
Finite Sample Prediction and Recovery Bounds for Ordinal Embedding

Lalit Jain
University of Michigan
Ann Arbor, MI 48109
lalitj@umich.edu

Kevin Jamieson
University of California, Berkeley
Berkeley, CA 94720
kjamieson@berkeley.edu

Robert Nowak
University of Wisconsin
Madison, WI 53706
rdnowak@wisc.edu

Abstract

The goal of ordinal embedding is to represent items as points in a low-dimensional Euclidean space given a set of constraints like “item i is closer to item j than item k ”. Ordinal constraints like this often come from human judgments. The classic approach to solving this problem is known as *non-metric multidimensional scaling*. To account for errors and variation in judgments, we consider the noisy situation in which the given constraints are independently corrupted by reversing the correct constraint with some probability. The ordinal embedding problem has been studied for decades, but most past work pays little attention to the question of whether accurate embedding is possible, apart from empirical studies. This paper shows that under a generative data model it is possible to learn the correct embedding from noisy distance comparisons. In establishing this fundamental result, the paper makes several new contributions. First, we derive prediction error bounds for embedding from noisy distance comparisons by exploiting the fact that the rank of a distance matrix of points in \mathbb{R}^d is at most $d + 2$. These bounds characterize how well a learned embedding predicts new comparative judgments. Second, we show that the underlying embedding can be recovered by solving a simple convex optimization. This result is highly non-trivial since we show that the linear map corresponding to distance comparisons is non-invertible, but there exists a nonlinear map that is invertible. Third, two new algorithms for ordinal embedding are proposed and evaluated in experiments.

1 Ordinal Embedding

Ordinal embedding aims to represent items as points in \mathbb{R}^d so that the distances between items agree as well as possible with a given set of ordinal comparisons such as item i is closer to item j than to item k . In other words, the goal is to find a geometric representation of data that is faithful to comparative similarity judgments. This problem has been studied and applied for more than 50 years, dating back to the classic *non-metric multidimensional scaling* (NMDS) [1, 2] approach, and it is widely used to gauge and visualize how people perceive similarities.

Despite the widespread application of NMDS and recent algorithmic developments [3, 4, 5, 6, 7], the fundamental question of whether an embedding can be learned from noisy distance/similarity comparisons had not been answered. This paper shows that if the data are generated according to a known probabilistic model, then accurate recovery of the underlying embedding is possible by solving a simple convex optimization, settling this long-standing open question. In the process of answering this question, the paper also characterizes how well a learned embedding predicts new distance comparisons and presents two new computationally efficient algorithms for solving the optimization problem.

1.1 Related Work

The classic approach to ordinal embedding is NMDS [1, 2]. Recently, several authors have proposed new approaches based on more modern techniques. *Generalized NMDS* [3] and *Stochastic Triplet Embedding (STE)* [6] employ hinge or logistic loss measures and convex relaxations of the low-dimensionality (i.e., rank) constraint based on the nuclear norm. These works are most closely related to the theory and methods in this paper. The *Linear partial order embedding (LPOE)* method is similar, but starts with a known Euclidean embedding and learns a kernel/metric in this space based distance comparison data [7]. The *Crowd Kernel* [4] and *t-STE* [6] propose alternative non-convex loss measures based on probabilistic generative models. The main contributions in these papers are new optimization methods and experimental studies, but did not address the fundamental question of whether an embedding can be recovered under an assumed generative model. Other recent work has looked at the asymptotics of ordinal embedding, showing that embeddings can be learned as the number of items grows and the items densely populate the embedding space [8, 9, 10]. In contrast, this paper focuses on the practical setting involving a finite set items. Finally, it is known that at least $2dn \log n$ distance comparisons are necessary to learn an embedding of n points in \mathbb{R}^d [5].

