

On the estimation of latent distances using graph distances*

Ery Arias-Castro

Department of Mathematics, University of California, San Diego (USA)
e-mail: eariasca@ucsd.edu

and

Antoine Channarond

LMRS – UMR 6085 CNRS – Université de Rouen (France)

and

Bruno Pelletier

IRMAR – UMR CNRS 6625 – Université Rennes II (France)

and

Nicolas Verzelen

INRAE – UMR 729 MISTEA – Montpellier (France)

Abstract: We are given the adjacency matrix of a geometric graph and the task of recovering the latent positions. We study one of the most popular approaches which consists in using the graph distances and derive error bounds under various assumptions on the link function. In the simplest case where the link function is proportional to an indicator function, the bound matches an information lower bound that we derive.

Keywords and phrases: Latent positions, multidimensional scaling, graph embedding, graph distances, random geometric graphs.

Received August 2020.

Contents

1	Introduction	723
1.1	Related work	724
1.2	Our contribution	725
1.3	The graph distance method	726
2	Simple setting	727
3	General setting	730
4	Discussion	731

arXiv: [1804.10611](https://arxiv.org/abs/1804.10611)

*This work was partially supported by the US National Science Foundation (DMS 1513465, DMS 1916071), by the French Région Normandie (RIN 17B01101GR) and Agence Nationale de la Recherche (ANR-18-CE40-0014).

4.1	Localization	732
4.2	Regularization	733
4.3	Number of common neighbors	733
5	Proofs	734
5.1	Proof of Theorem 1	734
5.2	Proof of Theorem 2	734
5.3	Proof of Theorem 3	737
	References	745

1. Introduction

Suppose that we observe a undirected graph with adjacency matrix $W = (W_{ij} : i, j \in [n])$ (where $[n] := \{1, \dots, n\}$ and $n \geq 3$) with $W_{ij} \in \{0, 1\}$ and $W_{ii} = 0$. We assume the existence of points, $x_1, \dots, x_n \in \mathbb{R}^v$, such that

$$\mathbb{P}(W_{ij} = 1 \mid x_1, \dots, x_n) = \phi(\|x_i - x_j\|), \quad (1)$$

for some non-increasing link function $\phi : [0, \infty) \mapsto [0, 1]$. The $(W_{ij}, i < j)$ are assumed to be independent given the point set (x_1, \dots, x_n) . We place ourselves in a setting where the adjacency matrix W is observed, but the underlying points are unknown. We will be mostly interested in settings where ϕ is unknown (and no parametric form is known). Our most immediate interest is in the pairwise distances

$$d_{ij} := \|x_i - x_j\|. \quad (2)$$

In general, when the link function is unknown, all we can hope for is to rank these distances. Indeed, the most information we can aspire to extract from W is the probability matrix $P := (p_{ij})$, where

$$p_{ij} := \mathbb{P}(W_{ij} = 1 \mid x_1, \dots, x_n), \quad (3)$$

and even with perfect knowledge of P , the distances can only be known up to a monotone transformation, since $p_{ij} = \phi(d_{ij})$ and ϕ is in principle an arbitrary non-increasing function. Recovering the points based on such a ranking amounts to a problem of ordinal embedding (aka, non-metric multidimensional scaling), which has a long history [19, 36, 37, 45].

Although this is true in general, we focus our attention on the ‘local setting’ where the link function has very small support. In that particular case, we are able to (approximately) recover the pairwise distances up to a scaling. By fixing the scale arbitrarily (since it cannot be inferred from the available data), recovering the underlying points amounts to a problem of metric multidimensional scaling [7]. Classical Scaling [42] is the most popular method for that problem, and comes with a perturbation bound [5] which can help translate an error bound for the estimation of the pairwise distances (up to scale) to an error bound for the estimation of the points (up to a similarity transformation). We thus focus our attention on the estimation of the pairwise distances (2).

1.1. Related work

The model we consider in (1) is an example of a latent graph model and the points are often called latent positions. In its full generality, the model includes the planted partition model popular in the area of graph partitioning. To see this, take $r = 1$ and let v denote the number of blocks and, with e_k denoting the k -th canonical basis vector, set $x_i = e_k$ if i belongs to block k . The planted partition model is a special case of the stochastic block model of Holland et al. [18]. This is also a special case of our model, as can be seen by changing e_k to z_s chosen so that $\phi(\|z_s - z_\ell\|) = p_{k\ell}$, where $p_{k\ell}$ denotes the connection probability between blocks k and ℓ . Mixed-membership stochastic block models as in [2, 4, 44] are also special cases of latent graph models, but of a slightly different kind. The literature on the stochastic block model is now substantial and includes results on the recovery of the underlying communities; see, e.g., [1, 11, 16, 21, 26, 30, 38] and references therein.

Our contribution here is of a different nature as we focus on the situation where the latent positions are well spread out in space, forming no obvious clusters. This relates more closely to the work of Holland et al. [17]. Although their setting is more general in that additional information may be available at each position, without that additional information their approach reduces to the following logistic regression model:

$$\log\left(\frac{p_{ij}}{1 - p_{ij}}\right) = -d_{ij},$$

which is clearly a special case of (1) with link function the logistic function. Sarkar et al. [31] consider this same model motivated by a link prediction problem where the nodes are assumed to be embedded in space with their Euclidean distances being the dissimilarity of interest. In fact, they assume that the points are uniformly distributed in some region. They study a method based on the number of neighbors that a pair of nodes have in common, which is one of the main methods for link prediction [22, 23]. Parthasarathy et al. [28] consider a more general setting where a noisy neighborhood graph is observed: if (x_i) are points in a metric space with pairwise distances (d_{ij}) , then an adjacency matrix, $W = (W_{ij})$, is observed, where $W_{ij} = 1$ with probability $1 - p$ if $d_{ij} \leq r$ and with probability q if $d_{ij} > r$, where $p, q \in [0, 1]$ are parameters of the model. Under fairly general conditions on the metric space and the sampling distribution, and additional conditions on (n, r, p) , they show that the graph distances computed based on W provide, with high probability, a 2-approximation to the underlying distances in the case where $q = 0$. In the case where $q > 0$, the same is true, under some conditions on (n, r, p, q) , if W is replaced by $\tilde{W} = (\tilde{W}_{ij})$ where $\tilde{W}_{ij} = 1$ exactly when $N_{ij}/(N_i + N_j - N_{ij}) \geq \tau$, where τ is a carefully chosen tuning parameter and where $N_i := \#\{j : W_{ij} = 1\}$ (number of neighbors of i) and $N_{ij} := \#\{k : W_{ik} = W_{jk} = 1\}$ (number of common neighbors of i and j).

Scheinerman and Tucker [32] and Young and Scheinerman [46] consider what they call a dot-product random graph model where $p_{ij} = \langle x_i, x_j \rangle$, where it is

implicitly assumed that $\langle x_i, x_j \rangle \in [0, 1]$ for all $i \neq j$. This model is a special case of (1), with $\phi(d) = 1 - \frac{1}{2}d^2$. Sussman et al. [39] consider recovering the latent positions in this model with full knowledge of the link function. They devise a spectral method which consists in embedding the items $\{1, \dots, n\}$ as points in \mathbb{R}^v , with v assumed known, as the row vectors of $U_{(v)}\Theta_{(v)}^{1/2}$, where $W = U\Theta V^\top$ is the SVD of W , and for a matrix $A = (A_{ij})$ and an integer $s \geq 1$, $A_{(s)} = (A_{ij} : i \vee j \leq s)$. They analyze their method in a context where the latent positions are in fact a sample from a possibly unknown distribution. The same authors extended their work in [40] to an arbitrary link function, which may be unknown, although the focus is on a binary classification task in a setting where for each $i \in [n]$ a binary label y_i is available.

Alamgir and von Luxburg [3, 43] consider the closely related problem of recovering the latent positions in a setting where a nearest-neighbor graph is available. They propose a method based on estimating the underlying density denoted f . If \hat{f}_i denotes the density estimate at x_i , a graph is defined on $[n]$ with weights $w_{ij} = (\hat{f}_i^{-1/v} + \hat{f}_j^{-1/v})/2$, and d_{ij} is estimated by the graph distance between nodes i and j .

Latent positions random graph models also play a role in the literature on rankings¹ [15, 25]. A typical parametric model represents each player $i \in [n]$ by a number x_i such that the probability that i wins against j in a single game is $\phi(x_i - x_j)$. Note that the link function is applied to the difference and not the absolute value of the difference. For example, the Bradley–Terry–Luce model [8, 24] uses the logistic link function. Suppose that multiple games are played between multiple pairs of players. The result of that can be summarized as $(W_{ij} : i \neq j)$, where W_{ij} is the number of games where i prevailed over j . This is the weight matrix of a directed latent positions graph where the positions are (x_1, \dots, x_n) . We refer the reader to [27, 33] and references therein for theoretical results developed for such models.

1.2. Our contribution

Graph distances are well-known estimates for the Euclidean distances in the context of graph drawing [20, 35], where the goal is to embed items in space based on an incomplete distance matrix. They also appear in the literature on link prediction [22, 23] and are part of the method proposed in [43]. We examine the use of graph distances for the estimation of the Euclidean distances (2). As we shall see, the graph distances are directly useful when the link function ϕ is compactly supported, which is for example the case in the context of a neighborhood graph where $\phi(d) = \mathbb{I}\{d \leq r\}$ for some connectivity radius $r > 0$. In fact, the method is shown to achieve a minimax lower bound in this setting (under a convexity assumption). This setting is discussed in Section 2. In Section 3, we extend the analysis to other (compactly supported) link functions. We end with Section 4, where we discuss some important limitations of the method based on graph distances and consider some extensions, including localization (to avoid

¹Thanks to Philippe Rigollet for pointing out this out to us.

the convexity assumption) and the use of the number of common neighbors (to accommodate non-compact link functions). Proofs are gathered in Section 5.

