}}{n} \right) \leq \alpha$$

for $\alpha \in (0, 1/3)$. Note that for n large enough, one has

$$\|\hat{T}_n - T_n\|_{op} \leq_{\alpha/4} C \max \left\{ \sqrt{\frac{\rho_n}{n}}, \frac{\sqrt{\log n}}{n} \right\}$$

by Theorem 12, because $D_0 = \max_{0 \leq i \leq n} \sum_{j=1}^n \Theta_{ij}(1 - \Theta_{ij})$ is $\mathcal{O}(n\rho_n)$. Thus, for n large enough we have

$$\frac{1}{\rho_n} \|\hat{T}_n - T_n\|_{op} \leq_{\alpha/4} C \max \left\{ \frac{1}{\sqrt{\rho_n n}}, \frac{\sqrt{\log n}}{\rho_n n} \right\} \leq \frac{(\Delta^*)^2}{2^{\frac{17}{2}} \sqrt{d}}, \quad (3)$$

provided that $\frac{\sqrt{\log n}}{\rho_n n} = o(1)$, which is the case when $\rho_n = \Omega(\log n/n)$, which we have called the *relatively sparse* case. Let $V \in \mathbb{R}^{n \times d}$ and $\hat{V} \in \mathbb{R}^{n \times d}$ be two matrices with columns corresponding to the eigenvectors associated to eigenvalues $\lambda_{i_1}, \lambda_{i_2}, \dots, \lambda_{i_d}$ and $\hat{\lambda}_{i_1}, \hat{\lambda}_{i_2}, \dots, \hat{\lambda}_{i_d}$ of T_n and \hat{T}_n respectively, as in Theorem 13. We use Lemma 10 with $A = \hat{V}\hat{O}$ and $B = V$, where \hat{O} is an orthogonal matrix, and Theorem 13 assuming that the right hand side of (9) is smaller than 1, obtaining

$$\begin{aligned} \|\hat{V}\hat{V}^T - VV^T\|_F &\leq 2\|\hat{V}\hat{O} - V\|_F \\ &\leq \frac{2^{\frac{5}{2}} \min \{ \sqrt{d}\|T_n - \hat{T}_n\|_{op}, \|T_n - \hat{T}_n\|_F \}}{\Delta} \end{aligned} \quad (4)$$

where $\Delta := \text{dist}(\{\lambda_{i_1}, \dots, \lambda_{i_d}\}, \lambda(T_n) \setminus \{\lambda_{i_1}, \dots, \lambda_{i_d}\})$. Then we have

$$\begin{aligned} \|\hat{V}\hat{V}^T - VV^T\|_F &\leq \frac{2^{\frac{5}{2}} \frac{\sqrt{d}}{\rho_n} \|T_n - \hat{T}_n\|_{op}}{\frac{1}{\rho_n} \Delta} \\ &\leq_{\alpha} \frac{\rho_n (\Delta^*)^2}{2^6 \Delta} \end{aligned} \quad (5)$$

Now, we use the δ_2 metric to quantify the convergence of the eigenvalues of the normalized probability matrix $\frac{1}{\rho_n} T_n$ to the eigenvalues of the integral operator T_W . From Theorem 19 we have that, when n is large enough

$$\delta_2 \left(\lambda \left(\frac{1}{\rho_n} T_n \right), \lambda(T_W) \right) \leq_{\alpha/4} C \left(\frac{\log n}{n} \right)^{\frac{s}{2s+d-1}} \leq \frac{\Delta^*}{8}, \quad (6)$$

where Δ^* is the spectral gap $\text{Gap}_1(W)$. This and (3) ensure that \mathcal{E} has probability $1 - \alpha/2$. In particular, it gives the following result proving that our algorithm find the right eigenvectors.