1.2 Ordinal Embedding from Noisy Data

Consider n points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$. Let $\mathbf{X} = [\mathbf{x}_1 \cdots \mathbf{x}_n] \in \mathbb{R}^{d \times n}$. The *Euclidean distance matrix* \mathbf{D}^* is defined to have elements $D_{ij}^* = \|\mathbf{x}_i - \mathbf{x}_j\|_2^2$. Ordinal embedding is the problem of recovering \mathbf{X} given ordinal constraints on distances. This paper focuses on “triplet” constraints of the form $D_{ij}^* < D_{ik}^*$, where $1 \leq i \neq j \neq k \leq n$. Furthermore, we only observe noisy indications of these constraints, as follows. Each triplet $t = (i, j, k)$ has an associated probability p_t satisfying

$$p_t > 1/2 \iff \|\mathbf{x}_i - \mathbf{x}_j\|^2 < \|\mathbf{x}_i - \mathbf{x}_k\|^2.$$

Let \mathcal{S} denote a collection of triplets drawn independently and uniformly at random. And for each $t \in \mathcal{S}$ we observe an independent random variable $y_t = -1$ with probability p_t , and $y_t = 1$ otherwise. The goal is to recover the embedding \mathbf{X} from these data. Exact recovery of \mathbf{D}^* from such data requires a known *link* between p_t and \mathbf{D}^* . To this end, our main focus is the following problem.

Ordinal Embedding from Noisy Data

Consider n points $\mathbf{x}_1, \mathbf{x}_2 \cdots, \mathbf{x}_n$ in d -dimensional Euclidean space. Let \mathcal{S} denote a collection of triplets and for each $t \in \mathcal{S}$ observe an independent random variable

$$y_t = \begin{cases} -1 & w.p. f(D_{ij}^* - D_{ik}^*) \\ 1 & w.p. 1 - f(D_{ij}^* - D_{ik}^*) \end{cases}.$$

where the link function $f : \mathbb{R} \rightarrow [0, 1]$ is known. Estimate \mathbf{X} from \mathcal{S} , $\{y_t\}$, and f .

For example, if f is the logistic function, then for triplet $t = (i, j, k)$

$$p_t = \mathbb{P}(y_t = -1) = f(D_{ij}^* - D_{ik}^*) = \frac{1}{1 + \exp(D_{ij}^* - D_{ik}^*)}, \quad (1)$$

then $D_{ij}^* - D_{ik}^* = \log\left(\frac{1-p_t}{p_t}\right)$. However, we stress that we only require the existence of a link function for exact recovery of \mathbf{D}^* . Indeed, if one just wishes to *predict* the answers to unobserved triplets, then the results of Section 2 hold for arbitrary p_t probabilities. Aspects of the statistical analysis are related to one-bit matrix completion and rank aggregation [11, 12, 13]. However, we use novel methods for the recovery of the embedding based on geometric properties of Euclidean distance matrices.

1.3 Organization of Paper

This paper takes the following approach to ordinal embedding.

1. Our samples are assumed to be independently generated according to a probabilistic model based on an underlying low-rank distance matrix. We use relatively standard statistically learning theory

techniques to analyze the minimizer of a bounded, Lipschitz loss with a nuclear norm constraint, and show that an embedding can be learned from the data that predicts nearly as well as the true embedding with $O(dn \log n)$ samples (Theorem 1).

2. Next, assuming the form of the probabilistic generative model is known (e.g., logistic), we show that if the learned embedding is a good predictor of the ordinal comparisons, then it must also be a good estimator of the true differences of distances between the embedding points (Theorem 2). This result hinges on the fact that the (linear) observation model acts approximately like an isometry on differences of distances.

3. While the true differences of distances can be estimated, the observation process is “blind” to the mean distance between embedding points. Despite this, we show that the mean is determined by the differences of distances, due to the special properties of Euclidean distance matrices. Specifically, the second eigenvalue of the “mean-centered” distance matrix (well-estimated by the data from the estimate of the differences of distances, Theorem 3) is proportional to the mean distance (Theorem 4). This allows us to show that the minimizer of the loss with a nuclear norm constraint indeed recovers an accurate estimate of the underlying true distance matrix.

1.4 Notation and Assumptions

We will use (D^*, G^*) to denote the distance and Gram matrices of the latent embedding, and (D, G) to denote an arbitrary distance matrix and its corresponding Gram matrix. The observations $\{y_t\}$ carry information about D^* , but distance matrices are invariant to rotation and translation, and therefore it may only be possible to recover X up to a rigid transformation. Without loss of generality, we assume the points $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^d$ are centered at the origin (i.e., $\sum_{i=1}^n \mathbf{x}_i = \mathbf{0}$).