1.3. The graph distance method

Given the adjacency matrix W , the graph distance (aka shortest-path distance) between nodes i and j is defined as

$$\delta_{ij} := \inf \left\{ \ell : \exists k_0, \dots, k_\ell \in [n] \text{ s.t.} \right. \\ \left. k_0 = i, k_\ell = j, \text{ and } W(k_{s-1}, k_s) = 1, \forall s \in [\ell] \right\}, \quad (4)$$

where $\inf \emptyset = \infty$ by convention. Here and elsewhere, we will sometimes use the notation $W(i, j)$ for W_{ij} , $d(i, j)$ for d_{ij} , etc.

We propose estimating, up to a scale factor, the Euclidean distances (2) with the graph distances (4). Indeed, since ϕ is assumed unknown, the scale factor cannot be recovered from the data, as is the case in ordinal embedding, for example. Therefore, estimates are necessarily up to an arbitrary scaling factor, so that the accuracy of an estimator $\hat{d} = (\hat{d}_{ij})$ for $d = (d_{ij})$ is measured according to how close we can make $s\hat{d}$ and d in some chosen way by choosing the scale $s > 0$ with (oracle) knowledge of d . For example, with mean squared error, this leads to quantifying the accuracy of \hat{d} as follows

$$\min_{s>0} \sum_{i<j} (s\hat{d}_{ij} - d_{ij})^2.$$

The graph distance method is the analog of the MDS-D method of Kruskal and Seery [20] for graph drawing, which is a setting where some of the distances (2) are known and the goal is to recover the missing distances. Let \mathcal{E} denote the set of pairs $i < j$ for which d_{ij} is known. MDS-D estimates the missing distances with the distances in the graph with node set $[n]$ and edge set \mathcal{E} , and with edge $(i, j) \in \mathcal{E}$ weighed by d_{ij} . This method was later rediscovered by Shang et al. [35], who named it MDS-MAP(P), and coincides with the IsoMap procedure of Tenenbaum et al. [41] for isometric manifold embedding. (For more on the parallel between graph drawing and manifold embedding, see [10].)

As we shall see, the graph distance method is most relevant when the positions are sufficiently dense in their convex hull, which is a limitation it shares with MDS-D. For $\Omega \subset \mathbb{R}^v$ and $x_1, \dots, x_n \in \mathbb{R}^v$, define

$$\Lambda_\Omega(x_1, \dots, x_n) = \sup_{x \in \Omega} \min_{i \in [n]} \|x - x_i\|, \quad (5)$$

which measures how dense the latent points are in Ω . We also let $\Lambda(x_1, \dots, x_n)$ denote (5) when Ω is the convex hull of $\{x_1, \dots, x_n\}$.

2. Simple setting

In this section we focus on the simple, yet emblematic case of a neighborhood (ball) graph, that is, a setting where the link function is given by $\phi(d) = \mathbb{I}\{d \leq r\}$ for some $r > 0$. When the positions are drawn iid from some distribution, the result is what is called a random geometric graph [29], but here we consider the positions to be deterministic. In particular, the setting is not random.

We start with a performance bound for the graph distance method and then establish minimax lower bound. Similar results are available in [6, 9, 28], among other places, and we only provide a proof for completeness, and also to pave the way to the more sophisticated Theorem 3.

Theorem 1. *Consider a set of points x_1, \dots, x_n that satisfy $\Lambda(x_1, \dots, x_n) \leq \varepsilon$. Assume that $\phi(d) = \mathbb{I}\{d \leq r\}$ for some $r > 0$, and define $\hat{d}_{ij} = r\delta_{ij}$. If the connectivity radius r is sufficiently larger than the density of the point set ε , specifically if $\varepsilon \leq r/4$, then*

$$0 \leq \hat{d}_{ij} - d_{ij} \leq 4(\varepsilon/r)d_{ij} + r, \quad \forall i, j \in [n],$$

which in particular implies that

$$\left[\frac{1}{\binom{n}{2}} \sum_{i < j} (\hat{d}_{ij} - d_{ij})^2 \right]^{1/2} \leq 4(\varepsilon/r)\rho + r,$$

where ρ is the diameter of $\{x_1, \dots, x_n\}$.

In the statement, \hat{d} is not a true estimator in general as it relies on knowledge of r , which may not be available, nor be estimable, if the link function is unknown. Nevertheless, the result says that, up to that scale parameter, the graph distances achieve a nontrivial level of accuracy. Compare with [28, Th 2.5], which in the context of points on a Euclidean space as considered here says that, in a stochastic setting where the points are generated iid from some distribution supported on a convex set, $\max_{i,j}(\hat{d}_{ij} - d_{ij})$ is bounded by r in the limit where $n \rightarrow \infty$ while r remains fixed.

For a numerical example, see Figure 1. In Figure 2 we confirm numerically that the method is biased when the underlying domain from which the positions are sampled is not convex. That said, the method is robust to mild violations of the convex constraint, as shown in Figure 3, where the positions correspond to $n = 3000$ US cities.² (Computations were done in R, with the graph distances computed using the `igraph` package, to which Classical Scaling was applied, followed by a procrustes alignment and scaling using the package `vegan`.)

²These were sampled at random from the dataset available at simplemaps.com/data/us-cities.

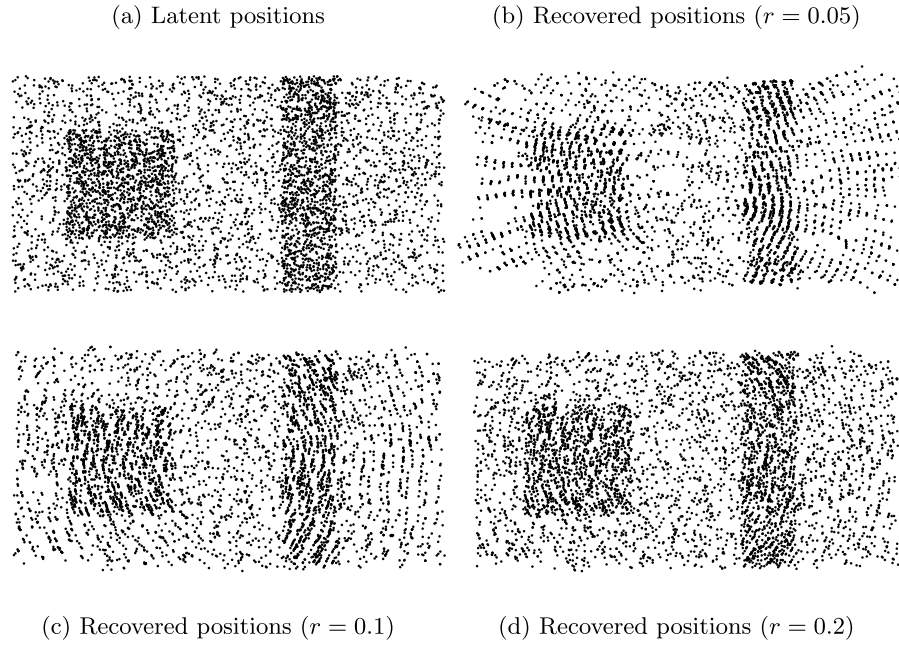


FIG 1. A numerical example illustrating the setting of Theorem 1. Here $n_0 = 3000$ positions were sampled uniformly at random from $\Omega_0 := [0, 2] \times [0, 1]$, $n_1 = 1000$ from $\Omega_1 := [0.25, 0.75] \times [0.25, 0.75]$, and $n_2 = 1000$ from $\Omega_2 := [1.25, 1.5] \times [0, 1]$, for a total of $n = 5000$ positions.

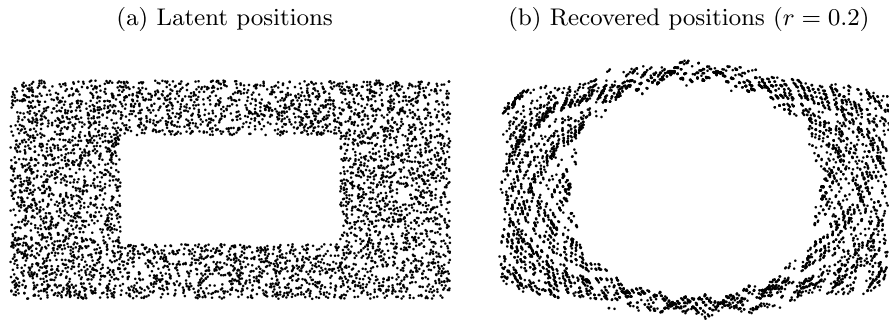


FIG 2. A numerical example illustrating the setting of Theorem 1 showing that the convexity constraint is indeed required for the graph distance method to be unbiased. Here $n = 5000$ positions were sampled uniformly at random from $\Omega := [0, 2] \times [0, 1] \setminus [0.5, 1.5] \times [0.25, 0.75]$.

Remark 1. If we apply Classical Scaling to \hat{d} , we obtain an embedding with arbitrary scaling and rigid positioning, which are not recoverable when r is unknown. Nevertheless, if we apply the perturbation bound recently established in [5, Cor 2], the recovery of the latent positions is of order at most $O(\varepsilon/r + r)$.