Proposition 8. *On the event \mathcal{E} , there exists one and only one set Λ_1 of d eigenvalues of $\frac{1}{\rho_n} \hat{T}_n$ separated by at least $\Delta^*/2$ from the other eigenvalues of \hat{T}_n . These eigenvalues are at a distance at most $\Delta^*/8$ of $\frac{1}{\rho_n} \lambda_1, \dots, \frac{1}{\rho_n} \lambda_d$, the eigenvalues of T_n whose eigenvectors define the matrix $\mathcal{G} := (1/c_1)VV^T$. Furthermore, on the event \mathcal{E} , our algorithm returns the matrix $\hat{\mathcal{G}} = (1/c_1)\hat{V}\hat{V}^T$ composed by the eigenvectors corresponding to the eigenvalues of Λ_1 .*

Proof. When $\Delta^* > 0$, we remark that $\lambda_1^* = \lambda_2^* = \dots = \lambda_d^*$ is the only eigenvalue of T_W with multiplicity $d_1 = d$, the others eigenvalues (except for λ_0^*) having multiplicity strictly greater than d . Now, using (6) we deduce that there exists a unique set $\frac{1}{\rho_n} \lambda_{i_1}, \frac{1}{\rho_n} \lambda_{i_2}, \dots, \frac{1}{\rho_n} \lambda_{i_d}$ of

d eigenvalues of T_n that can be separated from the other eigenvalues by a distance at least $3\Delta^*/4$, namely the triangular inequality gives

$$\frac{\Delta}{\rho_n} \geq \frac{3\Delta^*}{4}. \quad (7)$$

To these eigenvalues correspond the eigenvectors $V \in \mathbb{R}^{n \times d}$ defining $\mathcal{G} := (1/c_1)VV^T$.

Furthermore, using (5) we get that there exists eigenvalues $\hat{\lambda}_{i_1}, \hat{\lambda}_{i_2}, \dots, \hat{\lambda}_{i_d}$ and eigenvectors $\hat{V} \in \mathbb{R}^{n \times d}$ of \hat{T}_n such that $\|\hat{V}\hat{V}^T - VV^T\|_F \leq \Delta^*/48$. We define $\Lambda_1 := \{\hat{\lambda}_{i_1}, \dots, \hat{\lambda}_{i_d}\}$. By Hoffman-Wielandt inequality [3, Thm.VI.4.1], it holds

$$\left(\sum_{k=1}^d (\hat{\lambda}_k^{\text{sort}} - \lambda_k^{\text{sort}})^2 \right)^{1/2} \leq \|\hat{V}\hat{V}^T - VV^T\|_F \leq \Delta^*/8,$$

where $\hat{\lambda}_1^{\text{sort}} \geq \hat{\lambda}_2^{\text{sort}} \geq \dots \geq \hat{\lambda}_d^{\text{sort}}$ (resp. $\lambda_1^{\text{sort}} \geq \lambda_2^{\text{sort}} \geq \dots \geq \lambda_d^{\text{sort}}$) is the sorted version of the eigenvalues $\hat{\lambda}_{i_1}, \dots, \hat{\lambda}_{i_d}$ (resp. $\lambda_{i_1}, \dots, \lambda_{i_d}$). By triangular inequality, we deduce that

$$\hat{\Delta} := \text{dist}(\Lambda_1, \lambda(\hat{T}_n) \setminus \Lambda_1) \geq \frac{\Delta^*}{2},$$

namely $\hat{\lambda}_{i_1}, \hat{\lambda}_{i_2}, \dots, \hat{\lambda}_{i_d}$ is a set of d eigenvalues at distance at least $\Delta^*/2$ from the other eigenvalues of \hat{T}_n .

This analysis can be also done for the other eigenvalues as follows. Eq. (6) shows that there exists a set of d_k eigenvalues of T_n which concentrate around μ_k^* , and such that it has diameter smaller than $\Delta^*/4$. Recall that d_k is the size of the Spherical Harmonics space k and $d_k > d_1 = d$. Weyl's inequality [3, P.63] shows that there exists a set Λ_k of d_k eigenvalues of \hat{T}_n around μ_k^* of size $\Delta^*/4$. Now, consider a subset L of d eigenvalues which is different from Λ_1 then the previous discussion shows that there exists an eigenvalue $\hat{\lambda}$ which is not in L and that belongs to same cluster to one of the eigenvalues in L . In particular $\hat{\lambda}$ is at a distance less than $\Delta^*/4$ of L . By (7) we deduce that, on the event \mathcal{E} , Algorithm 1 returns $\hat{\mathcal{G}} = (1/c_1)\hat{V}\hat{V}^T$ composed by the eigenvectors corresponding to the eigenvalues of the aforementioned cluster of d eigenvalues. \square