Define the *centering matrix* $V := I - \frac{1}{n}\mathbf{1}\mathbf{1}^T$. If X is centered, $XV = X$. Note that D^* is determined by the Gram matrix $G^* = X^T X$. In addition, X can be determined from G^* up to a unitary transformation. Note that if X is centered, the Gram matrix is “centered” so that $VG^*V = G^*$. It will be convenient in the paper to work with both the distance and Gram matrix representations, and the following identities will be useful to keep in mind. For any distance matrix D and its centered Gram matrix G

$$G = -\frac{1}{2}VDV, \quad (2)$$

$$D = \text{diag}(G)\mathbf{1}^T - 2G + \mathbf{1}\text{diag}(G)^T, \quad (3)$$

where $\text{diag}(G)$ is the column vector composed of the diagonal of G . In particular this establishes a bijection between centered Gram matrices and distance matrices. We refer the reader to [14] for an insightful and thorough treatment of the properties of distance matrices. We also define the set of all unique triplets

$$\mathcal{T} := \{(i, j, k) : 1 \leq i \neq j \neq k \leq n, j < k\}.$$

Assumption 1. *The observed triplets in S are drawn independently and uniformly from \mathcal{T} .*

2 Prediction Error Bounds

For $t \in \mathcal{T}$ with $t = (i, j, k)$ we define \mathcal{L}_t to be the linear operator satisfying $\mathcal{L}_t(X^T X) = \|\mathbf{x}_i - \mathbf{x}_j\|^2 - \|\mathbf{x}_i - \mathbf{x}_k\|^2$ for all $t \in \mathcal{T}$. In general, for any Gram matrix G

$$\mathcal{L}_t(G) := G_{jj} - 2G_{ij} - G_{kk} + 2G_{ik}.$$

We can naturally view \mathcal{L}_t as a linear operator on \mathbb{S}_+^n , the space of $n \times n$ symmetric positive semidefinite matrices. We can also represent \mathcal{L}_t as a symmetric $n \times n$ matrix that is zero everywhere except on the submatrix corresponding to i, j, k which has the form

$$\begin{bmatrix} 0 & -1 & 1 \\ -1 & 1 & 0 \\ 1 & 0 & -1 \end{bmatrix}$$

and so we will write

$$\mathcal{L}_t(G) := \langle \mathcal{L}_t, G \rangle$$

where $\langle A, B \rangle = \text{vec}(A)^T \text{vec}(B)$ for any compatible matrices A, B . Ordering the elements of \mathcal{T} lexicographically, we arrange all the $\mathcal{L}_t(\mathbf{G})$ together to define the $n \binom{n-1}{2}$ -dimensional vector

$$\mathcal{L}(\mathbf{G}) = [\mathcal{L}_{123}(\mathbf{G}), \mathcal{L}_{124}(\mathbf{G}), \dots, \mathcal{L}_{ijk}(\mathbf{G}), \dots]^T. \quad (4)$$

Let $\ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle)$ denote a loss function. For example we can consider the 0 – 1 loss $\ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) = \mathbb{1}_{\{\text{sign}\{y_t \langle \mathcal{L}_t, \mathbf{G} \rangle\} \neq 1\}}$, the hinge-loss $\ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) = \max\{0, 1 - y_t \langle \mathcal{L}_t, \mathbf{G} \rangle\}$, or the logistic loss

$$\ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) = \log(1 + \exp(-y_t \langle \mathcal{L}_t, \mathbf{G} \rangle)). \quad (5)$$

Let $p_t := \mathbb{P}(y_t = -1)$ and take the expectation of the loss with respect to both the uniformly random selection of the triple t and the observation y_t , we have the risk of \mathbf{G}

$$R(\mathbf{G}) := \mathbb{E}[\ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle)] = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} p_t \ell(-\langle \mathcal{L}_t, \mathbf{G} \rangle) + (1 - p_t) \ell(\langle \mathcal{L}_t, \mathbf{G} \rangle).$$

