

Community Recovery in the Geometric Block Model

Sainyam Galhotra

SG@CS.CORNELL.EDU

*Department of Computer Science
Cornell University
Bill and Melinda Gates Hall,
Ithaca, NY 14850 USA*

Arya Mazumdar

ARYA@UCSD.EDU

*Halicioglu Data Science Institute
University of California, San Diego
9500 Gilman Dr
La Jolla, CA 92093 USA*

Soumyabrata Pal

SOUMYABRATA@GOOGLE.COM

*Google Research
Bengaluru 560016, India*

Barna Saha

BARNAS@UCSD.EDU

*Department of Computer Science and Halicioglu Data Science Institute
University of California, San Diego
9500 Gilman Dr
La Jolla, CA 92093 USA*

Editor: Alexandre Proutiere

Abstract

To capture the inherent geometric features of many community detection problems, we propose to use a new random graph model of communities that we call a *Geometric Block Model*. The geometric block model builds on the *random geometric graphs* (Gilbert, 1961), one of the basic models of random graphs for spatial networks, in the same way that the well-studied stochastic block model builds on the Erdős-Rényi random graphs. It is also a natural extension of random community models inspired by the recent theoretical and practical advancements in community detection. To analyze the geometric block model, we first provide new connectivity results for *random annulus graphs* which are generalizations of random geometric graphs. The connectivity properties of geometric graphs have been studied since their introduction, and analyzing them has been more difficult than their Erdős-Rényi counterparts due to correlated edge formation.

We then use the connectivity results of random annulus graphs to provide necessary and sufficient conditions for efficient recovery of communities for the geometric block model. We show that a simple triangle-counting algorithm to detect communities in the geometric block model is near-optimal. For this we consider the following two regimes of graph density.

In the regime where the average degree of the graph grows logarithmically with the number of vertices, we show that our algorithm performs extremely well, both theoretically and practically. In contrast, the triangle-counting algorithm is far from being optimum for the stochastic block model in the logarithmic degree regime. We simulate our results on both real and synthetic datasets to show superior performance of both the new model as well as our algorithm.

Keywords: Random graphs, Community recovery, Generative model, Graph clustering, Random geometric graphs.

1. Introduction

Clustering of graphs is a ubiquitous problem where the objective is to partition the vertices of a graph into disjoint clusters such that each cluster is more densely connected internally than across clusters. Many real-world datasets can be represented in the form of graphs where the vertices (nodes) represent data elements and the edges represent (noisy) interaction between the elements. It is of primary interest to recover latent clusters in the graph either as an end goal or as a means to some other learning problem. As an example, consider the graph induced by the collection of political blogs in internet where each vertex corresponds to a blog website, and an edge exists between two vertices if one of the corresponding blogs hyperlinks to the other. From this graph, it might be of interest to partition the blogs into two clusters. Achieving this objective allows a recommendation engine to recommend each blog to the correct audience or suggest correct political advertisements and thereby garner more views.

A very simple algorithm to perform this task is the well-known *min-cut* algorithm which intends to partition the graph into clusters such that minimum number of edges go across the clusters. The min-cut algorithm runs in polynomial time but it does not have any guarantee on the sizes of the clusters returned. If we constrain the output clusters to be equally sized, then this problem, also known as the *min-bisection* problem becomes NP-hard. However, most real-world datasets are not pathological and there exist many properties which are satisfied by the graph, for example, sparsity and transitivity. Hence, an approach to resolve this problem is to devise a simple and elegant modeling assumption according to which the observed real-world graphs are generated. Under this assumption, we can design efficient algorithms that can recover the latent clusters. Such modeling assumptions not only allow a rigorous theoretical treatment, but also allow us to benchmark and compare different heuristics for graph partitioning. Finally, note that these models must capture the inherent properties of real world graphs so that the algorithms designed under the corresponding assumption work well on real world datasets as well.

The *planted-partition* model or the *stochastic block model* (SBM) is an example of such a random graph model for community detection that generalizes the well-known Erdős-Rényi graphs (Holland et al., 1983; Dyer and Frieze, 1989; Decelle et al., 2011; Abbe and Sandon, 2015a; Abbe et al., 2016; Hajek et al., 2015; Chin et al., 2015; Mossel et al., 2015). Consider a graph $G(V, E)$, where $V = V_1 \sqcup V_2 \sqcup \dots \sqcup V_k$ is a disjoint union of k clusters denoted by $V_1, \dots, V_k$. The edges of the graph are drawn randomly: there is an edge between $u \in V_i$ and $v \in V_j$ with probability $q_{i,j}$, $1 \leq i, j \leq k$.

This model has been popular both in theoretical and practical domains of community detection, and the aforementioned references are just a small sample. Recent theoretical works focus on characterizing sharp threshold of recovering the partition in the SBM. For example, when there are only two communities of exactly equal size, and the inter-cluster edge probability $q_{i,j} = \frac{b \log n}{n}$, $i \neq j$ and intra-cluster edge probability is $q_{i,i} = \frac{a \log n}{n}$, it is known that perfect recovery is possible if and only if $\sqrt{a} - \sqrt{b} > \sqrt{2}$ (Abbe et al., 2016; Mossel et al., 2015). The regime of the probabilities being $\Theta\left(\frac{\log n}{n}\right)$ has been put forward as one of most interesting ones, because in an Erdős-Rényi random graph, this is the threshold for graph connectivity (Bollobás, 1998). This result has been subsequently generalized for k communities (Abbe and Sandon, 2015a,b; Hajek et al., 2016) (for constant k or when $k = o(\log n)$), and under the assumption that the communities are generated according to a probabilistic generative model (there is a prior probability p_i of an element being in the i th community) (Abbe and Sandon, 2015a). Note that, the results are not only of theoretical

interest, many real-world networks exhibit a “sparsely connected” community feature (Leskovec et al., 2008), and any efficient recovery algorithm for SBM has many potential applications.

One aspect that the SBM does not account for is a “transitivity rule” (‘friends having common friends’) inherent to many social and other community structures. To be precise, consider any three vertices x, y and z . If x and y are connected by an edge, and y and z are connected by an edge, then it is more likely than not that x and z are connected by an edge. This phenomenon can be seen in many network structures - predominantly in social networks, blog-networks and advertising. SBM, primarily a generalization of Erdős-Rényi random graph, does not take into account this characteristic, and in particular, the event that an edge exists between x and z is independent of the events that there exist edges between x and y and y and z . However, one needs to be careful such that by allowing such “transitivity”, the simplicity and elegance of the SBM is not lost.

Further, from an algorithmic point of view, it is well known that triangle based heuristic algorithms work really well for graph clustering tasks (see (Tsourakakis et al., 2009), (Tsourakakis et al., 2017), (Kolountzakis et al., 2012)). In particular, in (Tsourakakis et al., 2017), the authors proposed a triangle counting based heuristic algorithm TECTONIC that has better performance on Amazon, DBLP and YouTube datasets for graph partitioning than many popular competitors including Spectral Clustering, Girvan-Newman algorithm, Louvain method and the Clauset-Newman-Moore (CNM) to name a few. Triangle (and motifs in general) based analytics has been extensively useful in different areas including social networks (Newman et al., 2002), metabolic networks (Milo et al., 2002), protein networks (Alon, 2007; Vinayagam et al., 2016), transportation networks (Rosvall et al., 2014), neural networks (Park and Friston, 2013), and food webs (Stouffer et al., 2012). Specifically, clustering methods that use triangles capture different interaction patterns. For example, optimizing to minimize the number of triangles across the cut returns similar energy flow patterns across species in the food web data (Benson et al., 2016). Further, it was shown in (Watts and Strogatz, 1998) that triangles are stronger signals of community structure than edges alone. This begs the following question: “Is it possible to provide a theoretical justification of why simple triangle based heuristics perform well on real-world data?”. It turns out that triangle counting can not recover the latent clusters in SBM in the logarithmic degree regime, which is one of the interesting regimes. This further limits the applicability of algorithms designed using the SBM assumption thus raising the possibility of a different random graph model that captures significantly more properties of real-world graphs.

Inspired by the above questions, we propose a novel random graph community detection model analogous to the stochastic block model, that we call the *geometric block model* (GBM). The GBM depends on the basic definition of the well-known class of random graphs called *random geometric graph* (RGG) (Gilbert, 1961) that has found a lot of practical use in wireless networking because of its inclusion of the notion of proximity between nodes (Penrose, 2003). The GBM satisfies several desiderata of real networks, such as the degree associativity property (high degree nodes tend to connect). However, analyzing the GBM is significantly challenging due to the presence of correlated edges and in particular, the techniques developed for the SBM do not extend to this setting.

More precisely, the GBM is defined as follows. Let $V \equiv V_1 \sqcup V_2 \sqcup \dots \sqcup V_k$ be the set of vertices that is a disjoint union of k clusters, denoted by $V_1, \dots, V_k$. Let, $\beta_{i,j}, i, j \in \{1, \dots, k\}$ be unknown latent variables. Given an integer $t \geq 1$, for each vertex $u \in V$, define a random vector $Z_u \in \mathbb{R}^{t+1}$ that is uniformly distributed in $\mathcal{S}^t \subset \mathbb{R}^{t+1}$, the t -dimensional sphere. In this random graph, an edge exists between $v \in V_i$ and $u \in V_j$ if and only if $\langle Z_u, Z_v \rangle \geq \beta_{i,j}$. In this special case of $t = 1$, the above definition is equivalent to choosing random variable θ_u uniformly distributed in

$[0, 2\pi]$, for all $u \in V$. Then there will be an edge between two vertices $u \in V_i, v \in V_j$ if and only if $\cos \theta_u \cos \theta_v + \sin \theta_u \sin \theta_v = \cos(\theta_u - \theta_v) \geq \beta_{i,j}$ or $\min\{|\theta_u - \theta_v|, 2\pi - |\theta_u - \theta_v|\} \leq \arccos \beta_{i,j}$. This in turn, is equivalent to choosing a random variable X_u uniformly distributed in $[0, 1]$ for all $u \in V$, and there exists an edge between two vertices $u \in V_i, v \in V_j$ if and only if

$$d_L(X_u, X_v) \equiv \min\{|X_u - X_v|, 1 - |X_u - X_v|\} \leq r_{i,j},$$

where $r_{i,j} \in [0, \frac{1}{2}]$, $0 \leq i, j \leq k$, are a set of real numbers. This corresponds to the Geometric block model in 1-dimension which will be referred to as GBM_1 .

For the rest of this paper, we concentrate on the case when $r_{i,i} = r_s$ for all $i \in \{1, \dots, k\}$, which we call the “intra-cluster distance” and $r_{i,j} = r_d$ for all $i, j \in \{1, \dots, k\}, i \neq j$, which we call the “inter-cluster distance,” to simplify the analysis. To allow for edge density to be higher inside the clusters than across the clusters, assume $r_s \geq r_d$.

1.1 Validation of GBM

We next give two examples of datasets that motivate the geometric block model. In particular, this datasets validate our hypothesis about geometric block model and the role of distance in the formation of edges. The first one is a product purchase metadata from Amazon, and the second one is a dataset with academic collaboration.

AMAZON METADATA

The first dataset that we use in our experiments is the Amazon product metadata on SNAP (<https://snap.stanford.edu/data/amazon-meta.html>), that has 548552 products and each product is one of the following types {Books, Music CD’s, DVD’s, Videos}. Moreover, each product has a list of attributes, for example, a book may have attributes like {“General”, “Sermon”, “Preaching”}. We consider the co-purchase network over these products. We make two observations here: (1) edges get formed (that is items are co-purchased) more frequently if they are similar, where we measure similarity by the number of common attributes between products, and (2) two products that share an edge have more common neighbors (number of items that are bought along with both those products) than two products with no edge in between.




Figure 1: Histogram: similarity of products bought together (mean ≈ 6)




Figure 2: Histogram: similarity of products not bought together (mean ≈ 2)

Figures 1 and 2 show respectively average similarity of products that were bought together, and not bought together. From the distribution, it is quite evident that edges in a co-purchase network gets formed according to distance, a salient feature of random geometric graphs, and the GBM.




Figure 3: Histogram of common neighbors of edges and non-edges in the co-purchase network, from left to right: Book-DVD, Book-Book, DVD-DVD

We next take equal number of product pairs inside Book (also inside DVD, and across Book and DVD) that have an edge in-between and do not have an edge respectively. Figure 3 shows that the number of common neighbors when two products share an edge is much higher than when they do not. In fact, almost all product pairs that do not have an edge in between also do not share any common neighbor. This can be explainable in sparse graphs if we imagine the latent features to follow the rules of GBM. On the other hand, this also suggests that SBM is not a good model for this network, as in sparse SBM, two nodes having common neighbors is extremely rare, irrespective of whether they share an edge or not. Note that, in practice, additional covariate features like product metadata, reviews etc. are also observed, which could be used for community detection.

ACADEMIC COLLABORATION

We consider the collaboration network of academicians in Computer Science since 2016 (data obtained from csrankings.org)¹. According to area of expertise of the authors, we consider five different communities: Data Management (MOD), Machine Learning and Data Mining (ML), Artificial Intelligence (AI), Robotics (ROB), Architecture (ARCH). If two authors share the same affiliation, or shared affiliation in the past, we assume that they are geographically close. We would like to hypothesize that, two authors in the same communities might collaborate even when they are geographically far. However, two authors in different communities are more likely to collaborate only if they share the same affiliation (or are geographically close). Table 1 describes the number of edges across the communities. It is evident that the authors from same community are likely to collaborate irrespective of the affiliations and the authors of different communities collaborate much frequently when they share affiliations or are close geographically. This clearly indicates that the inter cluster edges are likely to form if the distance between the nodes is quite small, motivating the fact $r_d < r_s$ in the GBM.

2. Description of Our Results

In Section 2.1, we describe a new random graph model that we call the *Random Annulus Graphs* (RAG) that we study in this paper. The RAG is a variant of the RGG and interestingly, the connectivity properties of the RAG will allow us to design algorithms with provable guarantees for community detection in the GBM. Subsequently, in section 2.2, we introduce the Geometric Block Model more

1. We use author affiliations from the paper available at <https://www.aminer.org/citation>

Area 1	Area 2	same	different	Area	same	different
ML	AI	65	43	MOD	88	91
AI	ROB	21	8	ARCH	51	45
ML	MOD	22	14	ROB	63	30
ARCH	MOD	7	1	AI	194	110
ROB	ARCH	5	1	ML	152	141

Table 1: On the left we count the number of inter-cluster edges when authors shared same affiliation and different affiliations. On the right, we count the same for intra-cluster edges.

formally. In section 2.3, we provide a brief discussion on the algorithms that we design for recovery of clusters in a random graph sampled according to GBM.

2.1 Random graph models

Models of random graphs are ubiquitous with Erdős-Rényi graphs (Erdős and Rényi, 1959; Gilbert, 1959) at the forefront. Studies of the properties of random graphs have led to many fundamental theoretical observations as well as many engineering applications. In an Erdős-Rényi graph $G(n, p)$, $n \in \mathbb{Z}_+$, $p \in [0, 1]$, the randomness lies in how the edges are chosen: each possible pair of vertices forms an edge independently with probability p . It is also possible to consider models of graphs where randomness lies in the vertices.

Keeping up with the simplicity of the Erdős-Rényi model, let us define a random annulus graph in one dimension (RAG₁) in the following way.

Definition 1 (Random annulus graph). *Let the t -dimensional unit sphere be denoted by $S^t \equiv \{x \in \mathbb{R}^{t+1} \mid \|x\|_2 = 1\}$ where $\|\cdot\|_2$ denote the ℓ_2 norm. For $t \geq 1$, a t -dimensional random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ on n vertices has parameters $n, t \in \mathbb{Z}_+$, and $r_1, r_2 \in [0, 2]$, $r_1 \leq r_2$. It is defined by assigning $X_i \in S^t$ to vertex i , $1 \leq i \leq n$, where X_i 's are independent and identical random vectors uniformly distributed in S^t . There will be an edge between vertices i and j , $i \neq j$, if and only if $r_1 \leq \|X_i - X_j\|_2 \leq r_2$.*

When from the context it is clear that we are in high dimensions ($t > 1$), we use $d(u, v)$ to denote $\|X_u - X_v\|_2$ or just the ℓ_2 distance between the arguments.

We give the name random annulus graph (RAG) because two vertices are connected if one is within an ‘annulus’ centered at the other. One can think of the random variables X_i , $1 \leq i \leq n$, to be uniformly distributed on the perimeter of a circle with radius $\frac{1}{2\pi}$ and the distance $d_L(\cdot, \cdot)$ to be the geodesic distance (the length of the smaller arc between the two points). For clarity in the calculations, it will be helpful to consider the vertices as just random points on $[0, 1]$. Note that every point has a natural left direction (if we think of them as points on a circle then this is the counterclockwise direction) and a right direction.

This definition is by no means new. For the case of $r_1 = 0, t = 1$, this is the random geometric graphs (RGG) in one dimension. Random Geometric graphs were defined first by (Gilbert, 1961) and constitute the first and simplest model of spatial networks. The definition of RAG₁ has been previously mentioned in (Dettmann and Georgiou, 2016). The interval $[r_1, r_2]$ is called the connectivity interval in RAG₁. Random geometric graphs have several desirable properties that model real human social networks, such as vertices with the degree associativity property (high degree nodes tend to

connect). This has led RGGs to be used as models of disease outbreak in social network (Eubank et al., 2004) and flow of opinions (Zhang et al., 2014). RGGs are an excellent model for wireless (ad-hoc) communication networks (Dettmann and Georgiou, 2016; Haenggi et al., 2009). From a more mathematical stand-point, RGGs act as a bridge between the theory of classical random graphs and that of percolation (Bollobás, 2001, 2006). Recent works on RGGs also include hypothesis testing between an Erdős-Rényi graph and a random geometric graph (Bubeck et al., 2016).

Threshold properties of Erdős-Rényi graphs have been at the center of much theoretical interest, and in particular it is known that many graph properties exhibit sharp phase transition phenomena (Friedgut and Kalai, 1996). Random geometric graphs also exhibit similar threshold properties (Penrose, 2003).

Random Annulus Graphs in one dimension We describe the case of $t = 1$ separately from $t > 1$ because in the former case we have exact connectivity result. The base of the logarithm is e here and everywhere else in the paper unless otherwise mentioned.

Theorem 2 (Connectivity threshold of one dimensional random annulus graphs). *The $\text{RAG}_1(n, [2 \sin \frac{\pi b \log n}{n}, 2 \sin \frac{\pi a \log n}{n}])$ is i) connected with probability $1 - o(1)$ if $a > 1$ and $a - b > 0.5$ ii) disconnected with probability $1 - o(1)$ if $a < 1$ or $a - b < 0.5$*

For the 1-dimensional case, the definition of random annulus graphs can be simplified, and that simplification explains the parameterization in the above theorem. The $\text{RAG}_1(n, [2 \sin \frac{\pi b \log n}{n}, 2 \sin \frac{\pi a \log n}{n}])$ is equivalent to $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ (i.e., the two random graphs follow the same distribution), where the later graph is defined below.

Definition 3. A random annulus graph $\text{RAG}_1^*(n, [r_1, r_2])$ on n vertices has parameters n , and a pair of real numbers $r_1, r_2 \in [0, 1/2]$, $r_1 \leq r_2$. It is defined by assigning a number $X_i \in \mathbb{R}$ to vertex i , $1 \leq i \leq n$, where X_i s are independent and identical random variables uniformly distributed in $[0, 1]$. There will be an edge between vertices i and j , $i \neq j$, if and only if $r_1 \leq d_L(X_i, X_j) \leq r_2$ where $d_L(X_i, X_j) \equiv \min\{|X_i - X_j|, 1 - |X_i - X_j|\}$.

For the 1-dimensional case, for any two vertices u, v , let $d(u, v)$ denote $d_L(X_u, X_v)$ where X_u, X_v are corresponding random values to the vertices respectively. We can extend this notion to denote the distance $d(u, v)$ between a vertex u (or the embedding of that vertex in $[0, 1]$) and a number $v \in [0, 1]$ naturally.

Remark 4. For the analysis and proofs of the 1-dimensional case, we will use the definition of RAG_1^* (instead of RAG_1). However, due to the equivalence between the two random graph models, our results for RAG_1^* also hold for RAG_1 with appropriately modified parameters. This is because we can think of the vertices in $\text{RAG}_1^*(n, [r_1, r_2])$ being distributed randomly on the circumference of circle of radius $1/2\pi$ and distance between two vertices in $\text{RAG}_1^*(n, [r_1, r_2])$ is geodesic (measured along the circumference of the circle). It is easier to reason about the geodesic distance over Euclidean distance in 1-dimension; however, the advantage is no longer there in higher dimensions. Since handling geodesic distances is more cumbersome in the higher dimensions, we resorted to Euclidean distance for $t > 1$ in order to retain simplicity.