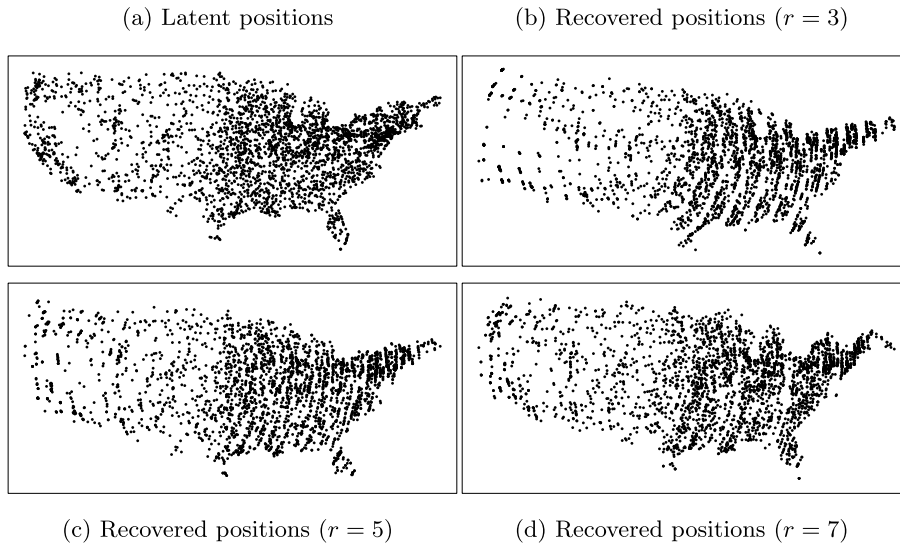


FIG 3. A numerical example illustrating the setting of Theorem 1. The latent positions are located at the coordinates of $n = 3000$ US cities and the connectivity radius varies (in degrees).

It turns out that the graph distance method comes close to achieving the best possible performance (understood in a minimax sense) in this particularly simple setting. Indeed, we are able to establish the following general lower bound that applies to any method.

Theorem 2. Assume that $\phi(d) = \mathbb{I}\{d \leq r\}$ with $r > 0$ known. Then there is a numeric constant $c_0 > 0$ with the property that, for any $\varepsilon > 0$ and any estimator³ \hat{d} , there is a point set x_1, \dots, x_n such that $\Lambda(x_1, \dots, x_n) \leq \varepsilon$ and, for at least half of the pairs $i \neq j$,

$$|\hat{d}_{ij} - d_{ij}| \geq c_0 \left(\frac{\varepsilon}{r \vee \varepsilon} d_{ij} + r \right), \tag{6}$$

and also, for another numeric constant $c_1 > 0$,

$$\left[\frac{1}{\binom{n}{2}} \sum_{i < j} (\hat{d}_{ij} - d_{ij})^2 \right]^{1/2} \geq c_1 \left(\frac{\varepsilon}{r \vee \varepsilon} \rho + r \right),$$

where ρ is the diameter of the point set.

Thus, in the strictest sense, the graph distance method is, for this particular link function, minimax optimal (in order of magnitude). It turns out that the point configurations that we consider in the proof are all embedded on the real

³An estimator here is simply a function on the set of n -by- n symmetric binary matrices with values in $\mathbb{R}_+^{n(n-1)/2}$.

line and thus, in principle, can be embedded in any Euclidean space. It is also the case that, for these particular configurations, it does not help if we know that $x_1 < \dots < x_n$.

We speculate that a better error rate can be achieved (in probability) under a stochastic model, for example, when x_1, \dots, x_n are drawn iid from the uniform distribution on some ‘nice’ domain of a Euclidean space. On the other hand, we anticipate that our performance analysis of the graph distance method is essentially tight even then. To achieve a better performance, more sophisticated methods need to be considered. Methods using neighbors-of-neighbors information [28, 31] are particularly compelling, but in principle require knowing (or perhaps estimating) the underlying density if it is unknown. We probe this question a little further in Section 4.3 with some simple but promising numerical experiments. (All we know about this approach is that it can lead to a 2-approximation [28].)

3. General setting

Beyond the setting of a neighborhood graph considered in Section 2, the graph distance method, in fact, performs similarly when the link function is discontinuous at the edge of its support, meaning when it drops abruptly to 0. A case in point is when $\phi(d) = p\mathbb{1}\{d \leq r\}$ for $p > 0$ constant, which corresponds to a random geometric graph with its edges independently deleted with probability $1 - p$. See Figure 4 for a numerical example illustrating this particular case.

More generally, we establish the performance of the graph distance method when the link function is compactly supported. The bound we obtain is in terms of how fast the function approaches 0 at the edge of its support. Note that, unlike the setting of a neighborhood graph, the model is truly random when the link function is not an indicator function, so that the statement below is in probability.

Theorem 3. *Assume that ϕ has support $[0, r]$, for some $r > 0$, and define $\hat{d}_{ij} = r\delta_{ij}$. Assume that, for some $C_0 > 0$ and $\alpha \geq 0$, $\phi(d) \geq C_0(1 - d/r)^\alpha$ for all $d \in [0, r]$. Then there are $C_1, C_2 > 0$ depending only on (α, C_0) such that, whenever $r/\varepsilon \geq C_1(\log n)^{1+\alpha}$, for any points x_1, \dots, x_n that satisfy $\Lambda(x_1, \dots, x_n) \leq \varepsilon$, with probability at least $1 - 1/n$,*

$$0 \leq \hat{d}_{ij} - d_{ij} \leq C_2[(\varepsilon/r)^{\frac{1}{1+\alpha}} d_{ij} + r], \quad \forall i \neq j, \quad (7)$$

which in particular implies that

$$\left[\frac{1}{\binom{n}{2}} \sum_{i < j} (\hat{d}_{ij} - d_{ij})^2 \right]^{1/2} \leq C_2[(\varepsilon/r)^{\frac{1}{1+\alpha}} \rho + r],$$

where ρ is the diameter of $\{x_1, \dots, x_n\}$.

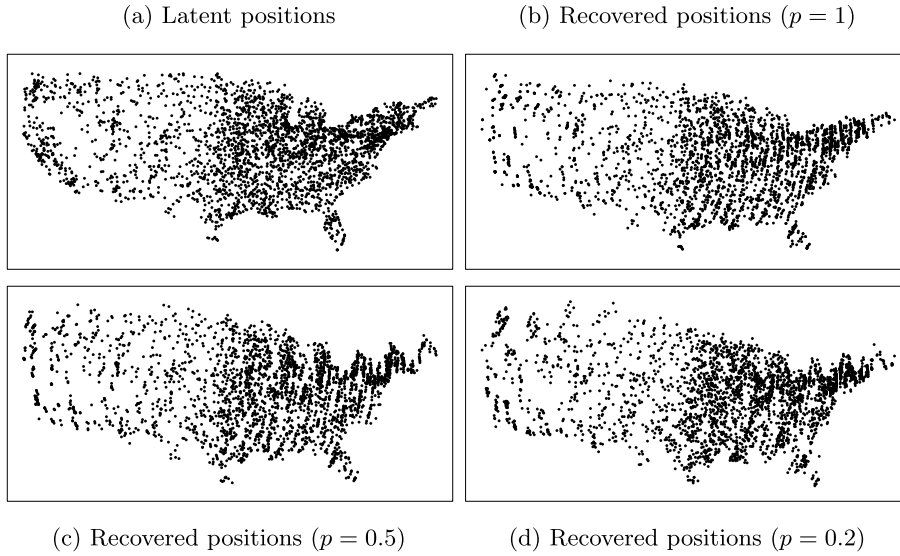


FIG 4. Same setting as in Figure 3. Here we set $r = 5$ and vary p . In fact, to ease the comparison, we coupled the different adjacency matrices in the sense that the $(p = 0.2)$ -matrix was built by erasing edges from the $(p = 0.5)$ -matrix independently with probability $0.2/0.5 = 0.4$.

Although we believe our performance analysis in (7) to be tight, we do not know whether it is minimax optimal in any way.

For the graph distance, we expect it to be less accurate the slower the link function ϕ approaches 0. This is borne out in some numerical experiments that we performed. In those experiments, $n = 5000$ points were drawn uniformly at random from $[0, 1]^2$ considered as a torus to avoid boundary effects. (Clearly, our results apply in this setting as well.) For each $\alpha \in \{0, 0.1, \dots, 0.9, 1, 2, 3, 4, 5\}$, we computed a realization of the adjacency matrix with link function

$$\phi_\alpha(d) := c_\alpha [1 \wedge (2 - 2d/r)_+^\alpha], \quad c_\alpha := \frac{1}{2} \frac{\alpha^2 + 3\alpha + 2}{\alpha^2 + 5\alpha + 8},$$

chosen so that $\mathbb{P}(W_{ij} = 1)$ is the same regardless of α as long as $r \leq 0.5$. In our experiments, we chose $r = 0.1$. This was repeated 100 times. The results are presented in Figure 5, where

$$\text{relative error} := \left[\frac{\sum_{i < j} (\hat{d}_{ij} - d_{ij})^2}{\sum_{i < j} d_{ij}^2} \right]^{1/2}. \tag{8}$$

4. Discussion

The method based on graph distances suffers from a number of serious limitations:

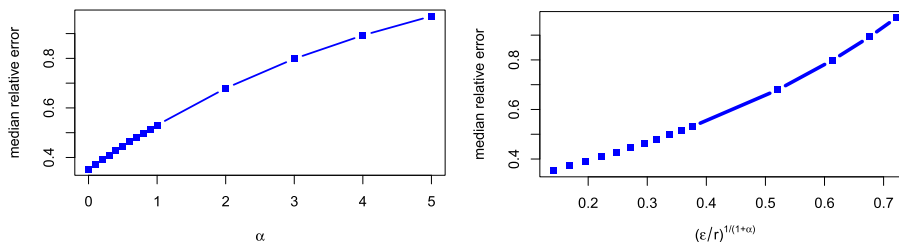


FIG 5. Here $n = 5000$ points were drawn uniformly at random from the 2D unit torus. The radius was set at $r = 0.1$ and the link function varied with α as specified in (8). The median relative error is over 100 repeats. The density ε set at 0.025, as determined by simulation.