We now prove the following lemma, which is stated in the article

Lemma 9. *On the event \mathcal{E} , the following equality holds*

$$\text{Gap}_1(\hat{T}_n) = \max \left\{ \max_{1 \leq i \leq n-d} \min \{ \text{left}(i), \text{right}(i+d) \}, \text{left}(n-d+1) \right\}$$

Proof. The lemma follows from Proposition 8. Indeed, on the event \mathcal{E} there exist only one set Λ_1 of eigenvalues of \hat{T}_n with cardinality d , whose distance to the rest of the spectrum is larger than Δ^* and its diameter is smaller than Δ^* . When sorting the eigenvalues of \hat{T}_n in decreasing order, those belonging to Λ_1 will appear in consecutive order. The lemma follows from this observation and from the fact $\text{Gap}_1(\hat{T}_n; i_{n-d-1}, \dots, i_{n-1}) = \text{left}(n-d-1)$. \square

C.2. Sampling error control

We have by (4) that

$$\|\hat{\mathcal{G}} - \mathcal{G}\|_F = \frac{1}{c_1} \|\hat{V}\hat{V}^T - VV^T\|_F \leq \alpha C \frac{(d-2)}{\sqrt{dn}}, \quad (8)$$

whenever n is large enough and $\Delta^* > 0$, where C may depend on W . In the last inequality we used that $c_1 = d/(d-2)$.

C.3. Sampled eigenvectors convergence

We are left to control $\|\mathcal{G}^* - \mathcal{G}\|_F$. We begin by recalling some basic definitions we have made through the paper and introducing some notation. Set $R = \mathcal{O}((n/\log n)^{\frac{1}{2s+d-1}})$ and $\tilde{R} := d_0 + d_1 + \dots + d_R$ the total size of the $R+1$ first Harmonic spaces. It is well known that $\tilde{R} = \mathcal{O}(R^{d-1}) = o(n)$ for $s > 0$. If W_R is the rank R' approximation of W , we have

$$T_R = \left(\frac{1}{n}W_R(X_i, X_j)\right)_{i,j} = \Phi_{0,R}\Lambda_{0,R}^*\Phi_{0,R}$$

where $\Phi_{0,R}$ is the matrix with columns $\Phi_k \in \mathbb{R}^n$, for $0 \leq k \leq R'$, such that $(\Phi_k)_i = \phi_k(X_i)$ and $\Lambda_{0,R}^* = \text{diag}(\lambda_0^*, \lambda_2^*, \dots, \lambda_{\tilde{R}}^*)$. Similarly $\Lambda_{0,R} = \text{diag}(\lambda_0, \lambda_2, \dots, \lambda_{\tilde{R}})$. Let \tilde{V} be the matrix that contains as columns the eigenvectors of the matrix T_n and \tilde{V}_R contains as columns the eigenvectors T_R so we have the eigenvalue decomposition

$$T_n = \tilde{V}\Lambda\tilde{V}^T$$

$$T_R = \tilde{V}_R\Lambda_R\tilde{V}_R^T$$

Let V be the matrix that contains the columns $1, \dots, d$ of \tilde{V} , V_R contains the columns $1, \dots, d$ of \tilde{V}_R and V^* contains ϕ_k for $1 \leq k \leq d$ as columns. Then $\mathcal{G}^*, \mathcal{G}, \mathcal{G}_R, \mathcal{G}_{proj}^*$ are defined by

$$\mathcal{G}^* := \frac{1}{c_1}V^*(V^*)^T$$

$$\mathcal{G} := \frac{1}{c_1}VV^T$$

$$\mathcal{G}_R := \frac{1}{c_1}V_RV_R^T$$

$$\mathcal{G}_{proj}^* := V^*(V^{*T}V^*)^{-1}V^{*T}$$

Note that \mathcal{G}_{proj}^* is the projection matrix for the column span of the matrix V^* , that is, it is the projection matrix onto the space $\text{span}\{\Phi_1, \dots, \Phi_d\}$.