Consider a $\text{RAG}_1^*(n, [0, r])$ defined above with $r = \frac{a \log n}{n}$. It is known that $\text{RAG}_1^*(n, [0, r])$ is connected with high probability if and only if $a > 1^2$ ((Muthukrishnan and Pandurangan, 2005; Pen-

2. That is, $\text{RAG}_1^*(n, [0, \frac{(1+\epsilon) \log n}{n}])$ is connected for any $\epsilon > 0$. We will ignore this ϵ and just mention connectivity threshold as $\frac{\log n}{n}$.

rose, 2003) See also (Penrose, 2016)). Now let us consider the graph $\text{RAG}_1^*(n, [\frac{\delta \log n}{n}, \frac{\log n}{n}])$, $\delta > 0$. Clearly this graph has less edges than $\text{RAG}_1^*(n, [0, \frac{\log n}{n}])$. **Is this graph still connected?** Surprisingly, we show that the above modified graph remains connected as long as $\delta \leq 0.5$. Note that, on the other hand, $\text{RAG}_1^*(n, [0, \frac{(1-\epsilon) \log n}{n}])$ is not connected for any $\epsilon > 0$.

This means the graphs $\text{RAG}_1^*(n, [0, \frac{0.99 \log n}{n}])$ and $\text{RAG}_1^*(n, [\frac{0.49 \log n}{n}, \frac{0.99 \log n}{n}])$ are not connected with high probability, whereas $\text{RAG}_1^*(n, [\frac{0.50 \log n}{n}, \frac{\log n}{n}])$ is connected. For a depiction of the connectivity regime for the random annulus graph $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ see Figure 4.




Figure 4: The shaded area in the a - b plot shows the regime where an $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is connected with high probability.

Can we explain this seemingly curious shift in connectivity interval, when one goes from $b = 0$ to $b > 0$? Compare the $\text{RAG}_1^*(n, [\frac{0.50 \log n}{n}, \frac{\log n}{n}])$ with the $\text{RAG}_1^*(n, [0, \frac{\log n}{n}])$. The former one can be thought of being obtained by deleting all the ‘short-distance’ edges from the latter. It turns out the ‘long-distance’ edges are sufficient to maintain connectivity, because they can connect points over multiple hops in the graph. Another possible explanation is that connectivity threshold for RAG_1^* is not dictated by isolated nodes as is the case in Erdős-Rényi graphs. Thus, after the connectivity threshold has been achieved, removing certain short edges still retains connectivity.

Random Annulus Graphs in Higher Dimension It is natural to ask similar question of connectivity for random annulus graphs in higher dimension. In a random annulus graph at dimension t , we may assign t -dimensional random vectors to each of the vertices, and use a standard metric such as the Euclidean distance to decide whether there should be an edge between two vertices.

The $\text{RAG}_t(n, [0, r])$ gives the standard definition of random geometric graphs in t dimensions (for example, see (Bubeck et al., 2016) or (Penrose, 2003)). Our main result here gives a condition that guarantees connectivity.

Theorem 5. *Let $t > 1$ and*

$$\psi(t) \equiv \frac{\sqrt{\pi}(t+1)\Gamma(\frac{t+2}{2})}{\Gamma(\frac{t+3}{2})},$$

where $\Gamma(x) = \int_0^\infty y^{x-1} e^{-y} dy$ is the gamma function. If $(a/2)^t - b^t \geq 8(t+1)\psi(t)$ and $a > 2b$, then a t -dimensional random annulus graph $\text{RAG}_t(n, [b(\frac{\log n}{n})^{1/t}, a(\frac{\log n}{n})^{1/t}])$ is connected with probability $1 - o(1)$.

Computing the connectivity threshold of RAG in high dimensions exactly is highly challenging, and we have to use several approximations of high dimensional geometry. Our arguments crucially rely on VC dimensions of sets of geometric objects such as intersections of high dimensional annuluses and hyperplanes.

In the process, we also derived the following results about the existence of isolated vertices in random annulus graphs.

Theorem 6 (Zero-One law for Isolated Vertices in RAG_t). *Let $t > 1$. For a t -dimensional random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ where $r_2 = a\left(\frac{\log n}{n}\right)^{\frac{1}{t}}$ and $r_1 = b\left(\frac{\log n}{n}\right)^{\frac{1}{t}}$, i) there exists isolated nodes with probability $1 - o(1)$ if $a^t - b^t < \psi(t)$, and ii) the graph is connected with probability $1 - o(1)$ if $a^t - b^t > \psi(t)$.*

An obvious deduction from this theorem is that an $\text{RAG}_t(n, [b\left(\frac{\log n}{n}\right)^{\frac{1}{t}}, a\left(\frac{\log n}{n}\right)^{\frac{1}{t}}])$ is not connected with probability $1 - o(1)$ if $a^t - b^t < \psi(t)$.

All these connectivity results find immediate application in analyzing the algorithm that we propose for the geometric block model (GBM). A GBM is a generative model for networks (graphs) with underlying community structure.

2.2 Geometric Block Model

Since for random geometric graphs in 1-dimensional case the distance $d_L(\cdot, \cdot)$ defined in Def. 3 gives a simpler expression, we use this for the 1-dimensional case. For higher dimensions we resort back to the Euclidean distance.

Definition 7 (Geometric Block Model in 1-dimension). *Given $V = V_1 \sqcup V_2, |V_1| = |V_2| = \frac{n}{2}$, choose a random variable X_u uniformly distributed in $[0, 1]$ for all $u \in V$. The geometric block model $\text{GBM}_1(r_s, r_d)$ with parameters $r_s > r_d$ is a random graph where an edge exists between vertices u and v if and only if,*

$$\begin{aligned} d_L(X_u, X_v) &\leq r_s \text{ when } u, v \in V_1 \text{ or } u, v \in V_2 \\ d_L(X_u, X_v) &\leq r_d \text{ when } u \in V_1, v \in V_2 \text{ or } u \in V_2, v \in V_1. \end{aligned}$$

As a consequence of the connectivity lower bound on RAG_1^* , we are able to show that recovery of the partition in $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$ is not possible with high probability by any means whenever $a - b < 0.5$ or $a < 1$ (see, Theorem 2). Another consequence of the random annulus graph results is that we show that if in addition to a GBM_1 graph, all the locations of the vertices are also provided, then recovery is possible if and only if $a - b > 0.5$ or $a > 1$ (formal statement in Theorem 19).

Coming back to the actual recovery problem, our main contribution for GBM_1 is to provide a simple and efficient algorithm that performs well in the sparse regime (see, Algorithm 1).

Theorem 8 (Recovery algorithm for GBM_1). *Suppose we have the graph $G(V, E)$ generated according to $\text{GBM}_1(r_s \equiv \frac{a \log n}{n}, r_d \equiv \frac{b \log n}{n})$, $a \geq 2b$. Define*

$$\begin{aligned} f_1 &= \min\{f : (2b + f) \log \frac{2b + f}{2b} - f > 1\}, \quad f_2 = \min\{f : (2b - f) \log \frac{2b - f}{2b} + f > 1\} \\ \theta_1 &= \max\{\theta : \frac{1}{2} \left((4b + 2f_1) \log \frac{4b + 2f_1}{2a - \theta} + 2a - \theta - 4b - 2f_1 \right) > 1 \text{ and } 0 \leq \theta \leq 2a - 4b - 2f_1\} \end{aligned}$$




Figure 5: Comparison of a and b in $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$. Red region denotes the values where the graph is disconnected, green region is recoverable by Algorithm 1 and yellow region denotes values of a and b for which there is no algorithm to recover the clusters.

$$\theta_2 = \min\{\theta : \frac{1}{2} \left((4b - 2f_2 \log \frac{4b - 2f_2}{2a - \theta} + 2a - \theta - 4b + 2f_2) > 1 \text{ and } a \geq \theta \geq \max\{2b, 2a - 4b + 2f_2\} \right\}.$$

Then, if $a - \theta_2 + \theta_1 > 2$ or $a > \max(1 + \theta_2, 2)$, there exists an efficient algorithm which will recover the correct partition in G with probability $1 - o(1)$. Moreover, if $a - b < 0.5$ or $a < 1$, any algorithm to recover the partition in $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$ will give incorrect output with probability $1 - o(1)$.

Some examples of the parameters when the proposed algorithm (Algorithm 1) can successfully recover are given in Table 2 and Figure 5 compares them with the lower bound.

b	0.01	1	2	3	4	5	6	7
Minimum value of a	3.18	8.96	12.63	15.9	18.98	21.93	24.78	27.57

Table 2: Minimum value of a , given b for which Algorithm 1 resolves clusters correctly in $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$.

As can be anticipated, the connectivity results for RAG applies to the ‘high dimensional’ geometric block model. In many applications, the latent feature space of nodes are high-dimensional. For example, road networks are two-dimensional whereas the number of features used in a social network may have much higher dimensions.

Definition 9 (The GBM in High Dimensions). *Let $t > 1$. Given $V = V_1 \sqcup V_2, |V_1| = |V_2| = \frac{n}{2}$, choose a random vector X_u independently uniformly distributed in $S^t \subset \mathbb{R}^{t+1}$ for all $u \in V$. The geometric block model $\text{GBM}_t(r_s, r_d)$ with parameters $r_s > r_d$ is a random graph where an edge exists between vertices u and v if and only if,*

$$\begin{aligned} \|X_u - X_v\|_2 &\leq r_s \text{ when } u, v \in V_1 \text{ or } u, v \in V_2 \\ \|X_u - X_v\|_2 &\leq r_d \text{ when } u \in V_1, v \in V_2 \text{ or } u \in V_2, v \in V_1. \end{aligned}$$

We extend the algorithmic results to high dimensions (details in Section 6).

Theorem 10. *Let $t > 1$. If $r_s = \Theta((\frac{\log n}{n})^{\frac{1}{t}})$ and $r_s - r_d = \Theta((\frac{\log n}{n})^{\frac{1}{t}})$, there exists a polynomial time efficient algorithm that recovers the partition from $\text{GBM}_t(r_s, r_d)$ with probability $1 - o(1)$.*

Moreover, if $r_s - r_d = o((\frac{\log n}{n})^{\frac{1}{t}})$ or $r_s = o((\frac{\log n}{n})^{\frac{1}{t}})$, any algorithm fails to recover the partition with probability at least $1/2$.

In the following section, we present the high level ideas of our algorithms to recover the clusters with sub-quadratic running time.

2.3 Clustering Algorithm

As we have mentioned already, it is well known that triangle based heuristics have good performance on real world graphs for cluster recovery. It turns out that for random graphs generated according to the GBM, simple triangle based algorithms are order optimal. A very simple and intuitive algorithm to cluster the vertices of a graph is as follows:

1. For a particular edge say (u, v) , count the number of triangles that contain the edge (u, v) and from this observed statistics, infer based on a threshold if the endpoints u and v belong to the same cluster or different cluster.
2. We perform this inference step on all edges of the graph and using the results, we can partition the nodes into their respective clusters. The analysis of the first step guarantees that our algorithm labels an edge as intra-cluster or inter-cluster with a probability of $1 - O(\frac{1}{n \log^2 n})$. On applying union bound, it guarantees that all edges are labelled correctly with probability $1 - O(\frac{1}{\log n})$.

This algorithm was analyzed in a preliminary conference version of this paper (Galhotra et al., 2018) and although the algorithm was order optimal, the guarantees for cluster recovery in the sparse logarithmic regime i.e. $r_s, r_d = \Theta(\frac{\log n}{n})$ hold only when $r_s \geq 4r_d$. When $r_s \leq 4r_d$, it is not possible to infer if the nodes forming an edge belong to the same cluster or not, based on a single threshold on the number of triangles the edge is part of. Therefore, we proceed to devise an improved algorithm whose main ingredient is still based on triangle counting but uses two thresholds on the triangle count instead of one. At a high level, if the number of triangles formed by an edge is between these two thresholds, we delete the edge because it is not possible to infer correctly whether the endpoints belong to the same or different cluster. We then show that all the remaining edges belong to same cluster and identify the condition when the surviving edges suffice to recover the clusters correctly. The connectivity of surviving edges is shown with the help of corresponding results for the Random Annulus graphs. A preliminary version of these results were published in (Galhotra et al., 2019).

Remark 11. *When the degree of a graph is logarithmic in the number of vertices (i.e., about $n \log n$ edges), a stochastic block model implies the existence of about $\log^3 n$ triangles, whereas in GBM there will be about $n \log^2 n$ triangles. As a result, in that regime, most vertices in SBM are not part of any triangles, and the above triangle counting algorithm will fail. However, in the same regime, the algorithm provably recover the communities in GBM, and also in practice.*

It is possible to generalize the GBM to include different distributions and different metric spaces. It is also possible to construct other type of spatial block models such as the one put forward in a parallel work (Sankararaman and Baccelli, 2018) which rely on the random dot product graphs (Young and Scheinerman, 2007). In (Sankararaman and Baccelli, 2018), edges are drawn between vertices randomly and independently as a function of the distance between the corresponding vertex

random variables. (Sankararaman and Baccelli, 2018) also considers the recovery scenario where in addition to the graph, values of the vertex random variables are provided. In GBM, we only observe the graph.

The rest of the paper is organized as follows. First, in Section 3, the sharp connectivity phase transition results for random annulus graphs are proven (details in Appendix A). In Section 4, the connectivity results are proven for high dimensional random annulus graphs (details in Appendix B). In Section 5, a lower bound for the geometric block model as well as the main recovery algorithm are presented. In Section 6, we present the extension of our recovery algorithm to high dimensional GBM. Finally, Section 7 empirically evaluates our techniques on various real world and synthetic datasets.

3. Connectivity of random annulus Graphs

In this section we give a sketch of the proof of sufficient condition for connectivity of RAG_1^* (as part of proving Theorem 2) since connectivity guarantees in RAG_1^* can be mapped to connectivity properties in RAG_1 (recall that the $\text{RAG}_1(n, [2 \sin \frac{\pi b \log n}{n}, 2 \sin \frac{\pi a \log n}{n}])$ is equivalent to $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$). The full details along with the proof of the necessary condition in Theorem 2 are given in Appendix A.

3.1 Sufficient condition for connectivity of RAG_1^*

Theorem. (Sufficient Condition in Theorem 2 for RAG_1^*) *If $a > 1$ and $a - b > 0.5$, the random annulus graph $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is connected with probability $1 - o(1)$.*

To prove this theorem we use two main technical lemmas that show two different events happen with high probability simultaneously.




Figure 6: Each vertex having two neighbors on either direction implies the graph is a union of cycles. The cycles can be interleaving in $[0, 1]$.

Lemma 12. *A set of vertices $\mathcal{C} \subseteq V$ is called a cover of $[0, 1]$, if for any point y in $[0, 1]$ there exists a vertex $v \in \mathcal{C}$ such that $d(v, y) \leq \frac{a \log n}{2n}$. If $a - b > 0.5$ and $a > 1$, the random annulus graph $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is a union of cycles such that every cycle forms a cover of $[0, 1]$ (see Figure 6) with probability $1 - o(1)$.*

This lemma also shows effectively the fact that ‘long-edges’ are able to connect vertices over multiple hops. Note that, the statement of Lemma 12 would be easier to prove if the condition were $a - b > 1$. In that case what we prove is that every vertex has neighbors (in the RAG_1^*) on both of the left and right directions. To see this for each vertex u , assign two indicator $\{0, 1\}$ -random variables A_u^l and A_u^r , with $A_u^l = 1$ if and only if there is no node x to the left of node u such that $d(u, x) \in [\frac{b \log n}{n}, \frac{a \log n}{n}]$. Similarly, let $A_u^r = 1$ if and only if there is no node x to the right of node u such that $d(u, x) \in [\frac{b \log n}{n}, \frac{a \log n}{n}]$. Now define $A = \sum_u (A_u^l + A_u^r)$. We have,

$$\Pr(A_u^l = 1) = \Pr(A_u^r = 1) = (1 - \frac{(a - b) \log n}{n})^{n-1},$$

and,

$$\mathbb{E}[A] = 2n(1 - \frac{(a - b) \log n}{n})^{n-1} \leq 2n^{1-(a-b)}.$$

If $a - b > 1$ then $\mathbb{E}[A] = o(1)$ which implies, by invoking Markov inequality, that with high probability every node will have neighbors (connected by an edge in the RAG_1^*) on either side. This results in the interesting conclusion that every vertex will lie in a cycle that covers $[0, 1]$. This is true for every vertex, hence the graph is simply a union of cycles each of which is a cover of $[0, 1]$. The main technical challenge is to show that this conclusion remains valid even when $a - b > 0.5$, which is proved in Lemma 12 in Appendix A.

Lemma 13. *Set two real numbers $k \equiv \lceil b/(a - b) \rceil + 1$ and $\epsilon < \frac{1}{2k}$. In an $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$, $0 < b < a$, with probability $1 - o(1)$ there exists a vertex u_0 and k nodes $\{u_1, u_2, \dots, u_k\}$ to the right of u_0 such that $d(u_0, u_i) \in [\frac{(i(a-b)-2i\epsilon) \log n}{n}, \frac{(i(a-b)-(2i-1)\epsilon) \log n}{n}]$ and k nodes $\{v_1, v_2, \dots, v_k\}$ to the right of u_0 such that $d(u_0, v_i) \in [\frac{((i(a-b)+b-(2i-1)\epsilon) \log n)}{n}, \frac{(i(a-b)+b-(2i-2)\epsilon) \log n}{n}]$, for $i = 1, 2, \dots, k$. The arrangement of the vertices is shown in Figure 13 (pg. 18).*

With the help of these two lemmas, we are in a position to prove the sufficient condition in Theorem 2. The proofs of the two lemmas are given in Appendix A and contain the technical essence of this section.

Proof [Proof of Sufficient condition in Theorem 2] We have shown that the two events mentioned in Lemmas 12 and 13 happen with high probability. Therefore they simultaneously happen under the condition $a > 1$ and $a - b > 0.5$. Now we will show that these events together imply that the graph is connected. To see this, consider the vertices $u_0, \{u_1, u_2, \dots, u_k\}$ and $\{v_1, v_2, \dots, v_k\}$ that satisfy the conditions of Lemma 13. We can observe that each vertex v_i has an edge with u_i and u_{i-1} , $i = 1, \dots, k$. This is because (see Figure 13 for a depiction)

$$d(u_i, v_i) \geq \frac{((i(a - b) + b - (2i - 1)\epsilon) \log n)}{n} - \frac{i(a - b) - (2i - 1)\epsilon \log n}{n} = \frac{b \log n}{n} \quad \text{and}$$

$$d(u_i, v_i) \leq \frac{i(a - b) + b - (2i - 2)\epsilon \log n}{n} - \frac{(i(a - b) - 2i\epsilon) \log n}{n} = \frac{(b + 2\epsilon) \log n}{n}.$$

Similarly,

$$d(u_{i-1}, v_i) \geq \frac{((i(a - b) + b - (2i - 1)\epsilon) \log n)}{n} - \frac{(i - 1)(a - b) - (2i - 3)\epsilon \log n}{n}$$

$$= \frac{(a - 2\epsilon) \log n}{n} \quad \text{and}$$

$$d(u_{i-1}, v_i) \leq \frac{i(a - b) + b - (2i - 2)\epsilon \log n}{n} - \frac{((i - 1)(a - b) - 2(i - 1)\epsilon) \log n}{n} = \frac{a \log n}{n}.$$

This implies that u_0 is connected to u_i and v_i for all $i = 1, \dots, k$. Using Lemma 12, the first event implies that the connected components are cycles spanning the entire line $[0, 1]$. Now consider two such disconnected components, one of which consists of the nodes $u_0, \{u_1, u_2, \dots, u_k\}$ and $\{v_1, v_2, \dots, v_k\}$. There must exist a node t in the other component (cycle) such that t is on the right of u_0 and $d(u_0, t) \equiv \frac{x \log n}{n} \leq \frac{a \log n}{n}$. If $x \leq b$, $\exists i \mid i \leq k$ and $i(a - b) + b - a - (2i - 2)\epsilon \leq x \leq i(a - b) - (2i - 1)\epsilon$ (see Figure 14). When $x \leq b$, we can calculate the distance between t and v_i as

$$d(t, v_i) \geq \frac{(i(a - b) + b - (2i - 1)\epsilon) \log n}{n} - \frac{(i(a - b) - (2i - 1)\epsilon) \log n}{n} = \frac{b \log n}{n}$$

and

$$d(t, v_i) \leq \frac{(i(a - b) + b - (2i - 2)\epsilon) \log n}{n} - \frac{(i(a - b) + b - a - (2i - 2)\epsilon) \log n}{n} = \frac{a \log n}{n}.$$

Therefore t is connected to v_i when $x \leq b$. If $x > b$ then t is already connected to u_0 . Therefore the two components (cycles) in question are connected. This is true for all cycles and hence there is only a single component in the entire graph. Indeed, if we consider the cycles to be disjoint super-nodes, then we have shown that there must be a star configuration. $\blacksquare$

The following result is an immediate corollary of the connectivity upper bound.

Corollary 14. *Consider a random graph $G(V, E)$ is being generated as a variant of the RAG_1^* where each $u, v \in V$ forms an edge if and only if $d(u, v) \in \left[0, c \frac{\log n}{n}\right] \cup \left[b \frac{\log n}{n}, a \frac{\log n}{n}\right]$, $0 < c < b < a$. If $a - b + c > 1$ or if $a - b > 0.5, a > 1$, the graph G is connected with probability $1 - o(1)$.*

The above corollary can be further improved for some regimes of a, b, c . In particular, we can get the following result (proof delegated to the appendix).