1. The positions need to span a convex set, although the method is robust to mild violations of this constraint as exemplified in Figure 3.
2. Even in the most favorable setting of Section 2, the relative error is still of order r , as established in (7). This is clearly tight for the graph distance method, and although it matches the lower bound established in Theorem 2, this bias could potentially be avoided when the positions are nicely spread out, for example, as a random sample from some nice distribution is expected to be.
3. The link function needs to be compactly supported. Indeed, the method can be grossly inaccurate in the presence of long edges, as in the interesting case where the link function is of the form

$$\phi(d) = p\mathbb{I}\{d \leq r\} + q\mathbb{I}\{d > r\}, \quad (9)$$

where $0 < q < p \leq 1$, as considered in [28].

We address each of these three issues in what follows.

4.1. Localization

A possible approach to addressing Issue 1 is to operate locally. This is well-understood and is what lead Shang and Ruml [34] to suggest MDS-MAP(P), which effectively localizes MDS-MAP [35]. (As we discussed earlier, the latter is essentially a graph-distance method and thus bound by the convexity constraint.) More recent methods for graph drawing based on ‘synchronization’ also operate locally [12, 13].

Experimentally, this strategy works well. See Figure 6 for a numerical example, which takes place in the context of the rectangle with a hole of Figure 2. We adopted a simple approach: we kept the graph distances that were below a threshold, leaving the other ones unspecified, and then applied a method for multidimensional scaling with missing values, specifically SMACOF [14] (initialized with the output of the graph distance method).


(a) Latent positions (b) Recovered positions ($r = 0.2, \delta_{ij} \leq 2$)

FIG 6. Same setting as in Figure 2.


(a) Latent positions (b) Recovered positions ($r = 0.5$)

FIG 7. Here $n = 2000$ positions were sampled uniformly at random from $\Omega := [0, 2] \times [0, 1]$. For all $i \neq j$, $\delta_{ij} \in \{1, 2, 3, 4, 5\}$, so that the graph distances are rather discrete, yet the embedding computed by classical multidimensional scaling is surprisingly accurate.

4.2. Regularization

Regarding Issue 2, in numerical experiments we have found that the graph distances, although grossly inaccurate, are nevertheless useful for embedding the points using (classical) multidimensional scaling. Thus, if one is truly interested in estimating the Euclidean distances, one may use graph distances as rough proxies for the underlying distances, apply multidimensional scaling, and then compute the distances between the embedded points. For a numerical illustration, see Figure 7. This phenomenon remains surprising to us and we do not have a good understanding of the situation.

4.3. Number of common neighbors

A possible approach to addressing Issue 3, as well as Issue 2, is to work with the number of common neighbors, which provides an avenue to ‘super-resolution’ in a way, at least when the positions are sampled iid from a known distribution such as the uniform distribution on a domain (known and convex). By this we

mean that, say in the simple setting of Section 2, although the adjacency matrix only tells whether two positions are within distance r , it is possible to gather all this information to refine this assessment. Similarly, in the setting where (9) is the link function, it is possible to tell whether two positions are nearby or not. This sort of concentration is well-known to the expert and seems to be at the foundation of spectral methods (see, e.g., [39, Prop 4.2]). We refer the reader to [28, 31], where such an approach is considered in greater detail.

5. Proofs

5.1. Proof of Theorem 1

Fix $i, j \in [n]$ distinct.

Let $m := \lfloor d_{ij}/(r - 2\varepsilon) \rfloor$ and note that $m(r - 2\varepsilon) \leq d_{ij} \leq (m + 1)(r - 2\varepsilon)$. For $s \in \{0, \dots, m + 1\}$, let $z_s = x_i + \frac{s}{m+1}(x_j - x_i)$. We have $z_0 = x_i$ and $z_{m+1} = x_j$, and z_0, z_1, \dots, z_{m+1} are on the line joining x_i and x_j and satisfy $\|z_s - z_{s+1}\| \leq r - 2\varepsilon$ for all s . Let x_{k_s} be such that $\|z_s - x_{k_s}\| \leq \varepsilon$, with $x_{k_0} = x_i$ and $x_{k_{m+1}} = x_j$. Note that x_{k_s} is well-defined since z_s belongs to the convex hull of $\{x_1, \dots, x_n\}$ and we have assumed that $\Lambda(x_1, \dots, x_n) \leq \varepsilon$. By the triangle inequality, for all $s \in \{0, \dots, m\}$,

$$\|x_{k_s} - x_{k_{s+1}}\| \leq \|x_{k_s} - z_s\| + \|z_s - z_{s+1}\| + \|z_{s+1} - x_{k_{s+1}}\| \leq \varepsilon + (r - 2\varepsilon) + \varepsilon \leq r.$$

Hence, $(x_{k_0}, x_{k_1}, \dots, x_{k_{m+1}})$ forms a path in the graph, and as a consequence, $\delta_{ij} \leq m + 1$. In turn, this implies that

$$\hat{d}_{ij} = r\delta_{ij} \leq rm + r \leq r \frac{d_{ij}}{r - 2\varepsilon} + r \leq d_{ij} + 4(\varepsilon/r)d_{ij} + r,$$

using the fact that $\varepsilon \leq r/4$.

Resetting the notation, let $k_0 = i, k_1, \dots, k_\ell = j$ denote a shortest path joining i and j , so that $\ell = \delta_{ij}$. By the triangle inequality,

$$d_{ij} = \|x_{k_0} - x_{k_\ell}\| \leq \sum_{s=0}^{\ell-1} \|x_{k_s} - x_{k_{s+1}}\| \leq \ell r = r\delta_{ij} = \hat{d}_{ij}, \quad (10)$$

using the fact that $\|x_{k_s} - x_{k_{s+1}}\| \leq r$ for all s .

5.2. Proof of Theorem 2

First term on the RHS of (6) We construct two point configurations that yield the same adjacency matrix and then measure the largest difference between the corresponding sets of pairwise distances. Assume that $r \leq 1/2$ (without loss of generality) and that $m := r(n - 1)$ is an integer for convenience. We define two configurations of points, both in $\Omega := [0, 1]$ (so that $v = 1$ here).

- *Configuration 1.* In this configuration,

$$x_i = \frac{i - 1}{n - 1}, \quad i \in [n].$$

Note that $\Lambda_\Omega(x_1, \dots, x_n) = 1/(2n - 2)$.

- *Configuration 2.* In this configuration,

$$x_i = \frac{(i - 1)(1 - \eta(i - 1))}{(n - 1)(1 - \eta(n - 1))}, \quad i \in [n].$$

for some $\eta > 0$ chosen small later on. When $\eta \leq 1/(2n - 3)$, which we assume, x_i is increasing with $x_1 = 0$ and $x_n = 1$. Note that

$$\Lambda_\Omega(x_1, \dots, x_n) = \frac{1 - \eta}{(2n - 2)(1 - \eta(n - 1))}.$$

The two configurations coincide when $\eta = 0$, but we will choose $\eta > 0$ in what follows. Under Configuration 1, the adjacency matrix W is given by $W_{ij} = \mathbb{I}\{|i - j| \leq m\}$. For the design matrix to be the same under Configuration 2, it suffices that x_1 have (exactly) m neighbors (to the right) and that x_n have (exactly) m neighbors (to the left); this is because $i \mapsto x_i - x_{i-1}$ is decreasing in this configuration. These two conditions correspond to four equations, given by

$$x_{m+1} - x_1 \leq r, \quad x_{m+2} - x_1 > r, \quad |x_n - x_{n-m}| \leq r, \quad x_n - x_{n-m-1} > r.$$

We need only consider the first and fourth as they imply the other two. After some simplifications, we see that the first one holds when $r \leq 1$, while the fourth holds when $r \leq 1 - 2/(n - 1)$ and $\eta \leq 1/(2n - 3 + m(n - m - 3))$. Since $r \leq 1/2$, $r \leq 1 - 2/(n - 1)$ when $n \geq 5$, and we choose $\eta = 1/(2n + m(n - m))$ for example. Then $\Lambda_\Omega(x_1, \dots, x_n) \sim 1/2n$ in Configuration 2 (same as in Configuration 1). We choose $n = n_\varepsilon$ just large enough that $\Lambda_\Omega(x_1, \dots, x_n) \leq \varepsilon$ in both configurations. In particular, $\varepsilon \sim 1/n_\varepsilon$ as $\varepsilon \rightarrow 0$. Since the result only needs to be proved for ε small, we may take n as large as we need.