We have by triangle inequality

$$\|\mathcal{G}^* - \mathcal{G}\|_F \leq \|\mathcal{G}^* - \mathcal{G}_{proj}^*\|_F + \|\mathcal{G}_{proj}^* - \mathcal{G}_R\|_F + \|\mathcal{G}_R - \mathcal{G}\|_F$$

We call truncation error to the last term in the right hand side, because it is related to the fact that W_R is a rank R' approximation of W .

To bound $\|\mathcal{G} - \mathcal{G}_R\|_F$ we will use Theorem 13 noting that \mathcal{G} and \mathcal{G}_R have as columns the eigenvectors of matrices T_n and T_R . So

$$\|\mathcal{G} - \mathcal{G}_R\|_F \leq \frac{2^{\frac{3}{2}}\|T_n - T_R\|_F}{\Delta} \leq C \frac{(n/\log n)^{-s/(2s+d-1)}}{\Delta}$$

where we recall that $R = \mathcal{O}((n/\log n)^{\frac{1}{2s+d-1}})$, which gives the optimal rate for this error term [9]. In order to bound $\|\mathcal{G}^* - \mathcal{G}_{proj}^*\|_F$ we use Lemma 11 with $B = V^*$ obtaining

$$\|\mathcal{G}^* - \mathcal{G}_{proj}^*\|_F \leq \|\text{Id}_d - V^{*T}V^*\|_F$$

On the other hand, we have

$$\begin{aligned} \|\text{Id}_d - V^{*T}V^*\|_F &\leq \sqrt{d}\|\text{Id}_d - V^{*T}V^*\|_{op} \\ &\leq \alpha \frac{d}{\sqrt{n}} \end{aligned}$$

where we used Theorem 20 to obtain the last inequality.

It only remains to bound the term $\|\mathcal{G}_{proj}^* - \mathcal{G}_R\|_F$. We concentrate first in bounding the term $\mathcal{G}_{proj}^*\mathcal{G}_R^\perp$. We use Theorem 15, with $E = \mathcal{G}_{proj}^*$, $F = \mathcal{G}_R^\perp$, $B = T_R$ and $A = T_R + H$, where

$$H := \tilde{\Phi}_{0,R}\Lambda_{0,R}^*\tilde{\Phi}_{0,R}^T - \Phi_{0,R}\Lambda_{0,R}^*\Phi_{0,R}$$

the matrix $\tilde{\Phi}_{0,R}$ has column $\tilde{\Phi}_k$ for $k \in \{1, \dots, R'\}$ where the $\tilde{\Phi}_k$ are obtained from Φ_k by a Gram-Schmidt orthonormalization process. In other words, there exists a matrix L such that $\tilde{\Phi}_{0,R} = \Phi_{0,R}(L^{-1})^T$. The matrix L comes from the Cholesky decomposition of $\Phi_{0,R}^T\Phi_{0,R}$, that is, L satisfy $\Phi_{0,R}^T\Phi_{0,R} = LL^T$.

Note that A and B are symmetric, hence normal matrices, so Theorem 15 applies. Also, in the event \mathcal{E} , we can take $S_1 = (\lambda_1 - \frac{\Delta^*}{8}, \lambda_1 + \frac{\Delta^*}{8})$ and $S_2 = \mathbb{R} \setminus (\lambda_1 - \frac{7\Delta^*}{8}, \lambda_1 + \frac{7\Delta^*}{8})$. By Theorem 15 we have

$$\|\mathcal{G}_{proj}^*\mathcal{G}_R^\perp\|_F \leq \frac{\|A - B\|_F}{\Delta^*} = \frac{\|H\|_F}{\Delta^*}$$

where $\Delta := \min_{k,\ell \neq 1, \dots, d} \{|\lambda_k^* - \lambda_1^*|, |\lambda_d^* - \lambda_\ell^*|\}$. It remains to bound H .