Corollary 15. *Consider a random graph $G(V, E)$ is being generated as a variant of the RAG_1^* where each $u, v \in V$ forms an edge if and only if $d(u, v) \in \left[0, c \frac{\log n}{n}\right] \cup \left[b \frac{\log n}{n}, a \frac{\log n}{n}\right]$, $0 < c < b < a$. If any of the following conditions are true:*

1. $2(a - b) + c/2 > 1$ when $a - b < c$ and $b > 3c/2$
2. $b - c > 1$ when $a - b < c$ and $b \leq 3c/2$
3. $a > 1$ when $a - b \geq c$ and $b \leq 3c/2$
4. $(a - b) + 3c/2 > 1$ when $a - b \geq c$ and $b > 3c/2$,

the graph G is connected with probability $1 - o(1)$.

4. Connectivity of High Dimensional Random Annulus Graphs: Proof of Theorem 5

In this section we show a proof sketch of Theorem 5 to establish the sufficient condition of connectivity of random annulus graphs. The details of the proof and the necessary conditions are provided in Appendix B.

Note, here $r_1 \equiv b \left(\frac{\log n}{n}\right)^{1/t}$ and $r_2 \equiv a \left(\frac{\log n}{n}\right)^{1/t}$. We show the upper bound for connectivity of a Random Annulus Graphs in t dimension as shown in Theorem 5. For this we first define a *pole* as a vertex which is connected to all vertices within a distance of r_2 from itself. In order to prove Theorem 5, we first show the existence of a pole with high probability in Lemma 16.

Lemma 16. *In a $\text{RAG}_t \left(n, \left[b \left(\frac{\log n}{n}\right)^{1/t}, a \left(\frac{\log n}{n}\right)^{1/t} \right] \right)$, $0 < b < a$, with probability $1 - o(1)$ there exists a pole.*

Next, Lemma 17 shows that for every vertex u and every hyperplane L passing through u and not too close to the tangent hyperplane at u , there will be a neighbor of u on either side of the plane. Therefore, there should be a neighbor towards the direction of the pole. In order to formalize this, let us define a few regions associated with a node u and a hyperplane $L : w^T x = \beta$ passing through u .

$$\begin{aligned}\mathcal{R}_L^1 &\equiv \{x \in S^t \mid r_1 \leq d(u, x) \leq r_2, w^T x \leq \beta\} \\ \mathcal{R}_L^2 &\equiv \{x \in S^t \mid r_1 \leq d(u, x) \leq r_2, w^T x \geq \beta\} \\ \mathcal{A}_L &\equiv \{x \mid x \in S^t, w^T x = \beta\}.\end{aligned}$$

Informally, $\mathcal{R}_L^1$ and $\mathcal{R}_L^2$ represent the partition of the annulus on either side of the hyperplane L and $\mathcal{A}_L$ represents the region on the sphere lying on L .

Lemma 17. *If we sample n nodes from S^t according to $\text{RAG}_t \left(n, \left[b \left(\frac{\log n}{n}\right)^{1/t}, a \left(\frac{\log n}{n}\right)^{1/t} \right] \right)$, then for every node u and every hyperplane L passing through u such that $\mathcal{A}_L$ is not all within distance r_2 of u , node u has a neighbor on both sides of the hyperplane L with probability at least $1 - \frac{1}{n}$ provided $(a/2)^t - b^t \geq \frac{8\sqrt{\pi}(t+1)^2\Gamma(\frac{t+2}{2})}{\Gamma(\frac{t+3}{2})}$ and $a > 2b$.*

We found the proof of this lemma to be challenging. Since, we do not know the location of the pole, we need to show that every point has a neighbor on both sides of the plane L no matter what the orientation of the plane. Since the number of possible orientations is uncountably infinite, we cannot use a union-bound type argument. To show this we have to rely on the VC Dimension of the family of sets $\{x \in S^t \mid r_1 \leq \|u - x\|_2 \leq r_2, w^T x \geq \beta, \mathcal{A}_{L:w^T x = \beta} \text{ not all within } r_2 \text{ of } u\}$ for all hyperplanes L (which can be shown to be less than $t + 1$). We rely on the celebrated result of (Haussler and Welzl, 1987) (we derived a continuous version of it), see Theorem 41, to deduce our conclusion.

For a node u and its corresponding location $X_u = (u_1, u_2, \dots, u_{t+1})$, define the particular hyperplane $L_u^* : x_1 = u_1$ which is normal to the line joining $u_0 \equiv (1, 0, \dots, 0)$ and the origin and passes through u . We now need one more lemma that will help us prove Theorem 5.

Lemma 18. *For a particular node u and corresponding hyperplane L_u^* , if every point in $\mathcal{A}_{L_u^*}$ is within distance r_2 from u , then u must be within r_2 of u_0 .*

For now, we assume that the Lemmas 16, 17 and 18 are true and show why these lemmas together imply the proof of Theorem 5.

Proof [Proof of Theorem 5] We consider an alternate (rotated but not shifted) coordinate system by multiplying every vector by an orthonormal matrix such that the new position of the pole is the $t + 1$ -dimensional vector $(1, 0, \dots, 0)$ where only the first co-ordinate is non-zero. Let the $t + 1$ dimensional vector describing any node u in this new coordinate system be $\hat{u} = (\hat{u}_1, \hat{u}_2, \dots, \hat{u}_{t+1})$. Now consider the hyperplane $L : x_1 = \hat{u}_1$ and if u is not connected to the pole already, then by Lemma 17 and Lemma 18, the node u has a neighbor u_2 which has a higher first coordinate ($\hat{u}_2 > \hat{u}_1$). The same analysis applies for u_2 and hence we have a path where the first coordinate of every node is higher than the previous node. Since the number of nodes is finite, this path cannot go on indefinitely and at some point, one of the nodes is going to be within r_2 of the pole and will be connected to the pole. Therefore every node is going to be connected to the pole and hence our theorem is proved. $\blacksquare$

5. The Geometric Block Model in one dimension (GBM₁)

In this section, we prove the necessary condition for exact cluster recovery of GBM₁ and give an efficient algorithm that matches that within a constant factor.

5.1 Immediate consequence of RAG₁^{*} connectivity

The following lower bound for GBM₁ can be obtained as a consequence of Theorem 2.

Theorem (Impossibility in GBM₁ in Theorem 8). *If $a - b < 0.5$ or $a < 1$, then any algorithm to recover the partition in GBM₁($\frac{a \log n}{n}, \frac{b \log n}{n}$) will give incorrect output with probability $1 - o(1)$.*

Proof Consider the scenario that not only the geometric block model graph GBM₁($\frac{a \log n}{n}, \frac{b \log n}{n}$) was provided to us, but also the random values $X_u \in [0, 1]$ for all vertex u in the graph were provided. We will show that we will still not be able to recover the correct partition of the vertex set V with probability at least 0.5 (with respect to choices of X_u , $u, v \in V$ and any randomness in the algorithm).

In this situation, the edge (u, v) where $d_L(X_u, X_v) \leq \frac{b \log n}{n}$ does not give any new information than X_u, X_v . However the edges (u, v) where $\frac{b \log n}{n} \leq d_L(X_u, X_v) \leq \frac{a \log n}{n}$ are informative, as existence of such an edge will imply that u and v are in the same part. These edges constitute a one-dimensional random annulus graph RAG₁^{*}($n, [\frac{b \log n}{n}, \frac{a \log n}{n}]$). But if there are more than two components in this random annulus graph, then it is impossible to separate out the vertices into the correct two parts, as the connected components can be assigned to any of the two parts and the RAG₁^{*} along with the location values $(X_u, u \in V)$ will still be consistent.

What remains to be seen that RAG₁^{*}($n, [\frac{b \log n}{n}, \frac{a \log n}{n}]$) will have $\omega(1)$ components with high probability if $a - b < 0.5$ or $a < 1$. This is certainly true when $a - b < 0.5$ as we have seen in Theorem 8, there can indeed be $\omega(1)$ isolated nodes with high probability. On the other hand, when $a < 1$, just by using an analogous argument it is possible to show that there are $\omega(1)$ vertices that do not have any neighbors on the left direction (counterclockwise). We delegate the proof of this claim as Lemma 43 in the appendix. If there are k such vertices, there must be at least $k - 1$ disjoint candidates. This completes the proof. $\blacksquare$

Indeed, when the locations X_u associated with every vertex u is provided, it is also possible to recover the partition exactly when $a - b > 0.5$ and $a > 1$, matching the above lower bound exactly (see Theorem 19).

Corollary 19 (GBM₁ with known vertex locations). *Suppose a geometric block model graph $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$ is provided along with the associated values of the locations X_u for every vertex u . Any algorithm to recover the partition in $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$ will give i) incorrect output with probability $1 - o(1)$ if $a - b < 0.5$ or $a < 1$ and ii) correct output probability $1 - o(1)$ if $a - b > 0.5$ and $a > 1$.*

Proof We need to only prove that it is possible to recover the partition exactly with probability $1 - o(1)$ when $a - b > 0.5$ and $a > 1$, since the other part is immediate from the impossibility guarantee in GBM₁ in Theorem 8 (When $a - b < 0.5$ or $a < 1$, there will be an isolated node with high probability. It is impossible to identify the cluster for this node.). For any pair of vertices u, v , we can verify if $d(u, v) \in [\frac{b \log n}{n}, \frac{a \log n}{n}]$. If that is the case then by just checking in the GBM₁ graph whether they are connected by an edge or not we can decide whether they belong to the same cluster or not respectively. What remains to be shown that all vertices can be covered by this procedure. However that will certainly be the case since $\text{RAC}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is connected with high probability. ■

5.2 A recovery algorithm for GBM₁

We now turn our attention to an efficient recovery algorithm for GBM₁. Intriguingly, we show a simple triangle counting algorithm works well for GBM₁ and recovers the clusters in the sparsity regime. Triangle counting algorithms are popular heuristics applied to social networks for clustering (Easley et al., 2012), however they fail in SBM. Hence, this serves as another validation why GBM₁ are well-suited to model community structures in social networks.

The algorithm is as follows. Suppose we are given a graph $G = (V : |V| = n, E)$ with two disjoint parts, $V_1, V_2 \subseteq V$ generated according to $\text{GBM}_1(r_s, r_d)$. The algorithm (Algorithm 1) goes over all edges $(u, v) \in E$. It counts the number of triangles containing the edge (u, v) by calling the `process` function that counts the number of common neighbors of u and v .

`process` outputs ‘true’ if it is confident that the nodes u and v belong to the same cluster and ‘false’ otherwise. More precisely, if the count is within some prescribed values E_S and E_D , it returns ‘false’³. The algorithm removes the edge on getting a ‘false’ from `process` function. After processing all the edges of the network, the algorithm is left with a reduced graphs (with certain edges deleted from the original). It then finds the connected components in the graph and returns them as the parts V_1 and V_2 .

Remark 20. *The algorithm can iteratively maintain the connected components over the processed edges (the pairs for which process function has been called and it returned true) like the union-find algorithm. This reduces the number of queries as the algorithm does not need to call the process function for the edges which are present in the same connected component.*

3. Note that, the thresholds E_S and E_D refer to the maximum and minimum value of triangle-count for an ‘inter cluster’ edge.

Algorithm 1: Cluster recovery in GBM_1

Require: $\text{GBM}_1 G = (V, E), r_s, r_d$

- 1: **for** $(u, v) \in E$ **do**
- 2: **if** $\text{process}(u, v, r_s, r_d)$ **then**
- 3: continue
- 4: **else**
- 5: $E.\text{remove}((u, v))$
- 6: **end if**
- 7: **end for**
- 8: **return** $\text{connectedComponent}(V, E)$

Algorithm 2: process

Require: u, v, r_s, r_d

Ensure: true/false

{Comment: When $a > 2b$, $f_1 = \min\{f : (2b + f) \log \frac{2b+f}{2b} - f > 1\}$, $f_2 = \min\{f : (2b - f) \log \frac{2b-f}{2b} + f > 1\}$ and $E_S = (2b + f_1) \frac{\log n}{n}$ and $E_D = (2b - f_2) \frac{\log n}{n}$ }

- 1: $\text{count} \leftarrow |\{z : (z, u) \in E, (z, v) \in E\}|$
- 2: **if** $\frac{\text{count}}{n} \geq E_S(r_d, r_s)$ or $\frac{\text{count}}{n} \leq E_D(r_d, r_s)$ **then**
- 3: **return** true
- 4: **end if**
- 5: **return** false

5.3 Analysis of Algorithm 1

Given a GBM_1 graph $G(V, E)$ with two clusters $V = V_1 \sqcup V_2$, and a pair of vertices $u, v \in V$, the events $\mathcal{E}_z^{u,v}, z \in V$ of any other vertex z being a common neighbor of both u and v given $(u, v) \in E$ are dependent; however given the distance between the corresponding random variables $d_L(X_u, X_v) = x$, the events are independent. This is a crucial observation which helps us to overcome the difficulty of handling correlated edge formation.

Moreover, given the distance between two nodes u and v are the same, the probabilities of $\mathcal{E}_z^{u,v} \mid (u, v) \in E$ are different when u and v are in the same cluster and when they are in different clusters. Therefore the count of the common neighbors are going to be different, and substantially separated with high probability for two vertices in cases when they are from the same cluster or from different clusters. However, this may not be the case, if we do not restrict the distance to be the same and look at the entire range of possible distances.

First, we quote two simple lemmas about the expected value of the common neighbors.

Lemma 21. *For any two vertices $u, v \in V_i : (u, v) \in E, i = 1, 2$ belonging to the same cluster with $d_L(X_u, X_v) = x$, the count of common neighbors $C_{u,v} \equiv |\{z \in V : (z, u), (z, v) \in E\}|$ is a random variable distributed according to $\text{Bin}(\frac{n}{2} - 2, 2r_s - x)$ when $r_s \geq x > 2r_d$ and according to $\text{Bin}(\frac{n}{2} - 2, 2r_s - x) + \text{Bin}(\frac{n}{2}, 2r_d - x)$ when $x \leq \min(2r_d, r_s)$, where $\text{Bin}(n, p)$ is a binomial random variable with mean np .*

Proof Without loss of generality, assume $u, v \in V_1$. For any vertex $z \in V$, let $\mathcal{E}_z^{u,v} \equiv \{(u, z), (v, z) \in E\}$ be the event that z is a common neighbor. For $z \in V_1$,

$$\begin{aligned} \Pr(\mathcal{E}_z^{u,v}) &= \Pr((z, u) \in E, (z, v) \in E) \\ &= 2r_s - x, \end{aligned}$$

since $d_L(X_u, X_v) = x$. For $z \in V_2$, we have,

$$\begin{aligned} \Pr(\mathcal{E}_z^{u,v}) &= \Pr((z, u), (z, v) \in E) \\ &= \begin{cases} 2r_d - x & \text{if } x < 2r_d \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

Now since there are $\frac{n}{2} - 2$ points in $V_1 \setminus \{u, v\}$ and $\frac{n}{2}$ points in V_2 , we have the statement of the lemma. $\blacksquare$

In a similar way, we can prove.

Lemma 22. *For any two vertices $u \in V_1, v \in V_2 : (u, v) \in E$ belonging to different clusters with $d_L(X_u, X_v) = x$, the count of common neighbors $C_{u,v} \equiv |\{z \in V : (z, u), (z, v) \in E\}|$ is a random variable distributed according to $\text{Bin}(n - 2, 2r_d)$ when $r_s > 2r_d$ and according to $\text{Bin}(n - 2, \min(r_s + r_d - x, 2r_d))$ when $r_s \leq 2r_d$ and $x \leq r_d$.*

The distribution of the number of common neighbors given $(u, v) \in E$ and $d(u, v) = x$ is given in Table 3. As throughout this paper, we have assumed that there are only two clusters of equal size. The functions change when the cluster sizes are different. In the table, $u \sim v$ means u and v are in the same cluster.

$(u, v) \in E$ $d(u, v) = x$	Distribution of count ($r_s > 2r_d$)		Distribution of count ($r_s \leq 2r_d$)	
	$u \sim v, x \leq r_s$	$u \not\sim v, x \leq r_d$	$u \sim v, x \leq r_s$	$u \not\sim v, x \leq r_d$
$z \mid (z, u) \in E, (z, v) \in E$	$\text{Bin}(\frac{n}{2} - 2, 2r_s - x) + \mathbb{1}\{x \leq 2r_d\} \text{Bin}(\frac{n}{2}, 2r_d - x)$	$\text{Bin}(n - 2, 2r_d)$	$\text{Bin}(\frac{n}{2} - 2, 2r_s - x) + \text{Bin}(\frac{n}{2}, 2r_d - x)$	$\text{Bin}(n - 2, \min(r_s + r_d - x, 2r_d))$

Table 3: Distribution of triangle count for an edge (u, v) conditioned on the distance between them $d(u, v) = d_L(X_u, X_v) = x$, when there are two equal sized clusters. Here $\text{Bin}(n, p)$ denotes a binomial random variable with mean np .

At this point note that, in a $\text{GBM}_1(r_s, r_d)$ for any edge u, v that do not belong to the same part, the expected total number of common neighbors of u and v does not depend on their distance. We will next show that in this case the normalized total number of common neighbors is concentrated around $2r_d$. Therefore, when Algorithm 1 finished removing all the edges, with high probability all the ‘inter-cluster’ edges are removed. However, some of the ‘in-cluster’ edges will also be removed in the process. This is similar to the case when from an $\text{RAG}_1^*(n, [0, r_s])$, all the edges that correspond to a distance close to $2r_d$ has been removed. This situation is shown for the case when $r_s \geq 2r_d$ in Figure 7. Finally we show that the edge-reduced $\text{RAG}_1^*(n, [0, r_s])$ is still connected under certain condition. In what follows we will assume the $\text{GBM}_1(r_s, r_d)$ with $r_s \geq 2r_d$. The other case of $r_s < 2r_d$ is similar.

In the next lemma, we show a concentration result for the count made in process.




Figure 7: Average number of common neighbors of $(u, v) \in E$ for varying values of $d(u, v) = x$ when $r_s \geq 2r_d$.

Lemma 23. Suppose we are given the graph $G(V, E)$ generated according to $\text{GBM}_1(r_s \equiv \frac{a \log n}{n}, r_d \equiv \frac{b \log n}{n})$, $a \geq 2b$. Our algorithm with $E_S = (2b + f_1) \frac{\log n}{n}$ and $E_D = (2b - f_2) \frac{\log n}{n}$, removes all the edges $(u, v) \in E$ such that u and v are in different parts with probability at least $1 - o(1)$, where

$$f_1 = \min\{f : (2b + f) \log \frac{2b + f}{2b} - f > 1\}$$

$$f_2 = \min\{f : (2b - f) \log \frac{2b - f}{2b} + f > 1\}.$$

Proof Here we will use the fact that for $a \geq 1$, the number of edges in $\text{GBM}_1(r_s \equiv \frac{a \log n}{n}, r_d \equiv \frac{b \log n}{n})$ is $O(n \log n)$ with probability $1 - \frac{1}{n^{\Theta(1)}}$. Consider any vertex $u \in V_1$ (symmetrically for $u \in V_2$), since the vertices are thrown uniformly at random in $[0, 1]$, the probability that a $v \in V_1$, $v \neq u$, is a neighbor of u is $\frac{a \log n}{n}$, and for $v \in V_2$, the corresponding probability is $\frac{b \log n}{n}$. Therefore, the expected degree of u is $\frac{(a+b)}{2} \log n$. By a simple Chernoff bound argument, the degree of u is therefore $O(\log n)$ with probability $1 - \frac{1}{n^c}$ for $c \geq 2$. By union bound over all the vertices, the total number of edges is $O(n \log n)$ with probability $1 - \frac{1}{n}$.

Let Z denote the random variable that equals the number of common neighbors of two nodes $u, v \in V : (u, v) \in E$ such that u, v are from different parts of the GBM. Using Lemma 22, we know that Z is sampled from the distribution $\text{Bin}(n - 2, 2r_d)$, where $r_d = \frac{b \log n}{n}$. Therefore,

$$\Pr(Z \geq nE_S) \leq \sum_{i=nE_S}^n \binom{n}{i} (2r_d)^i (n - 2r_d)^{n-i} \leq \exp\left(-nD\left((2b + f_1) \frac{\log n}{n} \parallel \frac{2b \log n}{n}\right)\right),$$

where $D(p \parallel q) \equiv p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ is the KL divergence between Bernoulli(p) and Bernoulli(q) distributions. It is easy to see that,

$$nD\left(\frac{\alpha \log n}{n} \parallel \frac{\beta \log n}{n}\right) = \left(\alpha \log \frac{\alpha}{\beta} + (\alpha - \beta)\right) \log n - o(\log n).$$

Therefore $\Pr(Z \geq nE_S) \leq \frac{1}{n(\log n)^2}$ because $(2b + f_1) \log \frac{2b+f_1}{2b} - f_1 > 1$. Similarly, we have that

$$\Pr(Z \leq nE_D) \leq \sum_{i=0}^{nE_D} \binom{n}{i} (2r_d)^i (n - 2r_d)^{n-i} \leq \exp(-nD((2b - t) \frac{\log n}{n} \parallel \frac{2b \log n}{n})) \leq \frac{1}{n(\log n)^2}.$$

So all of the inter-cluster edges will be removed by Algorithm 1 with probability $1 - O(\frac{n \log n}{n(\log n)^2}) = 1 - o(1)$, as with probability $1 - o(1)$ the total number of edges in the graph is $O(n \log n)$. $\blacksquare$

After Algorithm 1 finishes, in the edge-reduced $\text{GBM}_1(\frac{a \log n}{n}, \frac{b \log n}{n})$, all the edges are ‘in-cluster’ edges with high probability. However some of the ‘in-cluster’ edges are also deleted, namely, those that has a count of common neighbors between E_S and E_D . In the next two lemmas, we show the necessary condition on the ‘in-cluster’ edges such that they do not get removed by Algorithm 1.