Now that the two designs have the same adjacency matrix, we cannot distinguish them with the available information. It therefore suffices to look at the difference between the pairwise distances. Let $d_{ij}^{(k)}$ denote the distance between x_i and x_j in Configuration k . For $i < j$, we have

$$d_{ij}^{(1)} = \frac{j - i}{n - 1}, \quad d_{ij}^{(2)} = \frac{j - i}{n - 1} \frac{1 - \eta(j + i - 2)}{1 - \eta(n - 1)}.$$

In particular,

$$\frac{1}{2}d_{ij}^{(1)} \leq d_{ij}^{(2)} \leq 2d_{ij}^{(1)},$$

by the fact that $\eta n \leq 1/2$. Also,

$$d_{ij}^{(1)} - d_{ij}^{(2)} = \frac{j - i}{n - 1} \frac{\eta(i + j - n - 1)}{1 - \eta(n - 1)},$$

which leads to

$$\begin{aligned} |d_{ij}^{(1)} - d_{ij}^{(2)}| &= d_{ij}^{(1)} \frac{\eta|i+j-n-1|}{1-\eta(n-1)} \\ &\geq d_{ij}^{(1)} \frac{\eta n}{1-\eta(n-1)} \frac{|i+j-n-1|}{n} \\ &\geq C(d_{ij}^{(1)} \vee d_{ij}^{(2)}) \frac{\varepsilon}{\varepsilon \vee r} \frac{|i+j-n-1|}{n}, \end{aligned}$$

for some universal constant $C > 0$, using the fact that $\eta n \leq 1/2$ and $\eta n \asymp 1/(1 \vee rn) \asymp \varepsilon/(\varepsilon \vee r)$, the latter because $\varepsilon \asymp 1/n$ in our construction. Since $|i+j-n-1| \geq n/10$ for most pairs of indices $i < j$, the following is also true

$$|d_{ij}^{(1)} - d_{ij}^{(2)}| \geq (C/10) \frac{\varepsilon}{r \vee \varepsilon} (d_{ij}^{(1)} \vee d_{ij}^{(2)}).$$

To conclude, since the two configurations have the same adjacency matrix, they are indistinguishable solely based on that information, and so it must be that for any estimator \hat{d} , for most pairs $i < j$,

$$\begin{aligned} |\hat{d}_{ij} - d_{ij}| &\geq |\hat{d}_{ij} - d_{ij}^{(1)}| \vee |\hat{d}_{ij} - d_{ij}^{(2)}| \\ &\geq \frac{1}{2} |d_{ij}^{(1)} - d_{ij}^{(2)}| \\ &\geq (C/20) \frac{\varepsilon}{r \vee \varepsilon} d_{ij}, \end{aligned}$$

where d_{ij} denotes $d_{ij}^{(k)}$ if the true configuration is Configuration k .

Second term on the RHS of (6) We construct again two point configurations, also on the real line, that have the same adjacency matrix. We assume that $\varepsilon \leq r$ for otherwise the first term in the RHS of (6) is of order $\asymp 1$ for most pairs of indices. In fact, we assume that $q = r/\varepsilon$ is an integer for simplicity.

To any pattern $y_0 = 0 < y_1 < \dots < y_m = r$ with $y_j - y_{j-1} \leq 2\varepsilon$ for all j , associate the point set $x_1 < \dots < x_n$ where $x_i = y_{i \bmod m} + \lfloor i/m \rfloor r$. Note that the x point set is built by repeating the y pattern. As can be readily seen, all these point sets have the same adjacency matrix $W_{ij} = \mathbb{I}\{|i-j| \leq m\}$ and have Λ bounded by ε . We now consider two particular cases. Take m even so that $m = 2q$ for some integer $q > 0$, and define the following configurations.

- *Configuration 3.* Here x_1, \dots, x_n is defined based on $y_j = j\varepsilon$ for $j = 1, \dots, q-1$ and $y_j = (q-1)\varepsilon + (j-q+1)\eta$ for $j = q, \dots, 2q$, where $\eta := \varepsilon/(q+1)$.
- *Configuration 4.* Here the configuration is obtained by reversing the order of the previous one, namely, x_1, \dots, x_n is based on $y_j = j\eta$ for $j = 1, \dots, q$ and $y_j = (j-q)\varepsilon$ for $j = q+1, \dots, 2q$.

Let $d_{ij}^{(k)}$ denote the distance between x_i and x_j in Configuration k . Letting $b_i := i \bmod 2q$, we have

$$d_{ij}^{(3)} = (2q - b_i)\eta + (q - 1 - b_i)_+(\varepsilon - \eta)$$

$$\begin{aligned}
 &+ b_j \eta + b_j \wedge (q - 1)(\varepsilon - \eta) + r(\lfloor j/2q \rfloor - \lfloor i/2q \rfloor - 1), \\
 d_{ij}^{(4)} &= (2q - b_i)\eta + (2q - b_i) \wedge (q - 1)(\varepsilon - \eta) \\
 &+ b_j \eta + (b_j - q - 1)_+(\varepsilon - \eta) + r(\lfloor j/2q \rfloor - \lfloor i/2q \rfloor - 1),
 \end{aligned}$$

resulting in

$$\begin{aligned}
 &|d_{ij}^{(3)} - d_{ij}^{(4)}| \\
 &= (\varepsilon - \eta) \cdot \begin{cases} |b_j - b_i| & \text{if } b_i, b_j \leq q - 1 \text{ or } b_i, b_j \geq q + 1 \text{ or } b_i = b_j = q, \\ |2q - b_i - b_j| & \text{if } b_i \leq q - 1, b_j \geq q + 1 \text{ or } b_i \geq q + 1, b_j \leq q - 1, \\ |b_j - b_i - 1| & \text{if } b_i = q, b_j \geq q + 1 \text{ or } b_i \leq q - 1, b_j = q, \\ |b_j - b_i + 1| & \text{if } b_i = q, b_j \leq q - 1 \text{ or } b_i \geq q + 1, b_j = q, \end{cases}
 \end{aligned}$$

where for a real a , $a_+ := \max(0, a)$. We have $\varepsilon - \eta = q\varepsilon/(q + 1) = r/(q + 1)$, and elementary considerations confirm that for most $i < j$, the factor defined in the curly bracket above is $\geq C(q + 1)$ for some universal constant $C > 0$. Hence, $|d_{ij}^{(3)} - d_{ij}^{(4)}| \geq Cr$ for most pairs of indices, and this then implies as before that for any estimator \hat{d} , for most pairs $i < j$,

$$|\hat{d}_{ij} - d_{ij}| \geq (C/2)r,$$

where d_{ij} denotes $d_{ij}^{(k)}$ if the true configuration is Configuration k .

5.3. Proof of Theorem 3

In the following, C_0, C_1, C_2 refer to the constants appearing in the statement of Theorem 3, while c_1, c_2, \dots denote positive constants that only depend on (α, C_0) . Since the result only needs to be proved for large r/ε , we will take this quantity as large as needed. In what follows, we connect each node in the graph to itself. This is only for convenience and has no impact on the validity of the resulting arguments.

As before in (10), we have $\hat{d}_{ij} \geq d_{ij}$ for all $i \neq j$. Recall the definition of $p_{ij} \equiv p(i, j)$ in (3). Let $p_0 = \phi(r/2) > 0$ and note that $p_0 \geq C_0(1/2)^\alpha$.

Special case Suppose that $d_{kl} \leq r/2$ for all $k \neq l$. In that case, for all $i \neq j$, $p_{ij} = \phi(d_{ij}) \geq \phi(r/2) = p_0$. For (i, j, k) distinct, (x_i, x_k, x_j) forms a path in the graph if and only if $W_{ik}W_{kj} = 1$, which happens with probability $p_{ik}p_{kj} \geq p_0^2$. Therefore, by independence,

$$\mathbb{P}(\delta_{ij} > 2) \leq \mathbb{P}(W_{ik}W_{kj} = 0, \forall k \notin \{i, j\}) \leq (p_0^2)^{n-2}.$$

Therefore, by the union bound, with probability at least $1 - n^2 p_0^{2n-4} \geq 1 - n^2 \exp(-c_1 n)$, we have $\delta_{ij} \leq 2$, implying $\hat{d}_{ij} \leq 2r$, for all $i \neq j$.

Henceforth, we assume that

$$\max_{i \neq j} d_{ij} > r/2. \tag{11}$$

Claim 1. *By choosing C_1 large enough, the following event happens with probability at least $1 - 1/n^2$,*

$$\mathcal{A}_1 := \{\hat{d}_{ij} \leq 9d_{ij} + 2r, \text{ for all } i, j \in [n]\}.$$

Take i, j such that $d_{ij} \leq r/4$. We first note that there is j_* such that $d(i, j_*) > r/4$, for otherwise, for all $k, l \in [n]$, $d_{kl} \leq d_{ki} + d_{il} \leq r/4 + r/4 = r/2$, which would contradict our assumption (11).

Define

$$z_s = x_i + s \frac{\varepsilon(x_{j_*} - x_i)}{d(j_*, i)}, \quad s \in [m],$$

where $m := \lfloor (r/4 - \varepsilon)/\varepsilon \rfloor$. By construction each z_s is on the line segment joining x_i and x_{j_*} , and so belongs to the convex hull of x_1, \dots, x_n ; hence, by the fact that $\Lambda(x_1, \dots, x_n) \leq \varepsilon$, there is $i_s \in [n]$ be such that $\|x_{i_s} - z_s\| \leq \varepsilon$. By the triangle inequality,

$$d(i, i_s) = \|x_i - x_{i_s}\| \leq \|x_i - z_s\| + \|z_s - x_{i_s}\| \leq s\varepsilon + \varepsilon \leq m\varepsilon + \varepsilon \leq (r/4 - \varepsilon) + \varepsilon = r/4,$$

and

$$d(i_s, j) \leq d(i_s, i) + d(i, j) \leq r/4 + r/4 = r/2.$$

Therefore, for each $s \in [m]$, (x_i, x_{i_s}, x_j) forms a path with probability at least p_0^2 . By independence, therefore, there is such an $s \in [m]$ with probability at least $1 - (1 - p_0^2)^m$.

With the union bound and the fact that $m \geq r/5\varepsilon$ when r/ε is large enough, we may conclude that, if C_1 is chosen large enough, the event

$$\mathcal{A}_2 := \{\hat{d}_{ij} \leq 2r \text{ for all } i \neq j \text{ such that } d_{ij} \leq r/4\},$$

has probability at least $1 - 1/n^2$. Indeed,

$$\mathbb{P}(\mathcal{A}_2^c) \leq n^2(1 - p_0^2)^m \leq n^2 \exp(-c_2(r/5\varepsilon)) \leq 1/n^2,$$

eventually, when $r/\varepsilon \geq C_1(\log n)^{1+\alpha}$ with C_1 large enough.