We have that

$$\begin{aligned} \|H\|_F &\leq \|L^{-T}\Lambda_{0,R}^*L^{-1} - \Lambda_{0,R}^*\|_F \|\Phi_{0,R}^T\Phi_{0,R}\|_{op} \\ &\leq \|\Lambda_{0,R}^*\|_F \|L^{-1}L^{-T} - \text{Id}_{R'}\|_{op} \|\Phi_{0,R}^T\Phi_{0,R}\|_{op} \end{aligned}$$

where in the last line we used Corollary 17. It is easy to see that

$$\|L^{-1}L^{-T} - \text{Id}_{R'}\|_{op} = \|(\Phi_{0,R}^T\Phi_{0,R})^{-1} - \text{Id}_{R'}\|_{op}$$

which, using [9, Lem.12], implies that

$$\|Z\|_F \leq_{\alpha/4} 2C_1 \frac{R^{d-1}}{\sqrt{n}}$$

which, since $R = \mathcal{O}((n/\log n)^{\frac{1}{2s+d-1}})$, becomes

$$\|Z\|_F \leq_{\alpha/4} C' \left(\frac{\log n}{n}\right)^{\frac{s}{2s+d-1}}$$

for a constant $C' > 0$. Collecting terms we obtain

$$\|\mathcal{G}_{proj}^*\mathcal{G}_R^\perp\|_F \leq_{\alpha/4} \frac{C''}{\Delta^*} \left(\frac{\log n}{n}\right)^{\frac{s}{2s+d-1}}$$

Since \mathcal{G}_{proj}^* and \mathcal{G}_R are projectors we have, see [3, p.202]

$$\|\mathcal{G}_{proj}^* - \mathcal{G}_R\|_F = 2\|\mathcal{G}_{proj}^*\mathcal{G}_R^\perp\|_F$$

which implies that

$$\|\mathcal{G}_{proj}^* - \mathcal{G}_R\|_F \leq \alpha/4 \frac{2C''}{\Delta^*} \left(\frac{n}{\log n}\right)^{\frac{-s}{2s+d-1}}$$

To conclude, we have that

$$\begin{aligned} \|\text{Id}_d - V^{*T}V^*\|_F &\leq \sqrt{d} \|\text{Id}_d - V^{*T}V^*\|_{op} \\ &\leq \alpha/4 \frac{d}{\sqrt{n}} \end{aligned}$$

where we use Theorem 20 in the second inequality. Collecting terms we conclude that

$$\|\mathcal{G}^* - \mathcal{G}\|_F \leq \alpha/4 \frac{C_d}{\Delta^*} \left(\frac{\log n}{n}\right)^{\frac{s}{2s+d-1}}$$

where C_d is a constant that depends on d and α .

Appendix D: Useful results

Lemma 10. *Let A, B be two matrices in $\mathbb{R}^{n \times d}$ then*

$$\begin{aligned} \|AA^T - BB^T\|_F &\leq (\|A\|_{op} + \|B\|_{op})\|A - B\|_F \\ \|AA^T - BB^T\|_{op} &\leq (\|A\|_{op} + \|B\|_{op})\|A - B\|_{op}. \end{aligned}$$

If it holds that $A^T A = B^T B = I_d$ then

$$\|AA^T - BB^T\|_F \leq 2\|A - B\|_F$$

Proof. We begin with the first inequality

$$\begin{aligned} \|AA^T - BB^T\|_F &= \|(A - B)A^T + B(A^T - B^T)\|_F \\ &\leq \|(A \otimes I_n)\text{vec}(A - B)\|_2 + \|(I_d \otimes B)\text{vec}(A - B)^T\|_2 \\ &\leq (\|A \otimes I_n\|_{op} + \|I_d \otimes B\|_{op})\|A - B\|_F \\ &= (\|A\|_{op} + \|B\|_{op})\|A - B\|_F. \end{aligned}$$

Here $\text{vec}(\cdot)$ represent the vectorization of a matrix, that its transformation into a column vector. The second inequality is given by

$$\begin{aligned} \|AA^T - BB^T\|_{op} &= \|(A - B)A^T + B(A^T - B^T)\|_{op} \\ &\leq (\|A\|_{op} + \|B\|_{op})\|A - B\|_{op}. \end{aligned}$$