Lemma 24. *Suppose we have the graph $G(V, E)$ generated according to $\text{GBM}_1(r_s \equiv \frac{a \log n}{n}, r_d \equiv \frac{b \log n}{n})$, $a \geq 2b$. Define f_1, f_2, E_D, E_S as in Lemma 23. Consider an edge $(u, v) \in E$ where u, v belong to the same part of the GBM and let $d(u, v) \equiv x \equiv \frac{\theta \log n}{n}$. Suppose θ satisfies **either** of the following conditions:*

1. $\frac{1}{2} \left((4b + 2f_1) \log \frac{4b+2f_1}{2a-\theta} + 2a - \theta - 4b - 2f_1 \right) > 1$ and $\theta \leq 2a - 4b - 2f_1$
2. $\frac{1}{2} \left((4b - 2f_2) \log \frac{4b-2f_2}{2a-\theta} + 2a - \theta - 4b + 2f_2 \right) > 1$ and $a \geq \theta \geq \max\{2b, 2a - 4b + 2f_2\}$.

Then Algorithm 1 with $E_S = (2b + f_1) \frac{\log n}{n}$ and $E_D = (2b - f_2) \frac{\log n}{n}$ will not remove this edge with probability at least $1 - O(\frac{1}{n(\log n)^2})$.

Proof Let Z be the number of common neighbors of u, v . Recall that, u and v are in the same cluster. We know from Lemma 22 that Z is sampled from the distribution $\text{Bin}(\frac{n}{2} - 2, 2r_s - x) + \text{Bin}(\frac{n}{2}, 2r_d - x)$ when $x \leq 2r_d$, and from the distribution $\text{Bin}(\frac{n}{2} - 2, 2r_s - x)$ when $x \geq 2r_d$. We have,

$$\begin{aligned} \Pr(Z \leq nE_S) &= \begin{cases} \sum_{i=0}^{nE_S} \binom{\frac{n}{2}-2}{i} (2r_s - x)^i (1 - 2r_s + x)^{\frac{n}{2}-i-2} \sum_{j=0}^{nE_S-i} \binom{\frac{n}{2}}{j} (2r_d - x)^j (1 - 2r_d + x)^{\frac{n}{2}-j} & \text{if } x \leq 2r_d \\ \sum_{i=0}^{nE_S} \binom{\frac{n}{2}-2}{i} (2r_s - x)^i (1 - 2r_s + x)^{\frac{n}{2}-i} & \text{otherwise} \end{cases} \\ &\leq e^{-\frac{n}{2} D(2E_S \parallel \frac{(2a-\theta) \log n}{n})} \text{ since } 2a - \theta \geq 4b + 2f_1 \\ &\leq e^{-\frac{n}{2} D(\frac{(4b+2f_1) \log n}{n} \parallel \frac{(2a-\theta) \log n}{n})} \leq \frac{1}{n \log^2 n}, \end{aligned}$$

because of condition 1 of this lemma. Therefore, this edge will not be deleted with high probability. Similarly, let us find the probability of $Z \geq nE_D = (2b - f_2) \log n$. Let us just assume the worst case when $\theta \leq 2b$: that the edge is being deleted (see condition 2, this is prohibited if that condition is satisfied). Otherwise, $\theta > 2b$ and,

$$\Pr(Z \geq nE_D) = \sum_{i=nE_D}^n \binom{\frac{n}{2}-2}{i} (2r_s - x)^i (1 - 2r_s + x)^{\frac{n}{2}-i-2}$$

$$\begin{aligned}
&\leq e^{-\frac{n}{2}D(2E_D\|\frac{(2a-\theta)\log n}{n})} \text{ if } 2a - \theta \leq 4b - 2f_2 \\
&= e^{-\frac{n}{2}D(\frac{(4b-2f_2)\log n}{n}\|\frac{(2a-\theta)\log n}{n})} \leq \frac{1}{n \log^2 n},
\end{aligned}$$

because of condition 2 of this lemma. ■

Now we are in a position to prove our main theorem from this part. Let us restate this theorem.

Theorem. (*Sufficient Condition in Theorem 8*) Suppose we have the graph $G(V, E)$ generated according to $\text{GBM}_1(r_s \equiv \frac{a \log n}{n}, r_d \equiv \frac{b \log n}{n})$, $a \geq 2b$. Define,

$$f_1 = \min\{f : (2b + f) \log \frac{2b + f}{2b} - f > 1\}$$

$$f_2 = \min\{f : (2b - f) \log \frac{2b - f}{2b} + f > 1\}$$

$$\theta_1 = \max\{\theta : \frac{1}{2} \left((4b + 2f_1) \log \frac{4b + 2f_1}{2a - \theta} + 2a - \theta - 4b - 2f_1 \right) > 1 \text{ and } 0 \leq \theta \leq 2a - 4b - 2f_1\}$$

$$\theta_2 = \min\{\theta : \frac{1}{2} \left((4b - 2f_2) \log \frac{4b - 2f_2}{2a - \theta} + 2a - \theta - 4b + 2f_2 \right) > 1 \text{ and } a \geq \theta \geq \max\{2b, 2a - 4b + 2f_2\}\}.$$

Then, if $a - \theta_2 + \theta_1 > 2$ or $a - \theta_2 > 1, a > 2$, Algorithm 1 with $E_S = (2b + f_1) \frac{\log n}{n}$ and $E_D = (2b - f_2) \frac{\log n}{n}$ will recover the correct partition in the GBM with probability $1 - o(1)$.

Proof From Lemma 23, we know that after Algorithm 1 goes over all the edges, the edges with end-points being in different parts of the GBM are all removed with probability $1 - o(1)$. There are $O(n \log n)$ edges in the GBM with probability $1 - o(1)$. From Lemma 24, we can say that no edge with both ends at the same part is deleted with probability at least $1 - o(1)$ (by simply applying a union bound).

After Algorithm 1 goes over all the edges, the remaining edges form a disjoint union of two random annulus graphs of $\frac{n}{2}$ vertices each. For any two vertices u, v in the same part, there will be an edge if $d(u, v) \in [0, \theta_1] \cup [\theta_2, a]$. From Corollary 14, it is evident that each of these two parts (each part is of size $\frac{n}{2}$) will be connected if either $a - \theta_2 + \theta_1 > 2$ or $a - \theta_2 > 1, a > 2$. ■

It is also possible to incorporate the result of Corollary 15 as well to get somewhat stronger recovery guarantee for our algorithm.

6. High Dimensional GBM: Proof of Sufficient Condition in Theorem 10

In this section, we show that our algorithm for recovery of clusters in GBM, i.e., Algorithm 1 extends to higher dimensions. Recall the precise definition of the high-dimensional GBM:

Definition 25 (The GBM in High Dimensions). Given $V = V_1 \sqcup V_2, |V_1| = |V_2| = \frac{n}{2}$, choose a random vector X_u independently uniformly distributed in S^t for all $u \in V$. The geometric block model $\text{GBM}_t(r_s, r_d)$ with parameters $r_s > r_d$ is a random graph where an edge exists between vertices u and v if and only if,

$$\begin{aligned}
&\|X_u - X_v\|_2 \leq r_s \text{ when } u, v \in V_1 \text{ or } u, v \in V_2 \\
&\|X_u - X_v\|_2 \leq r_d \text{ when } u \in V_1, v \in V_2 \text{ or } u \in V_2, v \in V_1.
\end{aligned}$$

Indeed, for the higher dimensional case the algorithm remains exactly the same, except the value of E_D and E_S in the subroutine `process` needs to be changed. Recall that the algorithm proceeds by checking each edge and counting the number of triangle the edge is part of. If the count is between E_D and E_S the edge is removed. In this process we claim to remove all inter-cluster edges with high probability. The main difficulty lies in proving that the original communities remain connected in the redacted graph. For that we crucially use the connectivity results of the high dimensional random annulus graphs (from Section 4) in somewhat different way that what we do for the one dimensional case.

6.1 Analysis of Algorithm 1 in High Dimension

Let us define a few more terminologies to simplify the expressions for high dimensional space. The volume of a t -sphere with unit radius is $a_t = \frac{2\pi^{t+1/2}}{\Gamma(\frac{t+1}{2})}$. Let the spherical cap $B_t(O, r) \subset S^t$ define a region on the surface of this t -sphere S^t such that every point $u \in B_t(O, r)$ satisfies $\|u - O\|_2 \leq r$. Let us denote the volume of the spherical cap $B_t(O, r)$ normalized with a_t by $B_t(r)$. Similarly $B_t(O, [r_1, r_2])$ refers to a region on the t -sphere such that every point $u \in B_t(O, [r_1, r_2])$ satisfies $r_1 \leq \|u - O\|_2 \leq r_2$ and $B_t(r_1, r_2)$ refers to the volume normalized by a_t . Now consider two such spherical caps $B_t(O_1, r_1)$ and $B_t(O_2, r_2)$ such that $d(O_1, O_2) = \ell$. In that case let us define the volume of the intersection of the two aforementioned spherical caps (again normalized by a_t) by $\mathcal{V}_t(r_1, r_2, \ell)$.

Let us use $u \sim v$ ($u \approx v$) to denote u and v belong to the same cluster (different clusters). Let $\mathcal{E}_z^{u,v}$ denote the event that z is a common neighbor of u and v and $e(u, v)$ denote the event that there is an edge between u and v . Following are some simple observations.

Observation 26. $\Pr(e(u, v) \mid u \sim v) = B_t(r_s)$ and $\Pr(e(u, v) \mid u \approx v) = B_t(r_d)$.

Observation 27. $\Pr(\mathcal{E}_z^{u,v} \mid z \sim u, u \sim v \text{ and } \|u - v\|_2 = \ell) = \mathcal{V}_t(r_s, r_s, \ell)$ and $\Pr(\mathcal{E}_z^{u,v} \mid z \approx u, u \sim v \text{ and } \|u - v\|_2 = \ell) = \mathcal{V}_t(r_d, r_d, \ell)$.

In the following proof, we assume $r_s \leq 2r_d$. The other situation where the gap between r_s and r_d is higher is only easier to handle.

Lemma 28. *For any two vertices $u, v \in V_i : (u, v) \in E, i = 1, 2$ such that $d(u, v) = \ell$ belonging to the same cluster, the count of common neighbors $C_{u,v} \equiv |\{z \in V : (z, u), (z, v) \in E\}|$ is a random variable distributed according to $\text{Bin}(\frac{n}{2} - 2, \mathcal{V}_t(r_s, r_s, \ell))$ when $r_s \geq \ell > 2r_d$ and according to $\text{Bin}(\frac{n}{2} - 2, \mathcal{V}_t(r_s, r_s, \ell)) + \text{Bin}(\frac{n}{2}, \mathcal{V}_t(r_d, r_d, \ell))$ when $\ell \leq 2r_d$.*

Lemma 29. *For any two vertices $u \in V_1, v \in V_2 : (u, v) \in E$ such that $\|u - v\|_2 = \ell$ belonging to different clusters, the count of common neighbors $C_{u,v} \equiv |\{z \in V : (z, u), (z, v) \in E\}|$ is a random variable distributed according to $\text{Bin}(n - 2, B_t(r_d))$ when $r_s > 2r_d$ and according to $\text{Bin}(n - 2, \min(\mathcal{V}_t(r_s, r_d, \ell), B_t(r_d)))$ when $r_s \leq 2r_d$ and $\ell \leq r_d$.*

Proof [Proof of Lemma 28] Without loss of generality, assume $u, v \in V_1$. In order for $(u, v) \in E$, we must have $r_s \geq \ell$. Now there are two cases to consider, $\ell > 2r_d$ and $\ell \leq 2r_d$. In case 1, for z to be a common neighbor of u and v , z must be in V_1 by triangle inequality. Since, there are $\frac{n}{2} - 2$ points in $V_1 \setminus \{u, v\}$, from Observation 27, $\mathcal{E}_z^{u,v} \sim \text{Bin}(\frac{n}{2} - 2, \mathcal{V}_t(r_s, r_s, \ell))$. In case 2, z can also be part of V_2 and there are $\frac{n}{2}$ points in V_2 , thus again from Observation 27, $\mathcal{E}_z^{u,v} \sim \text{Bin}(\frac{n}{2} - 2, \mathcal{V}_t(r_s, r_s, \ell)) + \text{Bin}(\frac{n}{2}, \mathcal{V}_t(r_d, r_d, \ell))$. $\blacksquare$

The proof of Lemma 29 is similar. We now use the following version of the Chernoff bound to estimate the deviation on the number of common neighbors in the two cases: $u \sim v$ and $u \approx v$.

Lemma 30 (Chernoff Bound). *Let $X_1, \dots, X_n$ be iid random variables in $\{0, 1\}$. Let X denote the sum of these n random variables. Then for any $\delta > 0$,*

$$\begin{cases} \Pr(X > (1 + \delta)\mathbb{E}(X)) \leq e^{-\delta^2\mathbb{E}(X)/3} = \frac{1}{n \log^2 n}, & \text{when } \delta = \sqrt{\frac{3(\log n + 2 \log \log n)}{\mathbb{E}(X)}}, \\ \Pr(X < (1 - \delta)\mathbb{E}(X)) \leq e^{-\delta^2\mathbb{E}(X)/2} = \frac{1}{n \log^2 n}, & \text{when } \delta = \sqrt{\frac{2(\log n + 2 \log \log n)}{\mathbb{E}(X)}}. \end{cases}$$

We take $E_S = c_s^{(t)} \cdot (B_t(r_d)n + \sqrt{6B_t(r_d)n \log n})$ and $E_D = c_d^{(t)} \cdot (n\mathcal{V}_t(r_s, r_d, r_d) - \sqrt{2nB_t(r_d) \log n})$ where $c_s^{(t)} \geq 1$ and $c_d^{(t)} \leq 1$ are suitable constants that depend on t .

Lemma 31. *For any pair of nodes $(u, v) = e \in E$, $u \approx v$, the process algorithm removes the edge e with a probability of $1 - O\left(\frac{1}{n \log^2 n}\right)$ when $E_S \geq B_t(r_d)n + \sqrt{6B_t(r_d)n \log n}$ and $E_D \leq n\mathcal{V}_t(r_s, r_d, r_d) - \sqrt{2nB_t(r_d) \log n}$.*

Proof Let Z denote the random variable for the number of common neighbors of two nodes $u, v \in V : (u, v) \in E, \|u - v\|_2 = \ell, u \approx v$. From Lemma 29, $E[Z] \leq (n - 2)B_t(r_d)$. Using the Chernoff bound we know that with a probability of at least $1 - \frac{1}{n \log^2 n}$

$$\begin{aligned} Z &\leq F_\infty = (n - 2)B_t(r_d) + \sqrt{3(\log n + 2 \log \log n)(n - 2)B_t(r_d)} \\ &= B_t(r_d)n + \sqrt{3B_t(r_d)n \log n} + o(1) \\ &\leq E_S. \end{aligned}$$

Moreover again from Lemma 29, $E[Z] = (n - 2) \min(\mathcal{V}_t(r_s, r_d, \ell), B_t(r_d))$ as we assume $r_s \leq 2r_d$. Hence, with probability of at least $1 - \frac{1}{n \log^2 n}$

$$\begin{aligned} Z &\geq f_\infty = \min_{\ell: \ell \leq r_d, r_s \leq 2r_d} ((n - 2) \min(\mathcal{V}_t(r_s, r_d, \ell), B_t(r_d)) - \\ &\quad \sqrt{2(\log n + 2 \log \log n)(n - 2) \min(\mathcal{V}_t(r_s, r_d, \ell), B_t(r_d))}) \\ &\geq \min_{\ell: \ell \leq r_d, r_s \leq 2r_d} ((n - 2) \min(\mathcal{V}_t(r_s, r_d, \ell), B_t(r_d)) - \\ &\quad \sqrt{2(\log n + 2 \log \log n)(n - 2)B_t(r_d)}) \text{ since } \mathcal{V}_t(r_s, r_d, \ell) \subseteq B_t(r_d) \\ &> n\mathcal{V}_t(r_s, r_d, r_d) - \sqrt{2nB_t(r_d) \log n} \text{ since } \mathcal{V}_t(r_s, r_d, \ell) \text{ is a decreasing function of } \ell \\ &\geq E_D. \end{aligned}$$

Hence, $E_S \leq Z \leq E_D$ with a probability of $1 - \frac{2}{n \log^2 n}$ for $(u, v) \in E, u \approx v$. Hence (u, v) gets removed with high probability by the algorithm. $\blacksquare$

Applying a union bound, we therefore can assume all inter-cluster edges are removed with probability $1 - o(1)$ as there is $O(n \log n)$ edges.

In the next two lemmas, we provide two different conditions on $\|u - v\|_2$ when $u \sim v$ such that our algorithm does not remove the edge (u, v) . Then we obtain a sufficient condition for the two communities to remain connected by the edges that are not removed.

Lemma 32. *Given a pair of nodes u, v belonging to the same cluster such that $(u, v) \in E$, the process algorithm does not remove the edge e with probability of $1 - O\left(\frac{1}{n \log^2 n}\right)$ when $\|u - v\|_2 = \ell$ (say) satisfies the following:*

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell) \right) - \sqrt{2n \log n} \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right) > E_S.$$

Proof Let Z denote the random variable corresponding to the number of common neighbors of u, v . Let $\mu_s(\ell) = \mathbb{E}(Z | u \sim v, d(u, v) = \ell)$. From Lemma 28, $\mu_s(\ell) = (\frac{n}{2} - 2)\mathcal{V}_t(r_s, r_s, \ell) + \frac{n}{2}\mathcal{V}_t(r_d, r_d, \ell)$.

Using the Chernoff bound, with a probability of $1 - O\left(\frac{1}{n \log^2 n}\right)$

$$\begin{aligned} Z &> (n/2 - 2)\mathcal{V}_t(r_s, r_s, \ell) + n/2\mathcal{V}_t(r_d, r_d, \ell) - \sqrt{2(\log n + 2 \log \log n)\mathcal{V}_t(r_s, r_s, \ell)n/2} \\ &\quad - \sqrt{2(\log n + 2 \log \log n)\mathcal{V}_t(r_d, r_d, \ell)n/2} \\ &\geq n/2\mathcal{V}_t(r_s, r_s, \ell) + n/2\mathcal{V}_t(r_d, r_d, \ell) - \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right) \sqrt{2n \log n} \\ &= \frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell) \right) - \sqrt{2n \log n} \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right). \end{aligned}$$

Therefore, Algorithm 1 will not delete e if

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell) \right) - \sqrt{2n \log n} \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right) > E_S.$$

Note that there exists a maximum value of distance (referred to as ℓ_1) such that whenever $\|u - v\| \leq \ell_1$, the condition will be satisfied. $\blacksquare$

Lemma 33. *Given a pair of nodes u, v belonging to the same cluster such that $(u, v) \in E$, the process algorithm does not remove the edge e with probability of $1 - O\left(\frac{1}{n \log^2 n}\right)$ when $\ell \equiv \|u - v\|_2$ (say) satisfies the following:*

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell + \mathcal{V}_t(r_d, r_d, \ell)) \right) + \sqrt{n \log n} \sqrt{[\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell)]} \leq E_D.$$

Proof

Let Z denote the random variable corresponding to the number of common neighbors of u, v . Let $\mu_s(\ell) = \mathbb{E}(Z | u \sim v, \|u - v\|_2 = \ell)$. From Lemma 28, $\mu_s(\ell) = (\frac{n}{2} - 2)\mathcal{V}_t(r_s, r_s, \ell) + \frac{n}{2}\mathcal{V}_t(r_d, r_d, \ell)$.

Using the Chernoff bound, with a probability of $1 - O\left(\frac{1}{n \log^2 n}\right)$

$$\begin{aligned} Z &< n/2\mathcal{V}_t(r_s, r_s, \ell) + n/2\mathcal{V}_t(r_d, r_d, \ell) + \sqrt{2(\log n + 2 \log \log n)[\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell)](n/2)} \\ &\leq \frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell + \mathcal{V}_t(r_d, r_d, \ell)) \right) + \sqrt{n \log n} \sqrt{[\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell)]}. \end{aligned}$$

The process algorithm will not remove e if

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell + \mathcal{V}_t(r_d, r_d, \ell)) \right) + \sqrt{n \log n} \sqrt{[\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell)]} \leq E_D.$$

Note that there exists a minimum value of distance (referred to as ℓ_2) such that whenever $\|u - v\|_2 \geq \ell_2$, the condition will be satisfied. $\blacksquare$

Lemma 34. *Algorithm 1 can identify all edges (u, v) correctly for which $\ell \equiv \|u - v\|_2$ satisfies either of the following:*

$$\text{Cond.1 : } \frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell) \right) - \sqrt{2n \log n} \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right) > E_S$$

or

$$\text{Cond.2 : } \frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell) \right) + \sqrt{n \log n} \sqrt{[\mathcal{V}_t(r_s, r_s, \ell) + \mathcal{V}_t(r_d, r_d, \ell)]} \leq E_D.$$

with probability at least $1 - O\left(\frac{1}{\log n}\right)$.