Given a set of observations \mathcal{S} under the model defined in the problem statement, the empirical risk is,

$$\widehat{R}_{\mathcal{S}}(\mathbf{G}) = \frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} \ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) \quad (6)$$

which is an unbiased estimator of the true risk: $\mathbb{E}[\widehat{R}_{\mathcal{S}}(\mathbf{G})] = R(\mathbf{G})$. For any $\mathbf{G} \in \mathbb{S}_+^n$, let $\|\mathbf{G}\|_*$ denote the nuclear norm and $\|\mathbf{G}\|_{\infty} := \max_{ij} |\mathbf{G}_{ij}|$. Define the constraint set

$$\mathcal{G}_{\lambda, \gamma} := \{\mathbf{G} \in \mathbb{S}_+^n : \|\mathbf{G}\|_* \leq \lambda, \|\mathbf{G}\|_{\infty} \leq \gamma\}. \quad (7)$$

We estimate \mathbf{G}^* by $\widehat{\mathbf{G}}$, the solution of the program,

$$\widehat{\mathbf{G}} := \underset{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}}{\text{argmin}} \widehat{R}_{\mathcal{S}}(\mathbf{G}). \quad (8)$$

Since \mathbf{G}^* is positive semidefinite, we expect the diagonal entries of \mathbf{G}^* to bound the off-diagonal entries. So an infinity norm constraint on the diagonal guarantees that the points $\mathbf{x}_1, \dots, \mathbf{x}_n$ corresponding to \mathbf{G}^* live inside a bounded ℓ_2 ball. The ℓ_{∞} constraint in (7) plays two roles: 1) if our loss function is Lipschitz, large magnitude values of $\langle \mathcal{L}_t, \mathbf{G} \rangle$ can lead to large deviations of $\widehat{R}_{\mathcal{S}}(\mathbf{G})$ from $R(\mathbf{G})$; bounding $\|\mathbf{G}\|_{\infty}$ bounds $|\langle \mathcal{L}_t, \mathbf{G} \rangle|$. 2) Later we will define ℓ in terms of the link function f and as the magnitude of $\langle \mathcal{L}_t, \mathbf{G} \rangle$ increases the magnitude of the derivative of the link function f typically becomes very small, making it difficult to “invert”; bounding $\|\mathbf{G}\|_{\infty}$ tends to keep $\langle \mathcal{L}_t, \mathbf{G} \rangle$ within an invertible regime of f .

Theorem 1. Fix λ, γ and assume $\mathbf{G}^* \in \mathcal{G}_{\lambda, \gamma}$. If the loss function $\ell(\cdot)$ is L -Lipschitz (or $|\sup_y \ell(y)| \leq L \max\{1, 12\gamma\}$) then with probability at least $1 - \delta$,

$$R(\widehat{\mathbf{G}}) - R(\mathbf{G}^*) \leq \frac{4L\lambda}{|\mathcal{S}|} \left(\sqrt{\frac{18|\mathcal{S}| \log(n)}{n}} + \frac{\sqrt{3}}{3} \log n \right) + L\gamma \sqrt{\frac{288 \log 2/\delta}{|\mathcal{S}|}}$$

Proof. The proof follows from standard statistical learning theory techniques, see for instance [15]. By the bounded difference inequality, with probability $1 - \delta$

$$\begin{aligned} R(\widehat{\mathbf{G}}) - R(\mathbf{G}^*) &= R(\widehat{\mathbf{G}}) - \widehat{R}_{\mathcal{S}}(\widehat{\mathbf{G}}) + \widehat{R}_{\mathcal{S}}(\widehat{\mathbf{G}}) - \widehat{R}_{\mathcal{S}}(\mathbf{G}^*) + \widehat{R}_{\mathcal{S}}(\mathbf{G}^*) - R(\mathbf{G}^*) \\ &\leq 2 \sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} |\widehat{R}_{\mathcal{S}}(\mathbf{G}) - R(\mathbf{G})| \leq 2\mathbb{E} \left[\sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} |\widehat{R}_{\mathcal{S}}(\mathbf{G}) - R(\mathbf{G})| \right] + \sqrt{\frac{2B^2 \log 2/\delta}{|\mathcal{S}|}} \end{aligned}$$

where $\sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} \ell(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) - \ell(y_{t'} \langle \mathcal{L}_{t'}, \mathbf{G} \rangle) \leq \sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} L |y_t \langle \mathcal{L}_t, \mathbf{G} \rangle - y_{t'} \langle \mathcal{L}_{t'}, \mathbf{G} \rangle| \leq 12L\gamma =: B$ using the facts that \mathcal{L}_t has 6 non-zeros of magnitude 1 and $\|\mathbf{G}\|_{\infty} \leq \gamma$.