Next, we prove that \mathcal{A}_2 implies \mathcal{A}_1 , which will suffice to establish the claim. For this, we consider the remaining case where i, j are such that $d_{ij} > r/4$. Define $z_0 = x_i$ and

$$z_s = x_i + s \frac{\varepsilon(x_j - x_i)}{d(j, i)}, \quad \text{for } s \in [m],$$

where this time $m := \lfloor d_{ij}/\varepsilon \rfloor$. As before, for each $s \in [m]$, there is $i_s \in [n]$ such that $\|x_{i_s} - z_s\| \leq \varepsilon$. We let $i_0 = i$ and $i_m = j$. The latter is possible since $\|z_m - x_j\| \leq \varepsilon$.

We have

$$d(i_s, i_{s'}) = \|x_{i_s} - x_{i_{s'}}\| \leq \|z_s - z_{s'}\| + 2\varepsilon = \varepsilon|s - s'| + 2\varepsilon, \quad (12)$$

so that, under \mathcal{A}_2 ,

$$d(i_s, i_{s'}) \leq r/4 \quad \text{when} \quad |s - s'| \leq h := \lfloor r/4\varepsilon - 2 \rfloor,$$

implying that $\hat{d}(i_s, i_{s'}) \leq 2r$ when $|s - s'| \leq h$. Thus, by the triangle inequality, under \mathcal{A}_2 ,

$$\hat{d}(i, j) = \sum_{k=0}^{\lfloor m/h \rfloor - 1} \hat{d}(i_{kh}, i_{(k+1)h}) + \hat{d}(i_{\lfloor m/h \rfloor h}, j) \leq \lfloor m/h \rfloor 2r + \hat{d}(i_{\lfloor m/h \rfloor h}, j).$$

By the triangle inequality,

$$d(i_{\lfloor m/h \rfloor h}, j) \leq \|z_{\lfloor m/h \rfloor h} - x_j\| + \varepsilon = |d_{ij} - \lfloor m/h \rfloor h\varepsilon| + \varepsilon,$$

and it is not hard to verify that $0 \leq d_{ij} - \lfloor m/h \rfloor h\varepsilon \leq (h + 1)\varepsilon$, so that $d(i_{\lfloor m/h \rfloor h}, j) \leq (h + 2)\varepsilon \leq r/4$, implying under \mathcal{A}_2 that $\hat{d}(i_{\lfloor m/h \rfloor h}, j) \leq 2r$. Hence, under \mathcal{A}_2 ,

$$\hat{d}(i, j) \leq \lfloor m/h \rfloor 2r + 2r \leq ((d_{ij}/\varepsilon)/(r/4\varepsilon - 3) + 1)2r \leq 9d_{ij} + 2r,$$

when ε/r is small enough.

We have thus established Claim 1.

Claim 1, of course, falls quite short of what is stated in the theorem, but we use it in the remainder of the proof. That said, the claim takes care of all pairs (i, j) such that $d_{ij} \leq 2r$. Thus, for the remainder of the proof, we only need focus on i, j such that $d_{ij} > 2r$. Define m and z_0, \dots, z_m as before, and also the corresponding $i_s \in [n]$.

As before, (12) implies that

$$d(i_s, i_{s'}) \leq r \quad \text{when} \quad |s - s'| \leq h := \lfloor r/\varepsilon - 2 \rfloor,$$

and in particular

$$d(i_s, j) \leq r \quad \text{when} \quad s \geq m - h.$$

(Note that we changed the definition of h .) Similarly, we have

$$d(i_s, j) = \|x_{i_s} - x_j\| \geq \|z_s - z_m\| - \varepsilon = \varepsilon|s - s'| - \varepsilon,$$

so that

$$d(i_s, j) > r \quad \text{when} \quad s \leq m - h + 2.$$

For each $0 \leq s \leq m$, define the random variable H_s by

$$H_s = \max \{0 \leq k \leq h \wedge (m - s) : W(i_s, i_{s+k}) = 1\}.$$

This is a maximum since we have set $W(k, k) = 1$ for all $k \in [n]$. Note that H_0, \dots, H_m are jointly independent random variables with support included in $\{0, \dots, h\}$. Set $S_0 = 0$, and for $t \geq 1$, define recursively $S_t = S_{t-1} + H_{S_{t-1}}$. Importantly, if $H_{S_t} = 0$, then $S_{t'} = S_t$ for all $t' \geq t$. Based on $\{S_t\}$, define

$$T = \inf\{t : S_t > m - h\},$$

with the convention that $\inf \emptyset = \infty$.

Our first objective is to bound T in probability. Given $t \geq 1$, we have

$$\begin{aligned} \mathbb{P}(T > t) &= \mathbb{P}(S_t \leq m - h) \\ &= \mathbb{P}(S_t \leq m - h, H_{S_{t-1}} = 0) + \mathbb{P}(S_t \leq m - h, H_{S_{t-1}} > 0). \end{aligned}$$

On the one hand,

$$\begin{aligned} \mathbb{P}(S_t \leq m - h, H_{S_{t-1}} = 0) &= \mathbb{P}(S_{t-1} \leq m - h, H_{S_{t-1}} = 0) \\ &\leq \mathbb{P}(\exists 0 \leq s \leq m - h : H_s = 0) \\ &\leq m \max_{0 \leq s \leq m-h} \mathbb{P}(H_s = 0). \end{aligned}$$

On the other hand, we note that, when $H_{S_{t-1}} > 0$, we necessarily have $S_t \geq t$, so that

$$\mathbb{P}(S_t \leq m - h, H_{S_{t-1}} > 0) \leq \mathbb{P}(t \leq S_t \leq m - h).$$

Thus,

$$\mathbb{P}(T > t) \leq m \max_{s \in [m-h]} \mathbb{P}(H_s = 0) + \mathbb{P}(t \leq S_t \leq m - h). \quad (13)$$

In what follows, we bound each term in the right hand side of (13).

Claim 2.1. For any $s \in [m - h]$,

$$\mathbb{P}(H_s \leq a) \leq \mathbb{P}(\bar{H} \leq a), \quad a \geq 0,$$

where \bar{H} is a random variable supported on $\{0, 1, \dots, h\}$ with distribution function

$$\mathbb{P}(\bar{H} \leq a) = \prod_{k=0}^{h-a-1} (1 - \phi(\epsilon(h-k) + 2\epsilon)), \quad 0 \leq a \leq h - 1.$$

Indeed, by independence,

$$\begin{aligned} \mathbb{P}(H_s \leq a) &= \mathbb{P}(W(i_s, i_{s+a+1}) = \dots = W(i_s, i_{s+h}) = 0) \\ &= \prod_{k=0}^{h-a-1} (1 - p(i_s, i_{s+h-k})), \end{aligned}$$

and by the fact that ϕ is non-increasing and (12),

$$p(i_s, i_{s'}) \geq \phi(\epsilon|s - s'| + 2\epsilon).$$

Claim 2.2. We have

$$\mathbb{P}(\bar{H} \leq a) \leq \exp(-c_3(\epsilon/r)^\alpha (r/\epsilon - 3 - a)^{1+\alpha}), \quad 0 \leq a \leq h - 1. \quad (14)$$

For any $0 \leq a \leq h - 1$, we have

$$\mathbb{P}(\bar{H} \leq a) \leq \exp\left(-\sum_{k=0}^{h-a-1} \phi(\epsilon(h-k) + 2\epsilon)\right).$$

Since ϕ is non-increasing, and using the lower bound we assume in the statement of the theorem,

$$\begin{aligned} \sum_{k=0}^{h-a-1} \phi(\epsilon(h-k) + 2\epsilon) &\geq \sum_{k=0}^{h-a-1} \int_{h-k}^{h-k+1} \phi(\epsilon y + 2\epsilon) dy \\ &= \int_{a+1}^{h+1} \phi(\epsilon y + 2\epsilon) dy \\ &\geq \int_{a+1}^{h+1} C_0 \left(1 - \frac{\epsilon y + 2\epsilon}{r}\right)_+^\alpha dy \\ &= C_0(\epsilon/r)^\alpha \frac{1}{1+\alpha} (r/\epsilon - 3 - a)^{1+\alpha}. \end{aligned}$$

This proves (14).

With Claims 2.1 and 2.2, the first term on the right-hand side of (13) is bounded by

$$m \mathbb{P}(\bar{H} = 0) \leq m \exp(-c_3(\epsilon/r)^\alpha (r/\epsilon - 3)^{1+\alpha}) \leq m \exp(-c_4 r/\epsilon),$$

when r/ϵ is large enough.

Claim 2.3. For any $1 \leq t \leq m - h$, and any $t \leq a \leq m - h$, we have

$$\mathbb{P}(t \leq S_t \leq a) \leq \mathbb{P}(\bar{H}_0 + \dots + \bar{H}_{t-1} \leq a),$$

where $\{\bar{H}_t\}$ are iid copies of \bar{H} .