Proof Follows from combining Lemma 32 and Lemma 33, and noting that in the connectivity regime, the number of edges is $O(n \log n)$. ■

Let ℓ_1 be the maximum value of $\|u - v\|_2$ such that Cond 1 is satisfied and ℓ_2 is the minimum value of $\|u - v\|_2$ such that Cond 2 is satisfied. Also note that $\ell_1 \leq \ell_2$. We now give a condition on ℓ_1 and ℓ_2 such the two communities are each connected by the edges (u, v) that satisfy either $\|u - v\|_2 \leq \ell_1$ or $\ell_2 \leq \|u - v\|_2 \leq r_s$.

Lemma 35. *If $(\ell_1/2)^t > 8(t+1)\psi(t)\frac{\log n}{n}$ then the edges e that satisfy $\|u - v\|_2 \leq \ell_1$ constitute two disjoint connected components corresponding to the two original communities.*

Proof Proof of this lemma follows from the result of connectivity of random annulus graphs (RAG) in dimension t , i.e., Theorem 5. ■

We now find out the values of r_s and r_d such that ℓ_1 and ℓ_2 satisfy the condition of Lemma 35 as well as Cond. 1 and Cond 2. respectively.

Theorem. *(Sufficient Condition in Theorem 10) If $r_s = \Theta\left(\left(\frac{\log n}{n}\right)^{\frac{1}{t}}\right)$ and $r_s - r_d = \Omega\left(\left(\frac{\log n}{n}\right)^{\frac{1}{t}}\right)$, Algorithm 1 recovers the clusters with probability $1 - o(1)$.*

Proof Let us take $r_s = a_t \left(\frac{\log n}{n}\right)^{\frac{1}{t}}$, $r_d \leq b_t \left(\frac{\log n}{n}\right)^{\frac{1}{t}}$ for some large constants a_t and b_t that depends on t . Then to satisfy Lemma 35, we can take $\ell_1 = a'_t \left(\frac{\log n}{n}\right)^{\frac{1}{t}}$ and $\ell_2 = b'_t \left(\frac{\log n}{n}\right)^{\frac{1}{t}}$ again for suitable constants a'_t and b'_t . While it is possible to concisely compute $B_t(r)$ and $\mathcal{V}_t(r_1, r_2, x)$ (Li, 2011; Ellis et al., 2007), for the purpose of analysis it is sufficient to know $B_t(r) = \Theta(r^t)$ for fixed t . Moreover, $\mathcal{V}_t(r_1, r_2, x) = \Theta((r_1 + r_2 - x)^t)$ if $r_1 + r_2 \geq x \geq \max(r_1, r_2)$, $\Theta(\min\{r_1, r_2\}^t)$ if $x \leq \max(r_1, r_2)$ and 0 otherwise.

Then, Cond 1. requires

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell_1) + \mathcal{V}_t(r_d, r_d, \ell_1) \right) - \sqrt{2n \log n} \left(\sqrt{B_t(r_s)} + \sqrt{B_t(r_d)} \right) > E_S$$

and Cond 2. requires

$$\frac{n}{2} \left(\mathcal{V}_t(r_s, r_s, \ell_2) + \mathcal{V}_t(r_d, r_d, \ell_2) \right) + \sqrt{n \log n} \sqrt{[\mathcal{V}_t(r_s, r_s, \ell_2) + \mathcal{V}_t(r_d, r_d, \ell_2)]} \leq E_D.$$

By selecting the constants $c_s^{(t)}$ and c_d^t involving E_S and E_D suitably, both the conditions are satisfied. ■

7. Experimental Results

In addition to validation experiments in Section 1.1, we also conducted an in-depth experimentation of our proposed model and techniques over a set of synthetic and real world networks. Additionally, we compared the efficacy and efficiency of our triangle-counting algorithm with the popular spectral clustering algorithm using normalized cuts⁴ and the correlation clustering algorithm (Bansal et al., 2004).

Dataset	Total no. of nodes	T_1	T_2	T_3	Accuracy		Running Time (sec)	
					Triangle-Counting	Spectral clustering	Triangle-Counting	Spectral clustering
Political Blogs	1222	20	2	1	0.788	0.53	1.62	0.29
DBLP	12138	10	1	2	0.675	0.63	3.93	18.077
LiveJournal	2366	20	1	1	0.7768	0.64	0.49	1.54

Table 4: Performance on real world networks

Real Datasets. We use three real datasets described below.

- **Political Blogs.** (Adamic and Glance, 2005) It contains a list of political blogs from 2004 US Election classified as liberal or conservative, and links between the blogs. The clusters are of roughly the same size with a total of 1200 nodes and 20K edges.
- **DBLP.** (Yang and Leskovec, 2015) The DBLP dataset is a collaboration network where the ground truth communities are defined by the research community. The original graph consists of roughly 0.3 million nodes. We process it to extract the top two communities of size ~ 4500 and 7500 respectively. This is given as input to our algorithm.
- **LiveJournal.** (Leskovec et al., 2007) The LiveJournal dataset is a free online blogging social network of around 4 million users. Similar to DBLP, we extract the top two clusters of sizes 930 and 1400 which consist of around 11.5K edges.


We have not used the academic collaboration (Section 1.1) dataset here because it is quite sparse and below the connectivity threshold regime of both GBM and SBM.

Synthetic Datasets. We generate synthetic datasets of different sizes according to the GBM with $t = 2, k = 2$ and for a wide spectrum of values of r_s and r_d , specifically we focus on the sparse region where $r_s = \frac{a \log n}{n}$ and $r_d = \frac{b \log n}{n}$ with variable values of a and b .

Experimental Setting. For real networks, it is difficult to calculate an exact threshold as the exact values of r_s and r_d are not known. Hence, we follow a three step approach. Using a somewhat large threshold T_1 we sample a subgraph S such that u, v will be in S if there is an edge between u and v , and they have at least T_1 common neighbors. We now attempt to recover the subclusters inside this subgraph by following our algorithm with a small threshold T_2 . Finally, for nodes that are not part of S , say $x \in V \setminus S$, we select each $u \in S$ that x has an edge with and use a threshold of T_3 to decide if u and x should be in the same cluster. The final decision is made by taking a majority vote. We can employ sophisticated methods over this algorithm to improve the results further, which is beyond the scope of this work.

We use the popular f-score metric which is the harmonic mean of precision (fraction of number of pairs correctly classified to total number of pairs classified into clusters) and recall (fraction of number of pairs correctly classified to the total number of pairs in the same cluster for ground truth),

4. <http://scikit-learn.org/stable/modules/clustering.html#spectral-clustering>



(a) f-score with varying a , fixing b . (b) Fraction of nodes misclassified.

Figure 8: Results of the triangle-counting algorithm on a synthetic dataset with 5000 nodes.

as well as the node error rate for performance evaluation. A node is said to be misclassified if it belongs to a cluster where the majority comes from a different ground truth cluster (breaking ties arbitrarily). Following this, we use the above described metrics to compare the performance of different techniques on various datasets.

Results. We compared our algorithm with the spectral clustering algorithm where we extracted two eigenvectors in order to extract two communities. Table 4 shows that our algorithm gives an accuracy as high as 78%. The spectral clustering performed worse compared to our algorithm for all real world datasets. It obtained the worst accuracy of 53% on political blogs dataset. The correlation clustering algorithm generates various small sized clusters leading to a very low recall, performing much worse than the triangle-counting algorithm for the whole spectrum of parameter values.

We can observe in Table 4 that our algorithm is much faster than the spectral clustering algorithm for larger datasets (LiveJournal and DBLP). This confirms that triangle-counting algorithm is more scalable than the spectral clustering algorithm. The spectral clustering algorithm also works very well on synthetically generated SBM networks even in the sparse regime (Lei et al., 2015; Rohe et al., 2011). The superior performance of the simple triangle clustering algorithm over the real networks provide a further validation of GBM over SBM. Correlation clustering takes 8-10 times longer as compared to triangle-counting algorithm for the various range of its parameters. We also compared our algorithm with the Newman algorithm (Girvan and Newman, 2002) that performs really well for the LiveJournal dataset (98% accuracy). But it is extremely slow and performs much worse on other datasets. This is because the LiveJournal dataset has two well defined subsets of vertices with very few intercluster edges. The reason for the worse performance of our algorithm is the sparseness of the graph. If we create a subgraph by removing all nodes of degrees 1 and 2, we get 100% accuracy with our algorithm. Finally, our algorithm is easily parallelizable to achieve better improvements. This clearly establishes the efficiency and effectiveness of triangle-counting.

We observe similar gains on synthetic datasets. Figures 8a and 8b report results on the synthetic datasets with 5000 nodes. Empirically, our results demonstrate much superior performance of our algorithm as compared to theoretical guarantees. The empirical results are much better than the theoretical bounds because the concentration inequalities assume the worst value of the distance between the pair of vertices that are under consideration. We also see a clear threshold behavior on both f-score and node error rate in Figures 8a and 8b. We have also performed spectral clustering on




Figure 9: Results of the spectral clustering on a synthetic dataset with 5000 nodes.




Figure 10: Fraction of triangles (intra-cluster and inter-cluster) in the ground truth of a graph generated according to GBM_1 and the clusters returned by spectral clustering for varying $r_d = \frac{b \log n}{n}$. The number of nodes is 1000 and $r_s = 10 \log n / n$ is fixed. Since r_s is fixed, the fraction of intra-cluster triangles in the ground truth remains unchanged.

this 5000-node synthetic dataset (Figures 9a and 9b). Compared to the plots of figures 8a and 8b, they show suboptimal performance, indicating the relative ineffectiveness of spectral clustering in GBM compared to the triangle counting algorithm.

We ran another experiment to test whether spectral clustering techniques capture varied number of intra-cluster and inter-cluster triangles and if it can be used for datasets where higher order structures are more prevalent. We tested the fraction of intra-cluster triangles returned by spectral clustering for datasets generated by GBM with varying inter-cluster connectivity threshold (r_d) and compared it against the fraction of intra-cluster triangles in the ground truth clustering. Figure 10 shows that the clusters returned by spectral clustering have a very low fraction of intra-cluster triangles and thus it is not designed for clustering based on triangles (or higher-order structures in general). In fact the clusters returned by spectral clustering have similar fraction of intra-cluster and inter-cluster triangles. This evaluation also shows that GBM indeed has a large disparity between inter and intra cluster triangles and it may be useful to model datasets that contain higher order structures.

Acknowledgements: This work is supported by NSF CCF awards 2133484, 2217058, 2223282.

References

- Emmanuel Abbe and Colin Sandon. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In *56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 670–688. IEEE, 2015a.
- Emmanuel Abbe and Colin Sandon. Recovering communities in the general stochastic block model without knowing the parameters. In *Advances in Neural Information Processing Systems*, pages 676–684, 2015b.
- Emmanuel Abbe, Afonso S. Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Trans. Information Theory*, 62(1):471–487, 2016.
- Lada A Adamic and Natalie Glance. The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43. ACM, 2005.
- Uri Alon. Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6): 450–461, 2007.
- Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine learning*, 56(1-3): 89–113, 2004.
- Austin R Benson, David F Gleich, and Jure Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016.
- Béla Bollobás. Random graphs. In *Modern Graph Theory*, pages 215–252. Springer, 1998.
- Béla Bollobás. Random graphs. *Cambridge Press*, 2001.
- Béla Bollobás. Percolation. *Cambridge Press*, 2006.
- Sébastien Bubeck, Jian Ding, Ronen Eldan, and Miklós Z Rácz. Testing for high-dimensional geometry in random graphs. *Random Structures & Algorithms*, 2016.
- Peter Chin, Anup Rao, and Van Vu. Stochastic block model and community detection in the sparse graphs: A spectral algorithm with optimal rate of recovery. *arXiv:1501.05021*, 2015.
- Harold Scott Macdonald Coxeter. *Regular polytopes*. Courier Corporation, 1973.
- Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- Carl P Dettmann and Orestis Georgiou. Random geometric graphs with general connection functions. *Physical Review E*, 93(3):032313, 2016.
- Martin E. Dyer and Alan M. Frieze. The solution of some random np-hard problems in polynomial expected time. *Journal of Algorithms*, 10(4):451–489, 1989.
- David Easley, Jon Kleinberg, et al. Networks, crowds, and markets. *Cambridge Books*, 2012.

- Robert B Ellis, Jeremy L Martin, and Catherine Yan. Random geometric graph diameter in the unit ball. *Algorithmica*, 47(4):421–438, 2007.
- Paul Erdős and Alfréd Rényi. On random graphs, i. *Publicationes Mathematicae (Debrecen)*, 6: 290–297, 1959.
- Stephen Eubank, Hasan Guclu, VS Anil Kumar, Madhav V Marathe, Aravind Srinivasan, Zoltan Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180, 2004.
- Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proceedings of the American mathematical Society*, 124(10):2993–3002, 1996.
- Sainyam Galhotra, Arya Mazumdar, Soumyabrata Pal, and Barna Saha. The geometric block model. In *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, 2018.
- Sainyam Galhotra, Arya Mazumdar, Soumyabrata Pal, and Barna Saha. Connectivity of random annulus graphs and the geometric block model. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- Edgar N Gilbert. Random graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- Edward N Gilbert. Random plane networks. *Journal of the Society for Industrial and Applied Mathematics*, 9(4):533–543, 1961.
- Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the national academy of sciences*, 99(12):7821–7826, 2002.
- Martin Haenggi, Jeffrey G Andrews, François Baccelli, Olivier Dousse, and Massimo Franceschetti. Stochastic geometry and random graphs for the analysis and design of wireless networks. *IEEE Journal on Selected Areas in Communications*, 27(7), 2009.
- Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory*, 62(5):2788–2797, 2016.
- Bruce E. Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 899–928, 2015. URL <http://jmlr.org/proceedings/papers/v40/Hajek15.html>.
- David Haussler and Emo Welzl. epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2(2):127–151, 1987.
- Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- Mihail N Kolountzakis, Gary L Miller, Richard Peng, and Charalampos E Tsourakakis. Efficient triangle counting in large graphs via degree-based vertex partitioning. *Internet Mathematics*, 8 (1-2):161–185, 2012.

- Peter Mahler Larsen and Skuren Schmidt. Improved orientation sampling for indexing diffraction patterns of polycrystalline materials. *Journal of Applied Crystallography*, 50(6):1571–1582, 2017.
- Jing Lei, Alessandro Rinaldo, et al. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, 43(1):215–237, 2015.
- Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.
- Jure Leskovec, Kevin J Lang, Anirban Dasgupta, and Michael W Mahoney. Statistical properties of community structure in large social and information networks. In *17th international conference on World Wide Web*, pages 695–704. ACM, 2008.
- Shengqiao Li. Concise formulas for the area and volume of a hyperspherical cap. *Asian Journal of Mathematics and Statistics*, 2011.
- Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- Elchanan Mossel, Joe Neeman, and Allan Sly. Consistency thresholds for the planted bisection model. In *47th Annual ACM Symposium on Theory of Computing*, pages 69–75. ACM, 2015.
- S Muthukrishnan and Gopal Pandurangan. The bin-covering technique for thresholding random geometric graph properties. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 989–998. Society for Industrial and Applied Mathematics, 2005.
- Mark EJ Newman, Duncan J Watts, and Steven H Strogatz. Random graph models of social networks. *Proceedings of the national academy of sciences*, 99(suppl 1):2566–2572, 2002.
- Hae-Jeong Park and Karl Friston. Structural and functional brain networks: from connections to cognition. *Science*, 342(6158), 2013.
- Mathew Penrose. *Random geometric graphs*. Number 5. Oxford University Press, 2003.
- Mathew D Penrose. Connectivity of soft random geometric graphs. *The Annals of Applied Probability*, 26(2):986–1028, 2016.
- Karl Rohe, Sourav Chatterjee, Bin Yu, et al. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- Martin Rosvall, Alcides V Esquivel, Andrea Lancichinetti, Jevin D West, and Renaud Lambiotte. Memory in network flows and its effects on spreading dynamics and community detection. *Nature communications*, 5(1):1–13, 2014.
- Abishek Sankararaman and François Baccelli. Community detection on euclidean random graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2181–2200. SIAM, 2018.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

- Daniel B Stouffer, Marta Sales-Pardo, M Irmak Sirer, and Jordi Bascompte. Evolutionary conservation of species' roles in food webs. *Science*, 335(6075):1489–1492, 2012.
- Charalampos E Tsourakakis, Mihail N Kolountzakis, and Gary L Miller. Approximate triangle counting. *arXiv preprint arXiv:0904.3761*, 2009.
- Charalampos E Tsourakakis, Jakub Pachocki, and Michael Mitzenmacher. Scalable motif-aware graph clustering. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1451–1460, 2017.
- Arunachalam Vinayagam, Travis E Gibson, Ho-Joon Lee, Bahar Yilmazel, Charles Roesel, Yanhui Hu, Young Kwon, Amitabh Sharma, Yang-Yu Liu, Norbert Perrimon, et al. Controllability analysis of the directed human protein interaction network identifies disease genes and drug targets. *Proceedings of the National Academy of Sciences*, 113(18):4976–4981, 2016.
- Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998.
- Jaewon Yang and Jure Leskovec. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213, 2015.
- Stephen J Young and Edward R Scheinerman. Random dot product graph models for social networks. In *International Workshop on Algorithms and Models for the Web-Graph*, pages 138–149. Springer, 2007.
- Weituo Zhang, Chjan C Lim, Gyorgy Korniss, and Boleslaw K Szymanski. Opinion dynamics and influencing on random geometric graphs. *Scientific reports*, 4:5568, 2014.

Appendix A. Connectivity of one dimensional Random Annulus graphs: Details

In this section, we prove the necessary and sufficient condition for connectivity of RAG_1^* in full details. We provide the necessary conditions for connectivity in Theorem 2 for RAG_1^* because of the equivalence between RAG_1^* and RAG_1 .

A.1 Necessary condition for connectivity of RAG_1^*

Theorem (RAG_1^* connectivity lower bound in Theorem 2). *If $a < 1$ or $a - b < 0.5$, then with probability $1 - o(1)$ the random annulus graph $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is not connected.*

Proof First of all, it is known that $\text{RAG}_1^*(n, [0, \frac{a \log n}{n}])$ is not connected with high probability when $a < 1$ (Muthukrishnan and Pandurangan, 2005; Penrose, 2003). Therefore $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ must not be connected with high probability when $a < 1$ as the connectivity interval is a strict subset of the previous case, and $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ can be obtained from $\text{RAG}_1^*(n, [0, \frac{a \log n}{n}])$ by deleting all the edges that has the two corresponding random variables separated by distance less than $\frac{b \log n}{n}$.

Next we will show that if $a - b < 0.5$ then there exists an isolated vertex with high probability. It would be easier to think of each vertex as a uniform random point in $[0, 1]$. Define an indicator

variable A_u for every node u which is 1 when node u is isolated and 0 otherwise. We have,

$$\Pr(A_u = 1) = \left(1 - \frac{2(a-b) \log n}{n}\right)^{n-1}.$$

Define $A = \sum_u A_u$, and hence

$$\mathbb{E}[A] = n \left(1 - \frac{2(a-b) \log n}{n}\right)^{n-1} = n^{1-2(a-b)-o(1)}.$$

Therefore, when $a - b < 0.5$, $\mathbb{E}[A] = \Omega(1)$. To prove this statement with high probability we can show that the variance of A is bounded. Since A is a sum of indicator random variables, we have that

$$\text{Var}(A) \leq \mathbb{E}[A] + \sum_{u \neq v} \text{Cov}(A_u, A_v) = \mathbb{E}[A] + \sum_{u \neq v} (\Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1)).$$

Now, consider the scenario when the vertices u and v are at a distance more than $\frac{2a \log n}{n}$ apart (happens with probability $1 - \frac{4a \log n}{n}$). Then the region in $[0, 1]$ that is between distances $\frac{b \log n}{n}$ and $\frac{a \log n}{n}$ from both of the vertices is empty and therefore $\Pr(A_u = 1 \cap A_v = 1) = \left(1 - \frac{4(a-b) \log n}{n}\right)^{n-2}$. When the vertices are within distance $\frac{2a \log n}{n}$ of one another, then $\Pr(A_u = 1 \cap A_v = 1) \leq \Pr(A_u = 1)$. Therefore,

$$\begin{aligned} \Pr(A_u = 1 \cap A_v = 1) &\leq \left(1 - \frac{4a \log n}{n}\right) \left(1 - \frac{4(a-b) \log n}{n}\right)^{n-2} + \frac{4a \log n}{n} \Pr(A_u = 1) \\ &\leq \left(1 - \frac{4a \log n}{n}\right) n^{-4(a-b)+o(1)} + \frac{4a \log n}{n} n^{-2(a-b)+o(1)}. \end{aligned}$$

Consequently for large enough n ,

$$\begin{aligned} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1) &\leq \left(1 - \frac{4a \log n}{n}\right) n^{-4(a-b)+o(1)} \\ &\quad + \frac{4a \log n}{n} n^{-2(a-b)+o(1)} - n^{-4(a-b)+o(1)} \leq \frac{8a \log n}{n} \Pr(A_u = 1). \end{aligned}$$

Now,

$$\text{Var}(A) \leq \mathbb{E}[A] + \binom{n}{2} \frac{8a \log n}{n} \Pr(A_u = 1) \leq \mathbb{E}[A](1 + 4a \log n).$$

By using Chebyshev bound, with probability at least $1 - \frac{1}{\log n}$,

$$A > n^{1-2(a-b)} - \sqrt{n^{1-2(a-b)}(1 + 4a \log n) \log n},$$

which implies for $a - b < 0.5$, there will exist isolated nodes with high probability. ■

A.2 Proof of Lemma 12 and Lemma 13

Lemma (Restatement of Lemma 12). *A set of vertices $\mathcal{C} \subseteq V$ is called a cover of $[0, 1]$, if for any point y in $[0, 1]$ there exists a vertex $v \in \mathcal{C}$ such that $d(v, y) \leq \frac{a \log n}{2n}$. If $a - b > 0.5$ and $a > 1$, then a random annulus graph $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$ is a union of cycles such that every cycle forms a cover of $[0, 1]$ (see Figure 6) with probability $1 - o(1)$.*

Proof [Proof of Lemma 12] The proof of this lemma is somewhat easily explained if we consider a weaker result (a stronger condition) with $a - b > 2/3$. Let us first briefly describe this case.