The third statement is an elementary consequence of the above inequalities. \square

Lemma 11. *Let B a $n \times d$ matrix with full column rank. Then we have*

$$\|BB^T - B(B^T B)^{-1}B^T\|_F = \|\text{Id}_d - B^T B\|_F$$

Proof. We have

$$\|BB^T - B(B^T B)^{-1}B^T\|_F = \|B((B^T B)^{-1} - \text{Id}_d)B^T\|_F$$

and by definition of the Frobenious norm and cyclic property of the trace

$$\begin{aligned} \|B((B^T B)^{-1} - \text{Id}_d)B^T\|_F^2 &= \text{tr}(B((B^T B)^{-1} - \text{Id}_d)B^T B((B^T B)^{-1} - \text{Id}_d)B^T) \\ &= \text{tr}((\text{Id}_d - B^T B)^2) \\ &= \|\text{Id}_d - B^T B\|_F^2 \end{aligned}$$

□

D.1. Bandeira-Van Handel theorem

The following theorem is a slight reformulation of the [2, Cor.3.12]

Theorem 12 (Bandeira-Van Handel). *Let Y be a $n \times n$ symmetric random matrix whose entries Y_{ij} are independent centered random variables. There exists a universal constant C_0 such that for $\alpha \in (0, 1)$*

$$\mathbb{P}\left(\|Y\|_{op} \geq 3\sqrt{2D_0} + C_0\sqrt{\log n/\alpha}\right) \leq \alpha$$

where $D_0 = \max_{0 \leq i \leq n} \sum_{j=1}^n Y_{ij}(1 - Y_{ij})$.

Proof. By [2, Rmk.3.13] we have the tail concentration bound (taking their ϵ equal to $1/2$)

$$\mathbb{P}\left(\|Y\|_{op}\right) \geq 3\sqrt{2D_0} + \max_{ij} |Y_{ij}| C_0\sqrt{\log n/\alpha}$$

the result follows, because $\max_{ij} |Y_{ij}| \leq 1$. □

Using the previous theorem with $Y = \hat{T}_n - T_n$, which is centered and symmetric, we obtain the tail bound

$$\mathbb{P}\left(\|\hat{T}_n - T_n\|_{op} \geq \frac{3\sqrt{2D_0}}{n} + C_0\frac{\sqrt{\log n/\alpha}}{n}\right) \leq \alpha$$

D.2. Davis-Kahan $\sin \theta$ theorem

For n large enough, the eigenspace associated to the eigenvalue $\hat{\lambda}_1$ is close to the eigenspace associated to the eigenvalue λ_1 . This is precised by the Davis-Kahan *sin θ* theorem. We use the following version which is proved in [30]

Theorem 13. *Let Σ and $\hat{\Sigma}$ be two symmetric $\mathbb{R}^{n \times n}$ matrices with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_n$ respectively. For $1 \leq r \leq s \leq n$ fixed, we assume that $\min\{\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1}\} > 0$ where $\lambda_0 := \infty$ and $\lambda_{n+1} = -\infty$. Let $d = s - r + 1$ and V and \hat{V} two matrices in $\mathbb{R}^{n \times d}$ with columns $(v_r, v_{r+1}, \dots, v_s)$ and $(\hat{v}_r, \hat{v}_{r+1}, \dots, \hat{v}_s)$ respectively, such that $\Sigma v_j = \lambda_j v_j$ and $\hat{\Sigma} \hat{v}_j = \hat{\lambda}_j \hat{v}_j$. Then there exists an orthogonal matrix \hat{O} in $\mathbb{R}^{d \times d}$ such that*

$$\|\hat{V}\hat{O} - V\|_F \leq \frac{2^{3/2} \min\{\sqrt{d}\|\Sigma - \hat{\Sigma}\|_{op}, \|\Sigma - \hat{\Sigma}\|_F\}}{\min\{\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1}\}} \quad (9)$$

Also, we need the following perturbation result [3, Thm.VII.2.8]