First, if $t = 1$, we have

$$\mathbb{P}(1 \leq S_1 \leq a) = \mathbb{P}(1 \leq H_0 \leq a) \leq \mathbb{P}(H_0 \leq a) \leq \mathbb{P}(\bar{H}_0 \leq a).$$

Next, fix $t \geq 2$ and suppose that the claim is true at $t - 1$. Since $t \leq S_t \leq a$ implies that $t - 1 \leq S_{t-1} \leq a$, we have

$$\begin{aligned} \mathbb{P}(t \leq S_t \leq a) &= \mathbb{P}(t \leq S_t \leq a, t - 1 \leq S_{t-1} \leq a) \\ &= \mathbb{P}(t \leq S_{t-1} + H_{S_{t-1}} \leq a, t - 1 \leq S_{t-1} \leq a) \\ &= \sum_{k=t-1}^a \mathbb{P}(t \leq k + H_k \leq a, S_{t-1} = k) \\ &\leq \sum_{k=t-1}^a \mathbb{P}(H_k \leq a - k) \mathbb{P}(S_{t-1} = k) \\ &\leq \sum_{k=t-1}^a \mathbb{P}(\bar{H}_k \leq a - k) \mathbb{P}(S_{t-1} = k) \\ &\leq \sum_{k=t-1}^a \mathbb{P}(\bar{H}_{t-1} \leq a - k) \mathbb{P}(S_{t-1} = k), \end{aligned}$$

where we used the fact that $\{S_{t-1} = k\}$ is independent of H_k , and also the fact that the \bar{H}_t are iid. We may assume, and we do so, that the $\{\bar{H}_t\}$ are defined on the same probability space as the $\{S_t\}$, and are independent of them. Then,

$$\begin{aligned}
& \sum_{k=t-1}^a \mathbb{P}(\bar{H}_{t-1} \leq a - k) \mathbb{P}(S_{t-1} = k) \\
&= \mathbb{P}(S_{t-1} + \bar{H}_{t-1} \leq a, t-1 \leq S_{t-1} \leq a) \\
&= \sum_{k=0}^h \mathbb{P}(S_{t-1} + k \leq a, t-1 \leq S_{t-1} \leq a, \bar{H}_{t-1} = k) \\
&= \sum_{k=0}^h \mathbb{P}(t-1 \leq S_{t-1} \leq a - k) \mathbb{P}(\bar{H}_{t-1} = k) \\
&\leq \sum_{k=0}^h \mathbb{P}(\bar{H}_0 + \cdots + \bar{H}_{t-2} \leq a - k) \mathbb{P}(\bar{H}_{t-1} = k) \\
&= \mathbb{P}(\bar{H}_0 + \cdots + \bar{H}_{t-1} \leq a),
\end{aligned}$$

where the inequality comes from the recursion hypothesis. Thus the recursion proceeds, and the claim is proved.

Claim 2.4. *There exists a constant $b_0 > 0$ and a constant $c_5 > 0$ such that for any $1 \leq t \leq m - h$ and any $a \geq \nu b_0$, we have*

$$\mathbb{P}(\bar{H}_0 + \cdots + \bar{H}_{t-1} \leq th - ta) \leq \exp(-c_5 t \nu^{-1} a),$$

where $\nu := (r/\varepsilon)^{\alpha/(\alpha+1)}$.

Define $\bar{U} = h - \bar{H}$ and $\bar{U}_t = h - \bar{H}_t$. We have

$$\mathbb{P}(\bar{H}_0 + \cdots + \bar{H}_{t-1} \leq th - ta) = \mathbb{P}(\bar{U}_0 + \cdots + \bar{U}_{t-1} \geq ta)$$

For any $1 \leq u \leq h$, we have

$$\mathbb{P}(\bar{U} \geq u) = \mathbb{P}(\bar{H} \leq h - u) \leq \exp(-c_3(\varepsilon/r)^\alpha(u-1)^{1+\alpha}),$$

using (14) and the fact that $h \leq r/\varepsilon - 2$. Hence, \bar{U} is stochastically bounded by $1 + \nu Y$, where Y is a random variable with distribution

$$\mathbb{P}(Y \geq y) = \exp(-c_3 y^{1+\alpha}), \quad y \geq 0.$$

Let $\{Y_t : t \geq 1\}$ be iid with distribution that of Y . By construction, $\bar{U}_0 + \cdots + \bar{U}_{t-1}$ is stochastically bounded $t + \nu(Y_0 + \cdots + Y_{t-1})$, implying that

$$\mathbb{P}(\bar{H}_0 + \cdots + \bar{H}_{t-1} \leq th - ta) \leq \mathbb{P}(Y_0 + \cdots + Y_{t-1} \geq t(a-1)/\nu).$$

By Chernoff's bound,

$$\mathbb{P}(Y_1 + \cdots + Y_{t-1} > tb) \leq \exp(-t\zeta(b)), \quad \forall b \geq 0,$$

where ζ is the rate function of Y . In particular, there is $c > 0$ such that $\zeta(b) \geq cb$ for all $b \geq b_0 := \mathbb{E}(Y) + 1$. (Note that b_0 is just another constant.) Thus, for a such that $(a - 1)/\nu \geq b_0$,

$$\mathbb{P}(Y_0 + \dots + Y_{t-1} \geq t(a - 1)/\nu) \leq \exp(-tc(a - 1)/\nu),$$

which together with $a - 1 \geq a/2$ (when r/ε is large) proves the claim.

With Claims 2.3 and 2.4, by choosing $a = h - (m - h)/t$, we obtain that the second term on the right-hand side of (13) is bounded by

$$\exp(-c_5 t \nu^{-1} a),$$

whenever $a \geq \nu b_0$, which happens when $t \geq (m - h)/(h - \nu b_0)$. This is the case when $t \geq t^* := 2d_{ij}/r$, which may be seen using the fact that $m - h \leq m \leq d_{ij}/\varepsilon$, and that $h - \nu b_0 \geq r/\varepsilon - 3(r/\varepsilon)^{\alpha/(\alpha+1)}b_0 \geq r/2\varepsilon$ when r/ε is large enough. In fact, when $t \geq t^*$, the corresponding a satisfies $a \geq r/\varepsilon - 3 - (d_{ij}/\varepsilon)/(2d_{ij}/r) \geq r/2\varepsilon$ when r/ε is large enough, so that the right-hand side of (13) is bounded by

$$\exp(-c_5 t \nu^{-1} (r/2\varepsilon)) \leq \exp(-tc_6 (r/\varepsilon)^{\frac{1}{1+\alpha}}),$$

when $t \geq t^*$.

All combined, we have

$$\mathbb{P}(T > t) \leq n \exp(-c_4 r/\varepsilon) + \exp(-tc_6 (r/\varepsilon)^{\frac{1}{1+\alpha}}), \quad \forall t \geq t^*,$$

using the fact that $m \leq n$. In particular, this implies that

$$\mathbb{P}(T = \infty) \leq n \exp(-c_4 r/\varepsilon).$$

Let $t^\circ = t^* + (5/c_5 b_0) \log n$. Let \mathcal{C} be the probability event defined by

$$\mathcal{C} = \{S_T > Th - t^\circ \nu b_0 \text{ and } T \leq t^\circ\}.$$

We have

$$\begin{aligned} \mathbb{P}(\mathcal{C}^c) &= \mathbb{P}(S_T \leq Th - t^\circ \nu b_0, T \leq t^\circ) + \mathbb{P}(T > t^\circ) \\ &\leq \mathbb{P}(t^\circ \leq S_{t^\circ} \leq t^\circ h - t^\circ \nu b_0) + \mathbb{P}(T > t^\circ) \\ &\leq \exp(-c_5 t^\circ b_0) + n \exp(-c_4 r/\varepsilon) + \exp(-t^\circ c_6 (r/\varepsilon)^{\frac{1}{1+\alpha}}), \end{aligned}$$

using in the second line the fact that $S_t \geq t$ when $T \leq t$, and that $S_t - th - (S_{t-1} - (t-1)h) = H_{S_{t-1}} - h \leq 0$ for all t . The application of Claim 2.4 here is valid when $t^\circ \leq m - h$. This is the case eventually as $t^\circ = 2d_{ij}/r + (5/c_5 b_0) \log n$ and $m - h \geq d_{ij}/\varepsilon - r/\varepsilon + 1 \geq d_{ij}/3\varepsilon$, with

$$\frac{2d_{ij}/r + (5/c_5 b_0) \log n}{d_{ij}/3\varepsilon} \leq \varepsilon/r + c_7 b_0 (\varepsilon/r) \log n \leq 1,$$

using the fact that $d_{ij} \geq r$ and our lower bound on r/ε . When C_1 is large enough, \mathcal{C} holds with probability at least $1 - 1/n^4$, eventually.

A joint control on V_T and T is useful because of the following. Under $\{T < \infty\}$, $(i, i_{S_1}, \dots, i_{S_T})$ forms a path in the graph, so that $\hat{d}(i, i_{S_T}) \leq Tr$. We also have

$$d_{ij} \geq \sum_{t=1}^T \|z_{S_t} - z_{S_{t-1}}\| = \sum_{t=1}^T \varepsilon(S_t - S_{t-1}) = \varepsilon(S_T - S_0) = \varepsilon S_T.$$

Thus, under \mathcal{C} ,

$$d_{ij} \geq \varepsilon S_T \geq \varepsilon hT - \varepsilon b_0 \nu t^\circ \geq rT - (3 + b_0 \nu) \varepsilon t^\circ \geq rT - 2b_0 \nu \varepsilon t^\circ,$$

eventually, using the fact that $h \geq r/\varepsilon - 3$. Furthermore,

$$\begin{aligned} \nu \varepsilon t^\circ &= (r/\varepsilon)^{\alpha/(1+\alpha)} \varepsilon (2d_{ij}/r + (5/c_5 b_0) \log n) \\ &= 2(\varepsilon/r)^{1/(1+\alpha)} d_{ij} + [(5/c_5 b_0)(\varepsilon/r)^{1/(1+\alpha)} (\log n)] r \\ &\leq 2(\varepsilon/r)^{1/(1+\alpha)} d_{ij} + (5/c_5 b_0) C_1^{-1/(1+\alpha)} r \\ &\leq 2(\varepsilon/r)^{1/(1+\alpha)} d_{ij} + r, \end{aligned}$$

when C_1 is large enough. Thus, under \mathcal{C} ,

$$\hat{d}(i, i_{S_T}) \leq Tr \leq d_{ij} + 2b_0 \nu \varepsilon t^\circ \leq d_{ij} + c_8 ((\varepsilon/r)^{1/(1+\alpha)} d_{ij} + r).$$

And since, as in (12),

$$d(i_{S_T}, j) \leq \varepsilon(m - S_T) + \varepsilon \leq \varepsilon(h - 3) + \varepsilon \leq \varepsilon(r/\varepsilon - 5) + \varepsilon \leq r.$$

Using the union bound over all pairs suitable (i, j) , we have established the following.