Using standard symmetrization and contraction lemmas, we can introduce Rademacher random variables $\epsilon_t \in \{-1, 1\}$ for all $t \in \mathcal{S}$ so that

$$\mathbb{E} \sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} |\widehat{R}_{\mathcal{S}}(\mathbf{G}) - R(\mathbf{G})| \leq \mathbb{E} \sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} \left| \frac{2L}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} \epsilon_t \langle \mathcal{L}_t, \mathbf{G} \rangle \right|.$$

The right hand side is just the Rademacher complexity of $\mathcal{G}_{\lambda, \gamma}$. By definition,

$$\{\mathbf{G} : \|\mathbf{G}\|_* \leq \lambda\} = \lambda \cdot \text{conv}(\{uu^T : |u| = 1\}).$$

where $\text{conv}(U)$ is the convex hull of a set U . Since the Rademacher complexity of a set is the same as the Rademacher complexity of its closed convex hull,

$$\mathbb{E} \sup_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} \left| \sum_{t \in \mathcal{S}} \epsilon_t \langle \mathcal{L}_t, \mathbf{G} \rangle \right| \leq \lambda \mathbb{E} \sup_{|u|=1} \left| \sum_{t \in \mathcal{S}} \epsilon_t \langle \mathcal{L}_t, uu^T \rangle \right| = \lambda \mathbb{E} \sup_{|u|=1} \left| u^T \left(\sum_{t \in \mathcal{S}} \epsilon_t \mathcal{L}_t \right) u \right|$$

which we recognize is just $\lambda \mathbb{E} \|\sum_{t \in \mathcal{S}} \epsilon_t \mathcal{L}_t\|$. By [16, 6.6.1] we can bound the operator norm $\|\sum_{t \in \mathcal{S}} \epsilon_t \mathcal{L}_t\|$ in terms of the variance of $\sum_{t \in \mathcal{S}} \mathcal{L}_t^2$ and the maximal eigenvalue of $\max_t \mathcal{L}_t$. These are computed in Lemma 1 given in the supplemental materials. Combining these results gives,

$$\frac{2L\lambda}{|\mathcal{S}|} \mathbb{E} \|\sum_{t \in \mathcal{S}} \epsilon_t \mathcal{L}_t\| \leq \frac{2L\lambda}{|\mathcal{S}|} \left(\sqrt{\frac{18|\mathcal{S}| \log(n)}{n}} + \frac{\sqrt{3}}{3} \log n \right).$$

□

We remark that if \mathbf{G} is a rank $d < n$ matrix then

$$\|\mathbf{G}\|_* \leq \sqrt{d} \|\mathbf{G}\|_F \leq \sqrt{dn} \|\mathbf{G}\|_\infty$$

so if \mathbf{G}^* is low rank, we really only need a bound on the infinity norm of our constraint set. Under the assumption that \mathbf{G}^* is rank d with $\|\mathbf{G}^*\|_\infty \leq \gamma$ and we set $\lambda = \sqrt{dn}\gamma$, then Theorem 1 implies that for $|\mathcal{S}| > n \log n / 161$

$$R(\widehat{\mathbf{G}}) - R(\mathbf{G}^*) \leq 8L\gamma \sqrt{\frac{18dn \log(n)}{|\mathcal{S}|}} + L\gamma \sqrt{\frac{288 \log 2/\delta}{|\mathcal{S}|}}$$

with probability at least $1 - \delta$. The above display says that $|\mathcal{S}|$ must scale like $dn \log(n)$ which is consistent with known finite sample bounds [5].