Theorem 14. *Let A and B two the normal matrices and define $\delta = \text{dist}(\lambda(A), \lambda(B))$. If X satisfies the Sylvester equation $AX - XB = Y$, then*

$$\|X\|_F \leq \frac{1}{\delta} \|Y\|_F$$

Another useful perturbation theorem [3, Thm.VII.3.1]

Theorem 15. *Let A and B be two normal operators and S_1 and S_2 two sets separated by a strip of size δ . Let E be the orthogonal projection matrix of the eigenspaces of A with eigenvalues inside S_1 and F be the orthogonal projection matrix of the eigenspaces of B with eigenvalues inside S_2 . Then*

$$\|EF\|_F \leq \frac{1}{\delta} \|E(A - B)F\|_F \leq \frac{1}{\delta} \|A - B\|_F$$

D.3. Ostrowski theorem

The following eigenvalue perturbation theorem is due to Ostrowski [16, Thm.4.5.9] and [6, Cor.3.54]

Theorem 16. *Let $A \in \mathbb{R}^{n \times n}$ be a Hermitian matrix and $S \in \mathbb{R}^{n \times n}$ be a nonsingular matrix. Then for each $1 \leq i \leq n$ there exists $\theta_i > 0$ such that*

$$\lambda_i(SAS^*) = \theta_i \lambda_i(A)$$

In addition, it holds

$$|\lambda_i(SAS^*) - \lambda_i(A)| \leq |\lambda_i(A)| \|S^*S - \text{Id}_n\|_{op}$$

Remark 2. *The previous theorem is also valid for S singular [16, Cor.4.5.11].*

The previous theorem can be extended to the case where S is not necessarily a square matrix [6, Cor.3.59]

Corollary 17. *Let $A \in \mathbb{R}^{n \times n}$ be a Hermitian matrix and $S \in \mathbb{R}^{d \times n}$ matrix then*

$$|\lambda_i(SAS^*) - \lambda_i(A)| \leq |\lambda_i(A)| \|S^*S - \text{Id}_n\|_{op}$$

From the previous result we deduce the following corollary

Corollary 18. *Under the same conditions of Corollary 17 we have*

$$\|SAS^* - A\|_F \leq \|A\|_F \|S^*S - \text{Id}_n\|_{op}$$

D.4. Convergence rate of regular graphon estimation

We use the following result, which can be found in [9]

Theorem 19. *Let W be a graphon on the sphere of the form $W(x, y) = f(\langle x, y \rangle)$. If f belongs to the weighted Sobolev space $Z_{w_\gamma}^s((-1, 1))$ then we have*

$$\delta_2\left(\lambda\left(\frac{1}{\rho_n} T_n\right), \lambda(T_W)\right) \leq_\alpha C \left(\frac{\log n}{n}\right)^{\frac{s}{2s+d-1}}$$

where \leq_α means that the inequality holds with probability greater than $1 - \alpha$ for $\alpha \in (0, 1/3)$ and n large enough.

While Theorem 12 gives a bound for the difference of the eigenvalues of the observed matrix with respect to the eigenvalues of the probability matrix, Proposition 19 ensures that the eigenvalues of the empirical matrix are close to these of the integral operator.

D.5. Covariance matrix approximation

Given a set of independent random vectors X_1, \dots, X_n uniformly distributed on the sphere \mathbb{S}^{d-1} we are interested in the concentration properties of the quantity $\frac{1}{n} \sum_{k=1}^n X_k X_k^T$ around its mean, which is $\mathbb{E}(X_i X_i^T) = \text{Id}_d$ for $1 \leq i \leq n$ (in other words, the vectors X_i are isotropic). Since the uniform distribution on the sphere is sub-gaussian [28, Thm.3.4.6], we can use the following theorem [27, Prop.2.1].

Theorem 20. *If X_1, \dots, X_n are independent random vectors in \mathbb{R}^d with $d \leq n$ which have sub-gaussian distribution. Then for any $\alpha \in (0, 1)$ it holds*

$$\left\| \frac{1}{n} \sum_{k=1}^n X_k X_k^T - \text{Id}_d \right\|_{op} \leq \alpha \sqrt{\frac{d}{n}}$$