Consider a node u and assume without loss of generality that the position of u is 0 (i.e. $X_u = 0$). Associate four indicator $\{0, 1\}$ -random variables $A_u^i, i = 1, 2, 3, 4$ which take the value of 1 if and only if there does not exist any node x such that

1. $d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [0, \frac{a-b}{2} \frac{\log n}{n}]$ for $i = 1$
2. $d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [\frac{-a-b}{2} \frac{\log n}{n}, -b \frac{\log n}{n}]$ for $i = 2$
3. $d(u, x) \in [-a \frac{\log n}{n}, -b \frac{\log n}{n}] \cup [\frac{-a+b}{2} \frac{\log n}{n}, 0]$ for $i = 3$
4. $d(u, x) \in [-a \frac{\log n}{n}, -b \frac{\log n}{n}] \cup [b \frac{\log n}{n}, \frac{a+b}{2} \frac{\log n}{n}]$ for $i = 4$.

The intervals representing these random variables are shown in Figure 11.

Notice that $\Pr(A_u^i = 1) = \max\{(1 - 1.5(a - b) \frac{\log n}{n})^{n-1}, (1 - a \frac{\log n}{n})^{n-1}\}$ and therefore $\sum_{i,u} \mathbb{E} A_u^i \approx 4 \max\{n^{1-1.5(a-b)}, n^{1-a}\}$. This means that for $a - b \geq 0.67$ and $a \geq 1$, $\sum_{i,u} \mathbb{E} A_u^i = o(1)$. Hence there exist vertices in all the regions described above for every node u with high probability.

Now, A_u^1 and A_u^2 being zero implies that either there is a vertex in $[b \frac{\log n}{n}, a \frac{\log n}{n}]$ or there exists two vertices v_1, v_2 in $[0, \frac{a-b}{2} \frac{\log n}{n}]$ and $[\frac{-a-b}{2} \frac{\log n}{n}, -b \frac{\log n}{n}]$ respectively (see, Figure 11). In the second case, u is connected to v_2 and v_2 is connected to v_1 . Therefore u has nodes on left (v_2) and right (v_1) and u is connected to both of them through one hop in the graph.

Similarly, A_u^3 and A_u^4 being zero implies that either there exists a vertex in $[-a \frac{\log n}{n}, -b \frac{\log n}{n}]$ or again u will have vertices on left and right and will be connected to them. So, when all the four $A_u^i, i = 1, 2, 3, 4$ are zero together:

- $A_u^1 = A_u^2 = 0$ implies there is a neighbor of u on either sides or there is a single node in $[b \frac{\log n}{n}, a \frac{\log n}{n}]$
- $A_u^3 = A_u^4 = 0$ implies there is a neighbor of u on either sides or there is a single node in $[-a \frac{\log n}{n}, -b \frac{\log n}{n}]$

This shows that when $A_u^1 = A_u^2 = 0$ and $A_u^3 = A_u^4 = 0$ guarantee a node on only one side of u , there are nodes in $[b \frac{\log n}{n}, a \frac{\log n}{n}]$ and $[-a \frac{\log n}{n}, -b \frac{\log n}{n}]$. But in that case u has direct neighbors on both its left and right. We can conclude that every vertex u is connected to a vertex v on its right and a vertex w on its left such that $d(u, v) \in [0, a \frac{\log n}{n}]$ and $d(u, w) \in [-a \frac{\log n}{n}, 0]$; therefore every vertex is part of a cycle that covers $[0, 1]$.

We can now extend this proof to the case when $a - b > 0.5$.

Let c be large number to be chosen specifically later. Consider a node u and assume that the position of u is 0. Now consider the four different regions $[-a \frac{\log n}{n}, -b \frac{\log n}{n}]$, $[-(a-b) \frac{\log n}{n}, 0]$, $[b \frac{\log n}{n}, a \frac{\log n}{n}]$





Figure 11: Representation of four different random variables for Lemma 12.


 Figure 12: Pictorial representation of $I_u^i, J_u^i, K_u^i, M_u^i$ and their connectivity as described in Lemma 12. The colored lines show the regions that are connected to each other.

and $[0, a - b \frac{\log n}{n}]$ around u each divided into $L \equiv 2^c$ patches (intervals) of size $\theta = \frac{a-b}{2^c}$ in the following way:

1. $I_u^i = [\frac{(-a+(i-1)\theta) \log n}{n}, \frac{(-a+i\theta) \log n}{n}]$
2. $J_u^i = [\frac{(-(a-b)+(i-1)\theta) \log n}{n}, \frac{(-(a-b)+i\theta) \log n}{n}]$
3. $K_u^i = [\frac{(b+(i-1)\theta) \log n}{n}, \frac{(b+i\theta) \log n}{n}]$
4. $M_u^i = [\frac{((i-1)\theta) \log n}{n}, \frac{i\theta \log n}{n}]$

where $i = 1, 2, 3, \dots, L$. Note that any vertex in $\cup I_u^i \cup K_u^i$ is connected to u . See, Figure 12 for a depiction.

Consider a $\{0, 1\}$ -indicator random variable X_u that is 1 if and only if there does not exist any node in a region formed by union of any $2L - 1$ patches amongst the ones described above. Notice that when $a < 2b$, the patches do not overlap and the total size of $2L - 1$ patches is $\frac{2^{c+1}-1}{2^c} \frac{(a-b) \log n}{n}$ and when $a \geq 2b$, the patches can overlap and the total size of the $2L - 1$ patches is going to be more than $\min\{\frac{2^{c+1}-1}{2^c} \frac{(a-b) \log n}{n}, \frac{a \log n}{n}\}$. Since there are $\binom{4L}{2L-1} \leq n^{\frac{4L}{\log n}}$ possible regions that consists of $2L - 1$ patches,

$$\begin{aligned} \sum_u \mathbb{E} X_u &\leq n \binom{4L}{2L-1} \left(1 - \min\left\{\frac{2^{c+1}-1}{2^c} \frac{(a-b) \log n}{n}, \frac{a \log n}{n}\right\}\right)^{n-1} \\ &\leq \max\{n^{1-\frac{2^{c+1}-1}{2^c} \frac{(a-b) \log n}{\log n}}, n^{1-a+\frac{4L}{\log n}}\}. \end{aligned}$$

At this point we can choose $c = c_n = o(\log n)$ such that $\lim_n c_n = \infty$. Hence when $a - b > \frac{1}{2}$ and $a > 1$, for every vertex u there exists at least one patch amongst every $2L - 1$ patches in $\cup I_u^i \cup J_u^j \cup K_u^k, i, j, k = 1, 2, \dots, L$ that contains a vertex.

Consider a collection of patches $\cup_i I_u^i \cup_j J_u^j, i, j = 1, 2, \dots, L$. We know that there exist two patches amongst these I_u^i s and K_u^j s that contain at least one vertices. If one of I_u^i s and one of K_u^j s contain two vertices, we found one neighbor of u on both left and right directions (see, Figure 12).

We consider the other case now. Without loss of generality assume that there are no vertex in all I_u^i s and there exist at least two patches in K_u^i s that contain at least one vertex each. Hence, there exists at least one of $\{K_u^i \mid i \in \{1, 2, \dots, L - 1\}\}$ that contains a vertex. Similarly, we can also conclude in this case that there exists at least one of $\{J_u^i \mid i \in \{2, 3, \dots, L\}\}$ which contain a node. Assume J_u^ϕ to be the left most patch in $\cup J_u^i \mid i \in \{1, 2, \dots, L\}$ that contains a vertex (see, Figure 12). From our previous observation, we can conclude that $\phi \geq 2$.

We can observe that any vertex in J_u^j is connected to the vertices in patches $K_u^k, \forall k < j$. This is because for two vertices $v \in J_u^j$ and $w \in K_u^k$, we have

$$\begin{aligned} d(v, w) &\geq \frac{(b + (k - 1)\theta) \log n}{n} - \frac{(-(a - b) + j\theta) \log n}{n} = \frac{(a + (k - j - 1)\theta) \log n}{n}; \\ d(v, w) &\leq \frac{(b + k\theta) \log n}{n} - \frac{(-(a - b) + (j - 1)\theta) \log n}{n} = \frac{(a + (k - j + 1)\theta) \log n}{n}. \end{aligned}$$

Consider a collection of $2L - 1$ patches $\{\cup I_u^i \cup J_u^j \cup K_u^k \mid i, j, k \in \{1, \dots, L\}, j > \phi, k \leq \phi - 1\}$ where $\phi \geq 2$. This is a collection of $2L - 1$ patches out of which one must have a vertex and since none of $\{J_u^j \mid j > \phi\}$ and I_u^i can contain a vertex, one of $\{K_u^k \mid k \leq \phi - 1\}$ must contain the vertex. Recall that the vertex in J_u^ϕ is connected to any node in K_u^k for any $k \leq \phi - 1$ and therefore u has a node to the right direction and left direction that are connected to u . Therefore every vertex is part of a cycle and each of the circles covers $[0, 1]$. $\blacksquare$

Lemma (Restatement of Lemma 13). *Set two real numbers $k \equiv \lceil b/(a - b) \rceil + 1$ and $\epsilon < \frac{1}{2k}$. In an $\text{RAG}_1^*(n, [\frac{b \log n}{n}, \frac{a \log n}{n}])$, $0 < b < a$, with probability $1 - o(1)$ there exists a vertex u_0 and k nodes $\{u_1, u_2, \dots, u_k\}$ to the right of u_0 such that $d(u_0, u_i) \in [\frac{(i(a-b)-2i\epsilon) \log n}{n}, \frac{(i(a-b)-(2i-1)\epsilon) \log n}{n}]$ and k nodes $\{v_1, v_2, \dots, v_k\}$ to the right of u_0 such that $d(u_0, v_i) \in [\frac{((i(a-b)+b-(2i-1)\epsilon) \log n)}{n}, \frac{(i(a-b)+b-(2i-2)\epsilon) \log n}{n}]$, for $i = 1, 2, \dots, k$. The arrangement of the vertices is shown in Figure 13.*

Proof [Proof of Lemma 13] Recall that we want to show that there exists a node u_0 and k nodes $\{u_1, u_2, \dots, u_k\}$ to the right of u_0 such that $d(u_0, u_i) \in [\frac{(i(a-b)-2i\epsilon) \log n}{n}, \frac{(i(a-b)-(2i-1)\epsilon) \log n}{n}]$ and exactly k nodes $\{v_1, v_2, \dots, v_k\}$ to the right of u_0 such that $d(u_0, v_i) \in [\frac{((i(a-b)+b-(2i-1)\epsilon) \log n)}{n}, \frac{(i(a-b)+b-(2i-2)\epsilon) \log n}{n}]$, for $i = 1, 2, \dots, k$ and ϵ is a constant less than $\frac{1}{2k}$ (see Figure 13 for a depiction). Let A_u be an indicator $\{0, 1\}$ -random variable for every node u which is 1 if u satisfies the above conditions and 0 otherwise. We will show $\sum_u A_u \geq 1$ with high probability.

We have,

$$\Pr(A_u = 1) = n(n-1) \dots (n-(2k-1)) \left(\frac{\epsilon \log n}{n} \right)^{2k} \left(1 - 2k\epsilon \frac{\log n}{n} \right)^{n-2k}$$




Figure 13: The location of u_i and v_i relative to u scaled by $\frac{\log n}{n}$ in Lemma 13. Edges stemming put of v_1, v_2, v_3 are shown as blue, red and violet respectively.

$$\begin{aligned}
 &= c_0 n^{-2k\epsilon} (\epsilon \log n)^{2k} \prod_{i=0}^{2k-1} (1 - i/n) \\
 &= c_1 n^{-2k\epsilon} (\epsilon \log n)^{2k}
 \end{aligned}$$

where c_0, c_1 are just absolute constants independent of n (recall k is a constant). Hence,

$$\sum_u \mathbb{E} A_u = c_1 n^{1-2k\epsilon} (\epsilon \log n)^{2k} \geq 1$$

as long as $\epsilon \leq \frac{1}{2k}$. Now, in order to prove $\sum_u A_u \geq 1$ with high probability, we will show that the variance of $\sum_u A_u$ is bounded from above. This calculation is very similar to the one in the proof of the connectivity lower bound in Theorem 2. Recall that if $A = \sum_u A_u$ is a sum of indicator random variables, we must have

$$\text{Var}(A) \leq \mathbb{E}[A] + \sum_{u \neq v} \text{Cov}(A_u, A_v) = \mathbb{E}[A] + \sum_{u \neq v} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1).$$

Now first consider the case when vertices u and v are at a distance of at least $\frac{2(a+b) \log n}{n}$ apart (happens with probability $1 - \frac{4(a+b) \log n}{n}$). Then the region in $[0, 1]$ that is within distance $\frac{(a+b) \log n}{n}$ from both u and v is the empty-set. In this case, $\Pr(A_u = 1 \cap A_v = 1) = n(n-1) \dots (n - (4k-1)) \left(\frac{\epsilon \log n}{n}\right)^{4k} \left(1 - 4k\epsilon \frac{\log n}{n}\right)^{n-4k} = c_2 n^{-4k\epsilon} (\epsilon \log n)^{4k}$, where c_2 is a constant.

In all other cases, $\Pr(A_u = 1 \cap A_v = 1) \leq \Pr(A_u = 1)$. Therefore,

$$\Pr(A_u = 1 \cap A_v = 1) \leq \left(1 - \frac{4(a+b) \log n}{n}\right) c_2 n^{-4k\epsilon} (\epsilon \log n)^{4k} + \frac{4(a+b) \log n}{n} c_1 n^{-2k\epsilon} (\epsilon \log n)^{2k}$$

and

$$\begin{aligned}
 \text{Var}(A) &\leq c_1 n^{1-2k\epsilon} (\epsilon \log n)^{2k} + \binom{n}{2} \left(\Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1) \right) \\
 &\leq c_1 n^{1-2k\epsilon} (\epsilon \log n)^{2k} + c_3 n^{1-2k\epsilon} (\log n)^{2k+1}
 \end{aligned}$$




Figure 14: The line segments where v_1, v_2, v_3 can have neighbors (scaled by $\frac{\log n}{n}$) in the proof of Lemma 13. The point t has to lie in one of these regions.

$$\leq c_4 n^{1-2k\epsilon} (\log n)^{2k+1}$$

where c_3, c_4 are constants. Again invoking Chebyshev's inequality, with probability at least $1 - \frac{1}{\log n}$

$$A > c_1 n^{1-2k\epsilon} (\epsilon \log n)^{2k} - \sqrt{c_4 n^{1-2k\epsilon} (\log n)^{2k+2}}.$$

■

Proof [Proof of Corollary 15] Consider a node u and assume that the position of u is 0. Associate a random variable A_u^i for $i \in \{1, 2, 3, 4\}$ which takes the value of 1 when there does not exist any node x such that




Figure 15: The representation of different intervals corresponding to each random variable as described in Corollary 15

1. $d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [0, c \frac{\log n}{n}] \cup [-c \frac{\log n}{n}, -c/2 \frac{\log n}{n}]$ for $i = 1$
2. $d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [0, c \frac{\log n}{n}] \cup [b - c/2 \frac{\log n}{n}, (a - c) \frac{\log n}{n}]$ for $i = 2$
3. $d(u, x) \in [-a \frac{\log n}{n}, -b \frac{\log n}{n}] \cup [-c \frac{\log n}{n}, 0] \cup [c/2 \frac{\log n}{n}, c \frac{\log n}{n}]$ for $i = 3$
4. $d(u, x) \in [-a \frac{\log n}{n}, -b \frac{\log n}{n}] \cup [-c \frac{\log n}{n}, 0] \cup [(c - a) \frac{\log n}{n}, (c/2 - b) \frac{\log n}{n}]$ for $i = 4$

$$\Pr(A_u^i = 1) = \begin{cases} \left(1 - (c + a - b + (a - b - c/2)) \frac{\log n}{n}\right)^n & \text{when } a - c < b \text{ and } b - c/2 > c \\ \left(1 - (b - c) \frac{\log n}{n}\right)^n & \text{when } a - c < b \text{ and } b - c/2 < c \\ \left(1 - (a) \frac{\log n}{n}\right)^n & \text{when } a - c \geq b \text{ and } b - c/2 < c \\ \left(1 - (c + a - b + c/2) \frac{\log n}{n}\right)^n & \text{when } a - c \geq b \text{ and } b - c/2 \geq c \end{cases}$$

Notice that A_u^1 and A_u^2 being zero implies that either there is a node in $\{x \mid d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [0, c \frac{\log n}{n}]\}$ or there exists nodes (v_1, v_2) in $\{x \mid d(u, x) \in [\frac{-c \log n}{n}, \frac{-c/2 \log n}{n}]\}$ and $\{x \mid d(u, x) \in [\frac{b-c/2 \log n}{n}, \frac{(a-c) \log n}{n}]\}$. In the second case, u is connected to v_1 and v_1 is connected to v_2 . Therefore u has nodes on left and right and u is connected to both of them although not directly. Similarly A_u^3 and A_u^4 being zero implies that there exist nodes in $\{x \mid d(u, x) \in [-a \frac{\log n}{n}, -b \frac{\log n}{n}] \cup [-c \frac{\log n}{n}, 0]\}$ or again u will have nodes on left and right and will be connected to them. So, when all the 4 events happen together, the only exceptional case is when there are nodes in $\{x \mid d(u, x) \in [b \frac{\log n}{n}, a \frac{\log n}{n}] \cup [0, c \frac{\log n}{n}]\}$ and $\{x \mid d(u, x) \in [-a \frac{\log n}{n}, b \frac{\log n}{n}] \cup [-c \frac{\log n}{n}, 0]\}$. But in that case u has direct neighbors on both its left and right. So, we can conclude that for every node u , there exists a node v such that $d(u, v) \in [0, \frac{a \log n}{n}]$ and a node w such that $d(u, w) \in [\frac{-a \log n}{n}, 0]$ such that u is connected to both v and w . This implies that every node u has neighbors on both its left and right and therefore every node is part of a cycle that covers $[0, 1]$. $\blacksquare$

Appendix B. Connectivity of High Dimensional Random Annulus Graphs: Detailed Proofs of Theorems 6 and 5

In this section we first prove an impossibility result on the connectivity of random annulus graphs in t dimensions by showing a sufficient condition of existence of isolated nodes. Next, we show that if the gap between r_1 and r_2 is large enough then the RAG is fully connected. We will start by introducing a few notations. Let us define the regions $B_t(u, r)$ and $B_t(u, [r_1, r_2])$ for the any $u \in S^t$ in the following way:

$$B_t(u, r) = \{x \in S^t \mid \|u - x\|_2 \leq r\}$$

$$B_t(u, [r_1, r_2]) = \{x \in S^t \mid r_1 \leq \|u - x\|_2 \leq r_2\}.$$

First, we calculate $|B_t(u, r)|$ and show that it is proportional to r^t .

Lemma 36. $|B_t(u, r)| = c_t r^t$ for $r = o(1)$ where $c_t = \frac{\pi^{t/2}}{\Gamma(\frac{t}{2}+1)}$.