3 Maximum Likelihood Embedding

We now turn our attention to recovering metric information about \mathbf{G}^* . Let \mathcal{S} be a collection of triplets sampled uniformly at random with replacement and let $f : \mathbb{R} \rightarrow (0, 1)$ be a known probability function governing the observations. Any link function f induces a natural loss function ℓ_f , namely, the negative log-likelihood of a solution \mathbf{G} given an observation y_t defined as

$$\ell_f(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle) = \mathbb{1}_{y_t = -1} \log\left(\frac{1}{f(\langle \mathcal{L}_t, \mathbf{G} \rangle)}\right) + \mathbb{1}_{y_t = 1} \log\left(\frac{1}{1 - f(\langle \mathcal{L}_t, \mathbf{G} \rangle)}\right)$$

For example, the logistic link function of (1) induces the logistic loss of (5). Recalling that $\mathbb{P}(y_t = -1) = f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle)$ we have

$$\begin{aligned} \mathbb{E}[\ell_f(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle)] &= f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle) \log\left(\frac{1}{f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle)}\right) + (1 - f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle)) \log\left(\frac{1}{1 - f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle)}\right) \\ &= H(f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle)) + KL(f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle) | f(\langle \mathcal{L}_t, \mathbf{G} \rangle)) \end{aligned}$$

where $H(p) = p \log(\frac{1}{p}) + (1 - p) \log(\frac{1}{1-p})$ and $KL(p, q) = p \log(\frac{p}{q}) + (1 - p) \log(\frac{1-p}{1-q})$ are the entropy and KL divergence of Bernoulli RVs with means p, q . Recall that $\|\mathbf{G}\|_\infty \leq \gamma$ controls the magnitude of $\langle \mathcal{L}_t, \mathbf{G} \rangle$ so for the moment, assume this is small. Then by a Taylor series $f(\langle \mathcal{L}_t, \mathbf{G} \rangle) \approx \frac{1}{2} + f'(0) \langle \mathcal{L}_t, \mathbf{G} \rangle$ using the fact that $f(0) = \frac{1}{2}$, and by another Taylor series we have

$$\begin{aligned} KL(f(\langle \mathcal{L}_t, \mathbf{G}^* \rangle) | f(\langle \mathcal{L}_t, \mathbf{G} \rangle)) &\approx KL\left(\frac{1}{2} + f'(0) \langle \mathcal{L}_t, \mathbf{G}^* \rangle \middle| \frac{1}{2} + f'(0) \langle \mathcal{L}_t, \mathbf{G} \rangle\right) \\ &\approx 2f'(0)^2 (\langle \mathcal{L}_t, \mathbf{G}^* - \mathbf{G} \rangle)^2. \end{aligned}$$

Thus, recalling the definition of $\mathcal{L}(\mathbf{G})$ from (4) we conclude that if $\widetilde{\mathbf{G}} \in \arg \min_{\mathbf{G}} R(\mathbf{G})$ with $R(\mathbf{G}) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbb{E}[\ell_f(y_t \langle \mathcal{L}_t, \mathbf{G} \rangle)]$ then one would expect $\mathcal{L}(\widetilde{\mathbf{G}}) \approx \mathcal{L}(\mathbf{G}^*)$. Moreover, since $\widehat{R}_{\mathcal{S}}(\mathbf{G})$ is an unbiased estimator of $R(\mathbf{G})$, one expects $\mathcal{L}(\widehat{\mathbf{G}})$ to approximate $\mathcal{L}(\mathbf{G}^*)$. The next theorem, combined with Theorem 1, formalizes this observation; its proof is found in the appendix.

Theorem 2. Let $C_f = \min_{t \in \mathcal{T}} \inf_{\mathbf{G} \in \mathcal{G}_{\lambda, \gamma}} |f'(\langle \mathcal{L}_t, \mathbf{G} \rangle)|$ where f' denotes the derivative of f . Then for any \mathbf{G}

$$\frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{L}(\mathbf{G}) - \mathcal{L}(\mathbf{G}^*)\|_F^2 \leq R(\mathbf{G}) - R(\mathbf{G}^*).$$

Note that if f is the logistic link function of (1) then it is straightforward to show that $|f'(\langle \mathcal{L}_t, \mathbf{G} \rangle)| \geq \frac{1}{4} \exp(-|\langle \mathcal{L}_t, \mathbf{G} \rangle|) \geq \frac{1}{4} \exp(-6\|\mathbf{G}\|_\infty)$ for any t, \mathbf{G} so it suffices to take $C_f = \frac{1}{4} \exp(-6\gamma)$.