Claim 2. *By choosing C_1 large enough, with probability at least $1 - 1/n^2$, for all $i, j \in [n]$ such that $d_{ij} > 2r$, there is $k \in [n]$ such that $\hat{d}_{ik} \leq d_{ij} + c_8 ((\varepsilon/r)^{1/(1+\alpha)} d_{ij} + r)$ and $d_{jk} \leq r$. (We denote this event by \mathcal{A}_3 .)*

Assume that $\mathcal{A}_1 \cap \mathcal{A}_3$ holds, which happens with probability at least $1 - 2/n^2$ based on the two claims that we have established above. In that case, let $i, j, k \in [n]$ be as in the last claim. Because \mathcal{A}_1 holds, $d_{jk} \leq r$ implies that $\hat{d}_{jk} \leq 11r$. Thus, with the union bound,

$$\hat{d}_{ij} \leq \hat{d}_{ik} + \hat{d}_{jk} \leq d_{ij} + c_8 ((\varepsilon/r)^{1/(1+\alpha)} d_{ij} + r) + 11r,$$

so that

$$\hat{d}_{ij} - d_{ij} \leq c_9 ((\varepsilon/r)^{1/(1+\alpha)} d_{ij} + r).$$

This proves that (7) holds for all $i, j \in [n]$ such that $d_{ij} > 2r$.

References

- [1] ABBE, E. (2017). Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research* **18** 6446–6531. [MR3827065](#)
- [2] AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *The Journal of Machine Learning Research* **9** 1981–2014.
- [3] ALAMGIR, M. and VON LUXBURG, U. (2012). Shortest path distance in random k-nearest neighbor graphs. In *International Conference on Machine Learning (ICML)* 1251–1258.
- [4] ANANDKUMAR, A., GE, R., HSU, D. and KAKADE, S. M. (2013). A tensor spectral approach to learning mixed membership community models. *arXiv preprint arXiv:1302.2684*. [MR3231594](#)
- [5] ARIAS-CASTRO, E., JAVANMARD, A. and PELLETIER, B. (2020). Perturbation bounds for procrustes, classical scaling, and trilateration, with applications to manifold learning. *Journal of Machine Learning Research* **21** 15–1. [MR4071198](#)
- [6] BERNSTEIN, M., DE SILVA, V., LANGFORD, J. C. and TENENBAUM, J. B. (2000). Graph approximations to geodesics on embedded manifolds Technical Report, Technical report, Department of Psychology, Stanford University.
- [7] BORG, I. and GROENEN, P. (2013). *Modern multidimensional scaling: theory and applications*. Springer Science & Business Media. [MR2158691](#)
- [8] BRADLEY, R. A. and TERRY, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* **39** 324–345. [MR0070925](#)
- [9] CHAZAL, F., GUIBAS, L. J., OUDOT, S. Y. and SKRABA, P. (2013). Persistence-based clustering in riemannian manifolds. *Journal of the ACM* **60** 41. [MR3144911](#)
- [10] CHEN, L. and BUJA, A. (2009). Local multidimensional scaling for nonlinear dimension reduction, graph drawing, and proximity analysis. *Journal of the American Statistical Association* **104** 209–219. [MR2504374](#)
- [11] CHEN, Y. and XU, J. (2016). Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *The Journal of Machine Learning Research* **17** 882–938. [MR3491121](#)
- [12] CUCURINGU, M. (2013). ASAP: An eigenvector synchronization algorithm for the graph realization problem. In *Distance Geometry* 177–195. Springer. [MR3051954](#)
- [13] CUCURINGU, M., LIPMAN, Y. and SINGER, A. (2012). Sensor network localization by eigenvector synchronization over the Euclidean group. *ACM Transactions on Sensor Networks* **8** 19.
- [14] DE LEEUW, J. and MAIR, P. (2009). Multidimensional Scaling Using Majorization: SMACOF in R. *Journal of Statistical Software* **31**.

- [15] FLIGNER, M. A. and VERDUCCI, J. S. (1993). *Probability models and statistical analyses for ranking data* **80**. Springer. [MR1237197](#)
- [16] HAJEK, B., WU, Y. and XU, J. (2016). Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory* **62** 2788–2797. [MR3493879](#)
- [17] HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association* **97** 1090–1098. [MR1951262](#)
- [18] HOLLAND, P. W., BLACKMOND-LASKEY, K. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- [19] KRUSKAL, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* **29** 1–27. [MR0169712](#)
- [20] KRUSKAL, J. B. and SEERY, J. B. (1980). Designing network diagrams. In *Conference on Social Graphics* 22–50.
- [21] LEI, J. and RINALDO, A. (2015). Consistency of spectral clustering in stochastic block models. *The Annals of Statistics* **43** 215–237. [MR3285605](#)
- [22] LIBEN-NOWELL, D. and KLEINBERG, J. (2003). The link-prediction problem for social networks. In *Conference on Information and Knowledge Management (CIKM'03)* 556–559.
- [23] LIBEN-NOWELL, D. and KLEINBERG, J. (2007). The link-prediction problem for social networks. *Journal of the Association for Information Science and Technology* **58** 1019–1031.
- [24] LUCE, R. D. (2012). *Individual choice behavior: A theoretical analysis*. Courier Corporation. [MR0108411](#)
- [25] MARDEN, J. I. (1996). *Analyzing and modeling rank data*. CRC Press. [MR1346107](#)
- [26] MOSSEL, E., NEEMAN, J. and SLY, A. (2015). Consistency thresholds for the planted bisection model. In *ACM Symposium on the Theory of Computing* 69–75. [MR3388184](#)
- [27] PANANJADY, A., MAO, C., MUTHUKUMAR, V., WAINWRIGHT, M. J. and COURTADE, T. A. (2020). Worst-case vs average-case design for estimation from fixed pairwise comparisons. *The Annals of Statistics* **48** 1072–1097. [MR4102688](#)
- [28] PARTHASARATHY, S., SIVAKOFF, D., TIAN, M. and WANG, Y. (2017). A quest to unravel the metric structure behind perturbed networks. *arXiv preprint arXiv:1703.05475*. [MR3685725](#)
- [29] PENROSE, M. et al. (2003). *Random geometric graphs*. Oxford University Press. [MR1986198](#)
- [30] ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics* **39** 1878–1915. [MR2893856](#)
- [31] SARKAR, P., CHAKRABARTI, D. and MOORE, A. W. (2010). Theoretical justification of popular link prediction heuristics. In *International Conference on Learning Theory*.
- [32] SCHEINERMAN, E. R. and TUCKER, K. (2010). Modeling graphs using dot

- product representations. *Computational Statistics* **25** 1–16. [MR2586721](#)
- [33] SHAH, N. B., BALAKRISHNAN, S., GUNTUBOYINA, A. and WAINWRIGHT, M. J. (2016). Stochastically transitive models for pairwise comparisons: statistical and computational issues. *IEEE Transactions on Information Theory* **63** 934–959. [MR3604649](#)
- [34] SHANG, Y. and RUML, W. (2004). Improved MDS-based localization. In *INFOCOM 2004. Twenty-Third Annual Joint Conference of the IEEE Computer and Communications Societies* **4** 2640–2651.
- [35] SHANG, Y., RUML, W., ZHANG, Y. and FROMHERZ, M. P. (2003). Localization from mere connectivity. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing* 201–212.
- [36] SHEPARD, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. I. *Psychometrika* **27** 125–140. [MR0140376](#)
- [37] SHEPARD, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. II. *Psychometrika* **27** 219–246. [MR0173342](#)
- [38] SUSSMAN, D. L., TANG, M., FISHKIND, D. E. and PRIEBE, C. E. (2012). A consistent adjacency spectral embedding for stochastic block-model graphs. *Journal of the American Statistical Association* **107** 1119–1128. [MR3010899](#)
- [39] SUSSMAN, D. L., TANG, M. and PRIEBE, C. E. (2014). Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36** 48–57.
- [40] TANG, M., SUSSMAN, D. L. and PRIEBE, C. E. (2013). Universally consistent vertex classification for latent positions graphs. *The Annals of Statistics* **41** 1406–1430. [MR3113816](#)
- [41] TENENBAUM, J. B., DE SILVA, V. and LANGFORD, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science* **290** 2319–2323.
- [42] TORGERSON, W. S. (1952). Multidimensional scaling: I. Theory and method. *Psychometrika* **17** 401–419. [MR0054219](#)
- [43] VON LUXBURG, U. and ALAMGIR, M. (2013). Density estimation from unweighted k-nearest neighbor graphs: a roadmap. In *Advances in Neural Information Processing Systems* 225–233.
- [44] YANG, J. and LESKOVEC, J. (2013). Overlapping community detection at scale: a nonnegative matrix factorization approach. In *ACM International Conference on Web Search and Data Mining* 587–596.
- [45] YOUNG, F. W. and HAMER, R. M. E. (1987). *Multidimensional scaling: history, theory, and applications*. Lawrence Erlbaum Associates, Inc.
- [46] YOUNG, S. J. and SCHEINERMAN, E. R. (2007). Random dot product graph models for social networks. In *Algorithms and Models for the Web-Graph* 138–149. Springer. [MR2504912](#)