Proof We use the following fact from (Larsen and Schmidt (2017); Li (2011)) for the proof. For a t -dimensional unit sphere, the hyperspherical cap of angular radius $\theta = \max_{x \in B_t(u, r)} \arccos \langle x, u \rangle$ has a surface area $C_t(\theta)$ given by

$$C_t(\theta) = \int_0^{\tan \theta} \frac{S_{t-1}(r)}{(1 + r^2)^2} dr$$

where $S_{t-1}(\theta) = \frac{t\pi^{t/2}}{\Gamma(\frac{t}{2}+1)}\theta^{t-1}$. Note that $C_t(\theta)$ is nothing but $|B_t(u, r)|$ where $\cos \theta = 1 - \frac{r^2}{2}$ and therefore $\tan \theta = \frac{r\sqrt{4-r^2}}{2-r^2} \approx r$ for small r . Now since $r = o(1)$ and $1 + r^2$ is an increasing function of r , we must have that

$$\int_0^r \frac{t\pi^{t/2}}{(1+o(1))\Gamma(\frac{t}{2}+1)}\theta^{t-1}d\theta < C_t(\theta) < \int_0^r \frac{t\pi^{t/2}}{\Gamma(\frac{t}{2}+1)}\theta^{t-1}d\theta$$

and therefore $C_t(\theta)$ can be expressed as $c_t r^t$ where c_t lies in $\left(\frac{\pi^{t/2}}{(1+o(1))\Gamma(\frac{t}{2}+1)}, \frac{\pi^{t/2}}{\Gamma(\frac{t}{2}+1)}\right)$. $\blacksquare$

B.1 Impossibility Result

The following theorem proves the impossibility result for the connectivity of a random annulus graph by proving a tight threshold for the presence of an isolated node with high probability.

Theorem. (Restatement of Theorem 6) For a random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ where $r_1 = b\left(\frac{\log n}{n}\right)^{\frac{1}{t}}$ and $r_2 = a\left(\frac{\log n}{n}\right)^{\frac{1}{t}}$, there exists isolated nodes with high probability if and only if $a^t - b^t < \frac{\sqrt{\pi}(t+1)\Gamma(\frac{t+2}{2})}{\Gamma(\frac{t+3}{2})}$.

Proof Consider the random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ in t dimensions. In this graph, a node u is isolated if there are no nodes v such that $r_1 \leq \|u - v\|_2 \leq r_2$. Since all nodes are uniformly and randomly distributed on S^t , the probability of a node v being connected to a node u is the volume of $B_t(u, [r_1, r_2])$. Define the indicator random variable $A_u \in \{0, 1\}$ which is 1 if and only if the node u is isolated. Also define the random variable $A = \sum_u A_u$ which denotes the total number of isolated nodes. Since $|B_t(u, [r_1, r_2])| = c_t(r_2^t - r_1^t)$, we must have

$$\Pr(A_u = 1) = \left(1 - \frac{c_t(r_2^t - r_1^t)}{|S^t|}\right)^{n-1}.$$

Now, we know from (Coxeter (1973)) that $|S^t| = \frac{(t+1)\pi^{\frac{t+1}{2}}}{\Gamma(\frac{t+3}{2})}$. Plugging in, we get that $\frac{c_t}{|S^t|} = \frac{\Gamma(\frac{t+3}{2})}{\sqrt{\pi}(t+1)\Gamma(\frac{t+2}{2})}$. Hence, the expected number of isolated nodes $\mathbb{E}A$ is going to be

$$n \left(1 - (a^t - b^t) \frac{c_t}{|S^t|} \frac{\log n}{n}\right)^{n-1} \approx n^{1 - \frac{c_t(a^t - b^t)}{|S^t|}}.$$

Therefore $\mathbb{E}[A] \geq 1$ if $a^t - b^t < \frac{|S^t|}{c_t}$. In order to show that $A = \omega(1)$ with high probability we are going to show that the variance of A is bounded from above. Since A is a sum of indicator random variables, we have that

$$\text{Var}(A) \leq \mathbb{E}[A] + \sum_{u \neq v} \text{Cov}(A_u, A_v) = \mathbb{E}[A] + \sum_{u \neq v} (\Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1)).$$

Now, consider the scenario when the vertices u and v are at a distance more than $2r_2$ apart which happens with probability $1 - \frac{c_t(2r_2)^t}{|S^t|}$. Then the region in which every point is within a distance of r_2

and r_1 from both u, v is empty and therefore $\Pr(A_u = 1 \cap A_v = 1) = \left(1 - \frac{2c_t}{|S^t|} \left(a^t - b^t\right) \frac{\log n}{n}\right)^{n-2}$.

When the vertices are within distance $2r_2$ of one another, then $\Pr(A_u = 1 \cap A_v = 1) \leq \Pr(A_u = 1)$. Therefore,

$$\begin{aligned} \Pr(A_u = 1 \cap A_v = 1) &\leq \left(1 - \frac{c_t(2r_2)^t}{|S^t|}\right) \left(1 - \frac{2c_t}{|S^t|} \left(a^t - b^t\right) \frac{\log n}{n}\right)^{n-2} + \frac{c_t(2r_2)^t}{|S^t|} \Pr(A_u = 1) \\ &\leq \left(1 - \frac{c_t(2r_2)^t}{|S^t|}\right) n^{-\frac{2c_t}{|S^t|}(a^t - b^t) + o(1)} + \frac{c_t(2r_2)^t}{|S^t|} n^{-\frac{c_t(a^t - b^t)}{|S^t|} + o(1)}. \end{aligned}$$

Consequently for large enough n ,

$$\begin{aligned} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1) &\leq \left(1 - \frac{c_t(2r_2)^t}{|S^t|}\right) n^{-\frac{2c_t(a^t - b^t)}{|S^t|} + o(1)} \\ &\quad + \frac{c_t(2r_2)^t}{|S^t|} n^{-\frac{c_t(a^t - b^t)}{|S^t|} + o(1)} - n^{-\frac{2c_t(a^t - b^t)}{|S^t|} + o(1)} \leq \frac{2c_t(2r_2)^t}{|S^t|} \Pr(A_u = 1). \end{aligned}$$

Now,

$$\text{Var}(A) \leq \mathbb{E}[A] + 2 \binom{n}{2} \frac{c_t(2r_2)^t}{|S^t|} \Pr(A_u = 1) \leq \mathbb{E}[A] \left(1 + \frac{c_t(2a)^t}{|S^t|} \log n\right).$$

By using Chebyshev bound, with probability at least $1 - O\left(\frac{1}{\log n}\right)$,

$$A > n^{1 - \frac{c_t(a^t - b^t)}{|S^t|}} - \sqrt{n^{1 - \frac{c_t(a^t - b^t)}{|S^t|}} \left(1 + \frac{c_t(2a)^t}{|S^t|} \log n\right) \log n},$$

which implies that for $a^t - b^t < \frac{|S^t|}{c_t}$, $A > 1$ and hence there will exist isolated nodes with high probability. $\blacksquare$

B.2 Connectivity Bound

We show the upper bound for connectivity of a Random Annulus Graphs in t as per Theorem 5, rewritten below.

Theorem. (Restatement of Theorem 5) For t dimensional random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ where $r_2 = a\left(\frac{\log n}{n}\right)^t$ and $r_1 = b\left(\frac{\log n}{n}\right)^t$ with $a \geq b$ and t is a constant, the graph is connected with high probability if

$$a^t - b^t \geq \frac{8|S^t|(t+1)}{c_t\left(1 - \frac{1}{2^{1+1/t-1}}\right)} \text{ and } a > 2^{1+\frac{1}{t}} b.$$

Let us define a *pole* to be a vertex which is connected to all vertices within a distance of r_2 from itself. In order to prove Theorem 5, we first show the existence of a pole with high probability in Lemma 16. Next, Lemma 17 shows that for every vertex u and every hyperplane L passing through u and not too close to the tangential hyperplane at u , there will be a neighbor of u on either side of the plane. In order to formalize this, let us define a few regions associated with a node u and a hyperplane $L : w^T x = \beta$ passing through u .

$$\begin{aligned}\mathcal{R}_L^1 &\equiv \{x \in S^t \mid b\left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \leq \beta\} \\ \mathcal{R}_L^2 &\equiv \{x \in S^t \mid b\left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\} \\ \mathcal{A}_L &\equiv \{x \mid x \in S^t, w^T x = \beta\}.\end{aligned}$$

Informally, $\mathcal{R}_L^1$ and $\mathcal{R}_L^2$ represents the partition of the region $B_t(u, [r_1, r_2])$ on either side of the hyperplane L and $\mathcal{A}_L$ represents the region on the sphere lying on L .

Lemma. (Restatement of Lemma 16) In $\text{RAG}_t\left(n, \left[b\left(\frac{\log n}{n}\right)^{1/t}, a\left(\frac{\log n}{n}\right)^{1/t}\right]\right)$, $0 < b < a$, with probability $1 - o(1)$ there exists a pole.

The proof of Lemma 16 is delegated to Appendix B.3.

Lemma. (Restatement of Lemma 17) If we sample n nodes from S^t according to $\text{RAG}_t\left(n, \left[b\left(\frac{\log n}{n}\right)^{1/t}, a\left(\frac{\log n}{n}\right)^{1/t}\right]\right)$, then for every node u and every hyperplane L passing through u such that $\mathcal{A}_L \not\subset B_t(u, a\left(\frac{\log n}{n}\right)^{1/t})$, node u has a neighbor on both sides of the hyperplane L with probability at least $1 - \frac{1}{n}$ provided

$$a^t - b^t \geq \frac{8|S^t|(t+1)}{c_t\left(1 - \frac{1}{2^{1+1/t-1}}\right)}$$

and $a > 2^{1+\frac{1}{t}}b$.

For a node $u \equiv (u_1, u_2, \dots, u_{t+1})$, define the particular hyperplane $L_u^* : x_1 = u_1$ which is normal to the line joining $u_0 \equiv (1, 0, \dots, 0)$ and the origin and passes through u . We now have the following lemma.

Lemma. (Restatement of Lemma 18) For a particular node u and corresponding hyperplane L_u^* , if $\mathcal{A}_{L_u^*} \subseteq B_t(u, r_2)$ then u must be within r_2 of u_0 .

For now, we assume that the Lemmas 16, 17 and 18 are true and show why these lemmas together imply the proof of Theorem 5.

Proof [Proof of Theorem 5] We consider an alternate (rotated but not shifted) coordinate system by multiplying every vector by a orthonormal matrix R such that the new position of the pole is the $t+1$ -dimensional vector $(1, 0, \dots, 0)$ where only the first co-ordinate is non-zero. Let the $t+1$ dimensional vector describing a node u in this new coordinate system be $\hat{u}$. Now consider the

hyperplane $L : x_1 = \hat{u}_1$ and if u is not connected to the pole already, then by Lemma 17 and Lemma 18 the node u has a neighbor u_2 which has a higher first coordinate. The same analysis applies for u_2 and hence we have a path where the first coordinate of every node is higher than the previous node. Since the number of nodes is finite, this path cannot go on indefinitely and at some point, one of the nodes is going to be within r_2 of the pole and will be connected to the pole. Therefore every node is going to be connected to the pole and hence our theorem is proved. $\blacksquare$

We show the proofs of Lemma 16, 17 and 18 in the following sections.

B.3 Proof of Lemma 16

Lemma 37 is a helper lemma that shows the region of connectivity for a small ball of radius $\epsilon \left(\frac{\log n}{n}\right)^{1/t}$. Lemma 38 uses this lemma to show the existence of a point u_0 which is connected to various balls of radius $\epsilon \left(\frac{\log n}{n}\right)^{1/t}$.

Lemma 37. *For a t dimensional random annulus graph $\text{RAG}_t(n, [r_1, r_2])$ where $r_1 = b \left(\frac{\log n}{n}\right)^{1/t}$, $r_2 = a \left(\frac{\log n}{n}\right)^{1/t}$ and $a \geq b$, consider the region $B_t(O, \theta)$ centered at O and radius $\theta = \epsilon \left(\frac{\log n}{n}\right)^{1/t}$. Then, every vertex in $B_t(O, \theta)$ is connected to all vertices present in $B_t(O, [\theta_1, \theta_2])$ where $\theta_1 = (b + \epsilon) \left(\frac{\log n}{n}\right)^{1/t}$ and $\theta_2 = (a - \epsilon) \left(\frac{\log n}{n}\right)^{1/t}$.*

Proof For any point $A \in B_t(O, \theta)$, we have $0 < \|A - O\|_2 \leq \theta$ and for any point $X \in B_t(O, [\theta_1, \theta_2])$, we must have $\theta_1 \leq \|X - O\|_2 \leq \theta_2$. Hence

$$\begin{aligned} \|A - X\|_2 &\leq \|A - O\|_2 + \|X - O\|_2 \\ &\leq \theta + \theta_2 \\ &= a \left(\frac{\log n}{n}\right)^{1/t}, \\ \|A - X\|_2 &\geq \|X - O\|_2 - \|A - O\|_2 \\ &\geq \theta_1 - \theta \\ &= b \left(\frac{\log n}{n}\right)^{1/t}, \end{aligned}$$

and therefore the claim of the lemma is proved. $\blacksquare$

Lemma 38. *Set two real numbers $k \equiv \lceil b/(a - b) \rceil + 1$ and $\epsilon < \left(\frac{|S_t|}{2kct}\right)^{1/t}$. In an $\text{RAG}_t\left(n, \left[b \left(\frac{\log n}{n}\right)^{1/t}, a \left(\frac{\log n}{n}\right)^{1/t}\right]\right)$, $0 < b < a$, with probability $1 - o(1)$ there exists a vertex $u_0 \in S^t$ with the following property. Consider a homogeneous hyperplane L in $\mathbb{R}^{t+1}$ that pass through u_0 . There are k nodes $\mathcal{A} = \{u_1, u_2, \dots, u_k\}$ with $u_i \in B_t\left(O_{u_i}, \epsilon \left(\frac{\log n}{n}\right)^{1/t}\right)$ for some $O_{u_i} \in L \cap S^t$ such that $\|O_{u_i} - u_0\|_2 = (i(a - b) - (4i - 1)\epsilon) \left(\frac{\log n}{n}\right)^{1/t}$ and k*




Figure 16: Any node in the red region is connected to any node in the blue region.

nodes $\mathcal{B} = \{v_1, v_2, \dots, v_k\}$ with $v_i \in B_t \left(O_{v_i}, \epsilon \left(\frac{\log n}{n} \right)^{1/t} \right)$ for some $O_{v_i} \in L \cap S^t$ such that $\|O_{v_i} - u_0\|_2 = (i(a - b) + b - (4i - 3)\epsilon) \left(\frac{\log n}{n} \right)^{1/t}$, for $i = 1, 2, \dots, k$ with $\mathcal{A}$ and $\mathcal{B}$ separated by L .

Proof [Proof of Lemma 38] Let A_u be an indicator $\{0, 1\}$ -random variable for every node u which is 1 if u satisfies the conditions stated in the lemma and 0 otherwise. We will show $\sum_u A_u \geq 1$ with high probability.

We have,

$$\begin{aligned} \Pr(A_u = 1) &= \frac{1}{2k!} n(n-1) \dots (n - (2k-1)) \left(\frac{\epsilon^t c_t \log n}{n |S^t|} \right)^{2k} \left(1 - 2k c_t \epsilon^t \frac{\log n}{n |S^t|} \right)^{n-2k} \\ &= c_1 n^{-2k \epsilon^t c_t / |S^t|} (\epsilon^t \log n)^{2k} \prod_{i=0}^{2k-1} (1 - i/n) \\ &= c_2 n^{-2k \epsilon^t c_t / |S^t|} (\epsilon^t \log n)^{2k} \end{aligned}$$

where $c_1 = \frac{c_t^{2k}}{2k! |S^t|^{2k}}$, c_2 are just absolute constants independent of n (recall k is a constant). Hence,

$$\sum_u \mathbb{E} A_u = c_2 n^{1-2k \epsilon^t c_t / |S^t|} (\epsilon^t \log n)^{2k} \geq 1$$

as long as $\epsilon \leq \left(\frac{|S^t|}{2k c_t} \right)^{1/t}$. Now, in order to prove $\sum_u A_u \geq 1$ with high probability, we will show that the variance of $\sum_u A_u$ is bounded from above. Recall that if $A = \sum_u A_u$ is a sum of indicator random variables, we must have

$$\text{Var}(A) \leq \mathbb{E}[A] + \sum_{u \neq v} \text{Cov}(A_u, A_v) = \mathbb{E}[A] + \sum_{u \neq v} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1).$$

Now first consider the case when vertices u and v are at a distance of at least $2(a + b) \left(\frac{\log n}{n} \right)^{1/t}$ apart (happens with probability $1 - 4^t (a + b)^t c_t \left(\frac{\log n}{n |S^t|} \right)$). Then the region that is within distance $(a + b) \left(\frac{\log n}{n} \right)^{1/t}$ from both u and v is the empty-set. In that case, $\Pr(A_u = 1 \cap A_v = 1) =$




Figure 17: Representation of u_i and v_i in the $t + 1$ -dimensional sphere with respect to u_0 .

$n(n-1)\dots(n-(4k-1))c_3\left(\frac{\epsilon^t c_t \log n}{n|S^t|}\right)^{4k}\left(1-4k\epsilon^t \frac{c_t \log n}{n|S^t|}\right)^{n-4k} = c_4 n^{-4k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{4k}$,
 where c_3, c_4 are constants.

In all other cases, $\Pr(A_u = 1 \cap A_v = 1) \leq \Pr(A_u = 1)$. Therefore,

$$\Pr(A_u = 1 \cap A_v = 1) \leq \left(1 - 4^t(a+b)^t c_t \left(\frac{\log n}{n|S^t|}\right)\right) c_4 n^{-4k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{4k} + \frac{4^t(a+b)^t c_t \log n}{n|S^t|} c_2 n^{-2k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{2k}$$

and

$$\begin{aligned} \text{Var}(A) &\leq c_2 n^{1-2k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{2k} + \binom{n}{2} \left(\Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1) \right) \\ &\leq c_2 n^{1-2k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{2k} + c_5 n^{1-2k\epsilon^t c_t/|S^t|}(\log n)^{2k+1} \\ &\leq c_6 n^{1-2k\epsilon^t c_t/|S^t|}(\log n)^{2k+1} \end{aligned}$$

where c_5, c_6 are constants. Again invoking Chebyshev's inequality, with probability at least $1 - O\left(\frac{1}{\log n}\right)$

$$A > c_2 n^{1-2k\epsilon^t c_t/|S^t|}(\epsilon^t \log n)^{2k} - \sqrt{c_6 n^{1-2k\epsilon^t c_t/|S^t|}(\log n)^{2k+2}}$$

which implies that $A > 1$ with high probability. $\blacksquare$

Lemma. (Restatement of Lemma 16) In a $\text{RAG}_t\left(n, \left[b\left(\frac{\log n}{n}\right)^{1/t}, a\left(\frac{\log n}{n}\right)^{1/t}\right]\right)$, $0 < b < a$, with probability $1 - o(1)$ there exists a vertex u_0 such that any node v that satisfies $\|u - v\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}$ is connected to u_0 .

Proof Consider the vertices $u_0, \{u_1, u_2, \dots, u_k\}$ and $\{v_1, v_2, \dots, v_k\}$ that satisfy the conditions of Lemma 38 as shown in Fig 17. We can observe that each vertex v_i has an edge with u_i and u_{i-1} , $i = 1, \dots, k$.

$$\|u_i - v_i\|_2 \geq \|u_i - O_{v_i}\|_2 - \|v_i - O_{v_i}\|_2$$




Figure 18: Shaded regions represent the region of connectivity with u_i and v_i (red for u_i 's and gray for v_i 's).

$$\begin{aligned} &\geq \|O_{v_i} - O_{u_i}\|_2 - \|u_i - O_{u_i}\|_2 - \|v_i - O_{v_i}\|_2 \\ &\geq (b + 2\epsilon) \left(\frac{\log n}{n}\right)^{1/t} - 2\epsilon \left(\frac{\log n}{n}\right)^{1/t} = b \left(\frac{\log n}{n}\right)^{1/t} \quad \text{and} \end{aligned}$$

$$\begin{aligned} \|u_i - v_i\|_2 &\leq \|O_{v_i} - O_{u_i}\|_2 + \|u_i - O_{u_i}\|_2 + \|v_i - O_{v_i}\|_2 \\ &= (b + 2\epsilon) \left(\frac{\log n}{n}\right)^{1/t} + 2\epsilon \left(\frac{\log n}{n}\right)^{1/t} = (b + 4\epsilon) \left(\frac{\log n}{n}\right)^{1/t} \end{aligned}$$

Similarly,

$$\begin{aligned} \|u_{i-1} - v_i\|_2 &\geq \|u_{i-1} - O_{v_i}\|_2 - \|v_i - O_{v_i}\|_2 \\ &\geq \|O_{v_i} - O_{u_{i-1}}\|_2 - \|u_{i-1} - O_{u_{i-1}}\|_2 - \|v_i - O_{v_i}\|_2 \\ &\geq (a - 2\epsilon) \left(\frac{\log n}{n}\right)^{1/t} - 2\epsilon \left(\frac{\log n}{n}\right)^{1/t} \\ &= (a - 4\epsilon) \left(\frac{\log n}{n}\right)^{1/t} \quad \text{and} \end{aligned}$$

$$\begin{aligned} \|u_{i-1} - v_i\|_2 &\leq \|O_{v_i} - O_{u_{i-1}}\|_2 + \|u_{i-1} - O_{u_{i-1}}\|_2 + \|v_i - O_{v_i}\|_2 \\ &\leq (a - 2\epsilon) \left(\frac{\log n}{n}\right)^{1/t} + 2\epsilon \left(\frac{\log n}{n}\right)^{1/t} = a \left(\frac{\log n}{n}\right)^{1/t}. \end{aligned}$$

This implies that u_0 is connected to u_i and v_i for all $i = 1, \dots, k$. Next, we show that any point in the region $B_t(u_0, r_s = a(\frac{\log n}{n})^{1/t})$ is connected to u_0 . Now recall that any point in the region $B_t(x, \left[(b + \epsilon)\left(\frac{\log n}{n}\right)^t, (a - \epsilon)\left(\frac{\log n}{n}\right)^t\right])$ is connected to any point in the region $B_t(x, \epsilon)$.