It remains to see that we can recover \mathbf{G}^* even given $\mathcal{L}(\mathbf{G}^*)$, much less $\mathcal{L}(\widehat{\mathbf{G}})$. To do this, it is more convenient to work with distance matrices instead of Gram matrices. Analogous to the operators $\mathcal{L}_t(\mathbf{G})$ defined above, we define the operators Δ_t for $t \in \mathcal{T}$ satisfying,

$$\Delta_t(\mathbf{D}) := D_{ij} - D_{ik} \equiv \mathcal{L}_t(\mathbf{G}).$$

We will view the Δ_t as linear operators on the space of symmetric hollow $n \times n$ matrices \mathbb{S}_h^n , which includes distance matrices as special cases. As with \mathcal{L} , we can arrange all the Δ_t together, ordering the $t \in \mathcal{T}$ lexicographically, to define the $n \binom{n-1}{2}$ -dimensional vector

$$\Delta(\mathbf{D}) = [D_{12} - D_{13}, \dots, D_{ij} - D_{ik}, \dots]^T.$$

We will use the fact that $\mathcal{L}(\mathbf{G}) \equiv \Delta(\mathbf{D})$ heavily. Because $\Delta(\mathbf{D})$ consists of differences of matrix entries, Δ has a non-trivial kernel. However, it is easy to see that \mathbf{D} can be recovered given $\Delta(\mathbf{D})$ and any one off-diagonal element of \mathbf{D} , so the kernel is 1-dimensional. Also, the kernel is easy to identify by example. Consider the regular simplex in d dimensions. The distances between all $n = d + 1$ vertices are equal and the distance matrix can easily be seen to be $\mathbf{1}\mathbf{1}^T - I$. Thus $\Delta(\mathbf{D}) = \mathbf{0}$ in this case. This gives us the following simple result.

Lemma 2. Let \mathbb{S}_h^n denote the space of symmetric hollow matrices, which includes all distance matrices. For any $\mathbf{D} \in \mathbb{S}_h^n$, the set of linear functionals $\{\Delta_t(\mathbf{D}), t \in \mathcal{T}\}$ spans an $\binom{n}{2} - 1$ dimensional subspace of \mathbb{S}_h^n , and the 1-dimensional kernel is given by the span of $\mathbf{1}\mathbf{1}^T - I$.

So we see that the operator Δ is not invertible on \mathbb{S}_h^n . Define $\mathbf{J} := \mathbf{1}\mathbf{1}^T - I$. For any \mathbf{D} , let \mathbf{C} , the centered distance matrix, be the component of \mathbf{D} orthogonal to the kernel of \mathcal{L} (i.e., $\text{tr}(\mathbf{C}\mathbf{J}) = 0$). Then we have the orthogonal decomposition

$$\mathbf{D} = \mathbf{C} + \sigma_D \mathbf{J},$$

where $\sigma_D = \text{trace}(\mathbf{D}\mathbf{J})/\|\mathbf{J}\|_F^2$. Since \mathbf{G} is assumed to be centered, the value of σ_D has a simple interpretation:

$$\sigma_D = \frac{1}{2 \binom{n}{2}} \sum_{1 \leq i < j \leq n} D_{ij} = \frac{2}{n-1} \sum_{1 \leq i \leq n} \langle x_i, x_i \rangle = \frac{2\|\mathbf{G}\|_*}{n-1}, \quad (9)$$

the average of the squared distances or alternatively a scaled version of the nuclear norm of \mathbf{G} .