We can observe that the nodes $u_1, \dots, u_k, v_1, \dots, v_k$ form a cover of $B_t(u_0, r_s = a(\frac{\log n}{n})^{1/t})$ in the form of these annulus regions (A region corresponding to a particular node implies the

portion of the hypersphere such that any other node in that region is connected to it) translated by $(a - b - 4\epsilon) \left(\frac{\log n}{n}\right)^{1/t}$. This is because any node in $B_t\left(O_{u_i}, \left[(b + \epsilon) \left(\frac{\log n}{n}\right)^t, (a - \epsilon) \left(\frac{\log n}{n}\right)^t\right]\right)$ is connected to u_i and any node in $B_t\left(O_{v_i}, \left[(b + \epsilon) \left(\frac{\log n}{n}\right)^t, (a - \epsilon) \left(\frac{\log n}{n}\right)^t\right]\right)$ is connected to v_i respectively (Figure 18). Therefore any node falling in any of the aforementioned regions is connected with u_0 . Since the width of each region is $(a - b - 2\epsilon) \left(\frac{\log n}{n}\right)^{1/t}$, the regions overlap with each other. Additionally, the inner radius of a particular region is $(b + \epsilon) \left(\frac{\log n}{n}\right)^{1/t}$ which is greater than $b \left(\frac{\log n}{n}\right)^{1/t}$. Hence, there can not exist any point in $B_t\left(u_0, r_2 = a \left(\frac{\log n}{n}\right)^{1/t}\right)$ which is not covered by the union of these regions. $\blacksquare$

B.4 Proofs of Lemma 17 and Lemma 18

Assume that the $(t+1)$ -dimensional space is described by a coordinate system whose center coincides with the center of the sphere. In this coordinate system, let us denote the point $(1, 0, 0, \dots, 0)$ by u_0 . Lemma 39 shows that for any plane L with $\mathcal{A}_L \not\subset B_t(u, r_2)$, the region of connectivity of u on both sides differ by a constant fraction.

Lemma 39. *For a particular node u in $\text{RAG}_t(n, [r_1, r_2])$ where $r_1 = b \left(\frac{\log n}{n}\right)^{1/t}$, $r_2 = a \left(\frac{\log n}{n}\right)^{1/t}$, consider a hyperplane L passing through u such that $\mathcal{A}_L \not\subset B_t(u, a \left(\frac{\log n}{n}\right)^{1/t})$, then $\frac{\min(|\mathcal{R}_L^1|, |\mathcal{R}_L^2|)}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} \geq \delta$ if $a > 2b$, where δ is a constant.*

Proof

First, for a node u and a given hyperplane $L : w^T x = \beta$ passing through u , we try to evaluate the surface area of the region corresponding to

$$\{x \in S^t \mid \|u - x\|_2 \leq r_2 = a \left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}$$

such that the farthest point from u on the plane L and S^t is at distance r_2 . This region is a spherical cap corresponding to $B_t(u', r')$ where u' is the intersection of S^t with the normal from the origin to the plane and $r' = \|u - u'\|_2$. Suppose h is the height of this cap (perpendicular distance from u' to the hyperplane L). Using pythagoras theorem, we can see that $r'^2 = h^2 + (r_2/2)^2$ and $(1 - h)^2 + r_2^2/4 = 1$. Simplifying this, we get $h = \frac{r_2^2}{2}$ and hence $r' \approx \frac{r_2}{2}$. Hence the area of this region is $c_t (r_2/2)^t$.

Without loss of generality, assume $|\mathcal{R}_L^1| \geq |\mathcal{R}_L^2|$. Now,

$$\begin{aligned} |\mathcal{R}_L^1| &= |\{x \in S^t \mid b \left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a \left(\frac{\log n}{n}\right)^{1/t}, w^T x \leq \beta\}| \\ &= |\{x \in S^t \mid b \left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a \left(\frac{\log n}{n}\right)^{1/t}\}| - |\mathcal{R}_L^2| \\ &= c_t(r_2^t - r_1^t) - |\mathcal{R}_L^2|. \end{aligned}$$

$$\begin{aligned}
 |\mathcal{R}_L^2| &\equiv |\{x \in S^t \mid b\left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| \\
 &\equiv |\{x \in S^t \mid \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| - |\{x \in S^t \mid \|u - x\|_2 \leq b\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| \\
 &\geq |\{x \in S^t \mid \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| - \\
 &\quad \left[|\{x \in S^t \mid \|u - x\|_2 \leq b\left(\frac{\log n}{n}\right)^{1/t}, w^T x < \beta\}| + |\{x \in S^t \mid \|u - x\|_2 \leq b\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| \right] \\
 &\equiv |\{x \in S^t \mid \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta\}| - |\{x \in S^t \mid \|u - x\|_2 \leq b\left(\frac{\log n}{n}\right)^{1/t}\}| \\
 &\equiv c_t (r_2/2)^t - c_t r_1^t
 \end{aligned}$$

If $c_t (r_2/2)^t - c_t r_1^t > 0$, then,

$$\begin{aligned}
 1 \leq \frac{|\mathcal{R}_L^1|}{|\mathcal{R}_L^2|} &\leq \frac{c_t (r_2^t - r_1^t)}{c_t (r_2/2)^t - c_t r_1^t} - 1 \\
 &= \frac{a^t - b^t}{(a/2)^t - b^t} - 1 = \delta'
 \end{aligned}$$

Hence,

$$2 \leq \frac{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|}{|\mathcal{R}_L^2|} \leq 1 + \delta'$$

This gives us that $\frac{\min(|\mathcal{R}_L^1|, |\mathcal{R}_L^2|)}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} = \frac{|\mathcal{R}_L^2|}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} \geq \frac{1}{1 + \delta'} = \delta$. Hence, the claim of the lemma is satisfied if $(a/2)^t - b^t > 0$ i.e. $a > 2b$. $\blacksquare$

Corollary 40. For a particular node u in $\text{RAG}_t\left(n, \left[b\left(\frac{\log n}{n}\right)^{1/t}, a\left(\frac{\log n}{n}\right)^{1/t}\right]\right)$, consider a hyperplane L passing through u such that $\mathcal{A}_L \not\subseteq B_t(u, a\left(\frac{\log n}{n}\right)^{1/t})$, then $\frac{\min(|\mathcal{R}_L^1|, |\mathcal{R}_L^2|)}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} \geq \frac{(a/2)^t - b^t}{a^t - b^t}$.

Proof Using Lemma 39, $\frac{\min(|\mathcal{R}_L^1|, |\mathcal{R}_L^2|)}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} \geq \delta = \frac{1}{1 + \delta'}$ where $\delta' = \frac{(a)^t - b^t}{(a/2)^t - b^t} - 1$. $\blacksquare$

For a node u , recall that the hyperplane $L_u^* : x_1 = u_1$ is normal to the line joining u_0 and the origin and passes through u . We now have the following lemma, which tries to show that if the plane satisfies $\mathcal{A}_{L_u^*} \subseteq B_t(u, r_2)$ then the node u must be within r_2 distance of u_0 .

Lemma. (Restatement of Lemma 18) For a particular node u and corresponding hyperplane L_u^* , if $\mathcal{A}_{L_u^*} \subseteq B_t(u, r_2)$ then u must be within r_2 of u_0 .

Proof The reflection of u in x-axis (say v) is the farthest from u that lies on both $\mathcal{A}_{L_u^*}$ and S^t . Now we want to show that $v = (u_1, -u_2, \dots, -u_{t+1})$ has the following property: if $\|u - v\|_2 \leq r_2$ then $\|u - u_0\|_2 \leq r_2$. We are given that

$$u_1^2 + u_2^2 + \dots + u_{t+1}^2 = 1$$

$$d(u, v) = \sqrt{4(u_2^2 + \dots + u_{t+1}^2)} \leq r_2$$

$$4(1 - u_1^2) \leq r_2^2$$

We need to show that,

$$\begin{aligned} d(u, u_0)^2 &= (1 - u_1)^2 + (u_2^2 + \dots + u_{t+1}^2) \\ &= (1 - u_1)^2 + (1 - u_1^2) \\ &= 2 - 2u_1 \\ &\leq r_2^2 \end{aligned}$$

which holds if $\sqrt{\frac{4-r_2^2}{4}} \geq \frac{2-r_2^2}{2}$. Notice that,

$$\begin{aligned} \sqrt{\frac{4-r_2^2}{4}} &\geq \frac{2-r_2^2}{2} \\ \implies 4 - r_2^2 &\geq 4 - 4r_2^2 + r_2^4 \\ \implies r_2^4 - 3r_2^2 &\leq 0 \\ \implies r_2^2(r_2^2 - 3) &\leq 0 \end{aligned}$$

which is true since $0 \leq r_2 \leq 1$. ■

Since, we do not know the location of the pole, we need to show that every point has a neighbor on both sides of the plane L no matter what the orientation of the plane given that $\mathcal{A}_L \not\subset B_t(u, r_2)$. For this we need to introduce the concept of VC Dimension. Define (X, R) to be a range space if X is a set (possibly infinite) and R is a family of subsets of X . For any set $A \subseteq X$, we define $P_R(A) = \{r \cap A \mid r \in R\}$ to be the projection of R on A . Finally we define the VC dimension d of a range space (X, R) to be $d = \sup_{A \subseteq X} \{|A| \mid |P_R(A)| = 2^{|A|}\}$. Next we give a modified version of a well-known theorem about VC-dimension (Haussler and Welzl (1987)).

Lemma 41. *Let (X, R) be a range space of VC dimension d and let U be a uniform probability measure defined on X . In that case, if we sample a set $\mathcal{M}$ of m points according to U such that*

$$m \geq \max \left(\frac{8d}{\epsilon} \log \frac{8d}{\epsilon}, \frac{4}{\epsilon} \log \frac{2}{\eta} \right)$$

then with probability $1 - \eta$ for any set $r \in R$ such that $\Pr_{x \sim U} (x \in r) \geq \epsilon$, we have $|r \cap \mathcal{M}| \neq \Phi$.

Proof Define a set $r \in R$ to be *heavy* if $\Pr_{x \sim U} (x \in r) \geq \epsilon$. We pick two random samples N and T each of size m according to the uniform distribution defined on X . Consider the event E_1 (bad event) for which there exists a heavy $r \in R$ such that $r \cap N = \Phi$. Consider another event E_2 for which there exists a heavy $r \in R$ such that $r \cap N = \Phi$ and $|r \cap T| \geq \frac{\epsilon m}{2}$. Now, since r is heavy, assume that $\Pr_{x \sim U} (x \in r) = \alpha$ such that $\alpha > \epsilon$. In that case, $|r \cap T|$ is a Binomial random variable with mean αm and variance at most αm as well. Hence, we have that

$$\begin{aligned} \Pr(E_2 \mid E_1) &= \Pr(|r \cap T| \geq \frac{\epsilon m}{2}) = 1 - \Pr(|r \cap T| \leq \frac{\epsilon m}{2}) \\ &\geq 1 - \Pr(|r \cap T| \leq \frac{\alpha m}{2}) \geq 1 - \frac{\alpha m}{(\frac{\alpha m}{2})^2} \geq 1 - \frac{4}{m\alpha} \end{aligned}$$

Now for $m \geq \frac{8}{\epsilon} \geq \frac{8}{\alpha}$, we conclude that $\Pr(E_2 \mid E_1) \geq \frac{1}{2}$. Now consider the same experiment in a different way. Consider picking $2m$ samples according to the uniform distribution from X and then equally partition them randomly between N and T . Consider the following event for a particular set r .

$$E_r : r \cap N = \Phi \text{ and } |r \cap T| \geq \frac{\epsilon m}{2}$$

and therefore

$$E_2 = \bigcup_{r: \text{heavy}} E_r$$

Let us fix $N \cup T$ and define $p = |r \cap (N \cup T)|$. In that case, we have

$$\Pr(r \cap N = \Phi \mid |r \cap (N \cup T)| \geq \frac{\epsilon m}{2}) = \frac{(2m-p)(2m-p-1) \dots (m-p+1)}{2m(2m-1) \dots (2m-p+1)} \leq 2^{-p} \leq 2^{-\frac{\epsilon m}{2}}$$

The last inequality holds since $p \geq \frac{\epsilon m}{2}$. Now, since the VC dimension of the range space (X, R) is d , the cardinality of the set $\{r \cap (N \cup T) \mid r \in R\}$ is at most $\sum_{i \leq d} \binom{2m}{i} \leq (2m)^d$ (see Shalev-Shwartz and Ben-David (2014)). Notice that

$$\Pr(E_r) = \Pr(r \cap N = \Phi \mid |r \cap (N \cup T)| \geq \frac{\epsilon m}{2}) \Pr(|r \cap (N \cup T)| \geq \frac{\epsilon m}{2}) \leq 2^{-\frac{\epsilon m}{2}}$$

Therefore by using the union bound over the possible number of distinct events E_r , we have

$$\Pr(E_2) \leq (2m)^d 2^{-\frac{\epsilon m}{2}}$$

Since $\Pr(E_2 \mid E_1) \geq \frac{1}{2}$ and $\Pr(E_1 \mid E_2) = 1$, we must have

$$\Pr(E_1) \leq 2(2m)^d 2^{-\frac{\epsilon m}{2}} \leq \delta$$

which is ensured by the statement of the theorem. ■

In order to use this theorem consider the range space $(X, \mathcal{R}_u)$ where X is the set of points in $\mathcal{S}^t$ and $\mathcal{R}_u$ be the family of sets $\{x \in \mathcal{S}^t \mid b\left(\frac{\log n}{n}\right)^{1/t} \leq \|u - x\|_2 \leq a\left(\frac{\log n}{n}\right)^{1/t}, w^T x \geq \beta, \mathcal{A}_{L:w^T x = \beta} \not\subset B_t(u, r_2)\}$. We now have the following lemma about the VC Dimension of the above range space which is a straightforward extension of VC dimension of half-spaces (Shalev-Shwartz and Ben-David (2014)):

Lemma 42. *VC dimension of the range space $(X, \mathcal{R}_u) \leq t + 1$.*

Proof In order to show this, consider a set $\mathcal{S}$ of $t + 2$ points. Recall that the convex hull of a set $\mathcal{S}$ of points $\{x_i\}_{i=1}^n$ is the set

$$C(\mathcal{S}) = \left\{ \sum \lambda_i x_i \mid \sum \lambda_i = 1, \lambda_i \geq 0 \right\}.$$

By Radon's lemma (Shalev-Shwartz and Ben-David (2014)) we have that the set of points $\mathcal{S}$ can be partitioned into two sets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that their convex hulls intersect. Let $p \in \mathcal{S}_1$ be a point in that intersection. Assume there exist a hyperplane such that

$$\begin{aligned} w^T x_i &\leq w_0, \forall x_i \in \mathcal{S}_1 \\ w^T x_i &\geq w_0, \forall x_i \in \mathcal{S}_2. \end{aligned}$$

Since p is in the convex hull of $\mathcal{S}_1$ we must have that $w^T p \leq w_0$. But then,

$$w^T p = \sum_{i: x_i \in \mathcal{S}_2} \lambda_i w^T x_i > \left(\sum_{i \in \mathcal{S}_2} \lambda_i \right) \min_{i: x_i \in \mathcal{S}_2} w^T x_i = \min(w^T x_i) > w_0.$$

which is a contradiction. Hence it is not possible to shatter $t + 2$ elements and therefore the VC dimension of this range space is at most $t + 1$. $\blacksquare$

Using the results shown above, we are ready to prove the following Lemma.

Lemma. (Restatement of Lemma 17) *If we sample n nodes from S^t according to $\text{RAG}_t(n, [r_1, r_2])$ with $r_1 = b \left(\frac{\log n}{n} \right)^{1/t}$, $r_2 = a \left(\frac{\log n}{n} \right)^{1/t}$, then for every node u and every hyperplane L passing through u such that $\mathcal{A}_L \not\subseteq B(u, r_2)$, node u has a neighbor on both sides of the hyperplane L with probability at least $1 - \frac{1}{n}$ provided*

$$(a/2)^t - b^t \geq \frac{8|S^t|(t+1)}{c_t}$$

and $a > 2b$.

Proof Recall that the volume of $B(u, r_1, r_2)$ is $c_t(r_2^t - r_1^t) = \frac{c_t \log n}{n}(a^t - b^t)$. According to Corollary 40, whenever $a > 2b$, $\frac{\min(|\mathcal{R}_L^1|, |\mathcal{R}_L^2|)}{|\mathcal{R}_L^1| + |\mathcal{R}_L^2|} \geq \delta$ where $\delta = \frac{(a/2)^t - b^t}{a^t - b^t}$ when the hyperplane L satisfies the conditions of the lemma. In that case, we have that for all $r \in R_u$,

$$\Pr_{x \sim_U X}(x \in r) \geq \frac{\delta c_t \log n}{n|S^t|}(a^t - b^t).$$

Since VC Dimension of $(X, R_u) \leq t + 1$ and n points are sampled from X , the conditions of Lemma 41 is satisfied for $\eta = \frac{2}{n^2}$ if

$$n \geq \max \left(\frac{8n|S^t|(t+1)}{\delta c_t \log n(a^t - b^t)} \log \frac{8n|S^t|(t+1)}{\delta c_t \log n(a^t - b^t)}, \frac{8n|S^t|}{\delta c_t \log n(a^t - b^t)} \log n \right)$$

Since $\lim_{n \rightarrow \infty} \frac{1}{\log n} \log \frac{8|S^t|n(t+1)}{\delta c_t \log n(a^t - b^t)} \rightarrow 1$ for constant t , hence we have that

$$a^t - b^t \geq \frac{8|S^t|(t+1)}{c_t \delta}.$$

By taking a union bound over all the n range spaces (X, R_u) corresponding to the n nodes and applying the statement of Lemma 41, we have proved the lemma. $\blacksquare$

Appendix C. Proof of Lemma 43

Lemma 43. *A random geometric graph $G(n, \frac{a \log n}{n})$ will have $\omega(1)$ disconnected components for $a < 1$.*

Proof Define an indicator random variable A_u for a node u which is 1 if it does not have a neighbor on its left. We must have that

$$\Pr(A_u) = \left(1 - \frac{a \log n}{n}\right)^{n-1}.$$

Therefore we must have that $\sum_u \mathbb{E}A_u = n^{1-a} = \Omega(1)$ if $a < 1$. This statement also holds true with high probability. To show this we need to prove that the variance of $\sum_u \mathbb{E}A_u$ is bounded. We have that

$$\text{Var}(A) < \mathbb{E}[A] + \sum_{u \neq v} \text{Cov}(A_u, A_v) = \mathbb{E}[A] + \sum_{u \neq v} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1)$$

Now, consider the scenario when the vertices u and v are at a distance more than $\frac{2a \log n}{n}$ apart (happens with probability at least $1 - \frac{4a \log n}{n}$). Then the region in $[0, 1]$ that is within distance $\frac{a \log n}{n}$ from both of the vertices is empty and therefore $\Pr(A_u = 1 \cap A_v = 1) = \Pr(A_u = 1) \Pr(A_v = 1) = \Pr(A_u = 1) \Pr(A_v = 1) \leq \Pr(A_u = 1) \Pr(A_v = 1) = (\Pr(A_u = 1))^2$. When the vertices are within distance $\frac{2a \log n}{n}$ of one another, then $\Pr(A_u = 1 \cap A_v = 1) \leq \Pr(A_u = 1)$. Therefore,

$$\Pr(A_u = 1 \cap A_v = 1) \leq \left(1 - \frac{4a \log n}{n}\right)(\Pr(A_u = 1))^2 + \frac{4a \log n}{n} \Pr(A_u = 1).$$

Consequently,

$$\begin{aligned} \Pr(A_u = 1 \cap A_v = 1) - \Pr(A_u = 1) \Pr(A_v = 1) &\leq \left(1 - \frac{4a \log n}{n}\right)(\Pr(A_u = 1))^2 \\ &\quad + \frac{4a \log n}{n} \Pr(A_u = 1) - (\Pr(A_u = 1))^2 \leq \frac{4a \log n}{n} \Pr(A_u = 1). \end{aligned}$$

Now,

$$\text{Var}(A) \leq \mathbb{E}[A] + \binom{n}{2} \frac{4a \log n}{n} \Pr(A_u = 1) \leq \mathbb{E}[A](1 + 2a \log n).$$

By using Chebyshev bound, with probability at least $1 - \frac{1}{\log n}$,

$$A > n^{1-a} - \sqrt{n^{1-a}(1 + 2a \log n) \log n},$$

Now, observe that if there exists k nodes which do not have a neighbor on one side, then there must exist $k - 1$ disconnected components. Hence the number of disconnected components in $G(n, \frac{a \log n}{n})$ is $\omega(1)$. ■