Let $\widehat{\mathbf{D}}$ and $\widehat{\mathbf{C}}$ be the corresponding distance and centered distance matrices corresponding to $\widehat{\mathbf{G}}$ the solution to 8. Though Δ is not invertible on all \mathbb{S}_h^n , it is invertible on the subspace orthogonal to the kernel, namely \mathbf{J}^\perp . So if $\Delta(\widehat{\mathbf{D}}) \approx \Delta(\mathbf{D}^*)$, or equivalently $\mathcal{L}(\widehat{\mathbf{G}}) \approx \mathcal{L}(\mathbf{G}^*)$, we expect $\widehat{\mathbf{C}}$ to be close to \mathbf{C}^* . The next theorem quantifies this.

Theorem 3. Consider the setting of Theorems 1 and 2 and let $\widehat{\mathbf{C}}, \mathbf{C}^*$ be defined as above. Then

$$\frac{1}{2 \binom{n}{2}} \|\widehat{\mathbf{C}} - \mathbf{C}^*\|_F^2 \leq \frac{L\lambda}{4C_f^2|\mathcal{S}|} \left(\sqrt{\frac{18|\mathcal{S}|\log(n)}{n}} + \frac{\sqrt{3}}{3} \log n \right) + \frac{L\gamma}{4C_f^2} \sqrt{\frac{288 \log 2/\delta}{|\mathcal{S}|}}$$

Proof. By combining Theorem 2 with the prediction error bounds obtained in 1 we see that

$$\frac{2C_f^2}{n \binom{n-1}{2}} \|\mathcal{L}(\widehat{\mathbf{G}}) - \mathcal{L}(\mathbf{G}^*)\|_F^2 \leq \frac{4L\lambda}{|\mathcal{S}|} \left(\sqrt{\frac{18|\mathcal{S}|\log(n)}{n}} + \frac{\sqrt{3}}{3} \log n \right) + L\gamma \sqrt{\frac{288 \log 2/\delta}{|\mathcal{S}|}}.$$

Next we employ the following *restricted isometry property* of Δ on the subspace \mathbf{J}^\perp whose proof is in the supplementary materials.

Lemma 3. Let \mathbf{D} and \mathbf{D}' be two different distance matrices of n points in \mathbb{R}^d and $\mathbb{R}^{d'}$. Let \mathbf{C} and \mathbf{C}' be the components of \mathbf{D} and \mathbf{D}' orthogonal to \mathbf{J} . Then

$$n\|\mathbf{C} - \mathbf{C}'\|_F^2 \leq \|\Delta(\mathbf{C}) - \Delta(\mathbf{C}')\|^2 = \|\Delta(\mathbf{D}) - \Delta(\mathbf{D}')\|^2 \leq 2(n-1)\|\mathbf{C} - \mathbf{C}'\|_F^2.$$

The result then follows. \square

This implies that by collecting enough samples, we can recover the centered distance matrix. By applying the discussion following Theorem 1 when \mathbf{G}^* is rank d , we can state an upperbound of $\frac{1}{2\binom{n}{2}}\|\widehat{\mathbf{C}} - \mathbf{C}^*\|_F^2 \leq O\left(\frac{L\gamma}{C_f^2} \sqrt{\frac{dn \log(n) + \log(1/\delta)}{|\mathcal{S}|}}\right)$. However, it is still not clear that this is enough to recover \mathbf{D}^* or \mathbf{G}^* . Remarkably, despite this unknown component being in the kernel, we show next that it can be recovered.

Theorem 4. Let \mathbf{D} be a distance matrix of n points in \mathbb{R}^d , let \mathbf{C} be the component of \mathbf{D} orthogonal to the kernel of Δ , and let $\lambda_2(\mathbf{C})$ denote the second largest eigenvalue of \mathbf{C} . If $n > d + 2$, then

$$\mathbf{D} = \mathbf{C} + \lambda_2(\mathbf{C})\mathbf{J}. \quad (10)$$

This shows that \mathbf{D} is uniquely determined as a function of \mathbf{C} . Therefore, since $\Delta(\mathbf{D}) = \Delta(\mathbf{C})$ and because \mathbf{C} is orthogonal to the kernel of Δ , the distance matrix \mathbf{D} can be recovered from $\Delta(\mathbf{D})$, even though the linear operator Δ is non-invertible.

We now provide a proof of Theorem 4 in the case where $n > d + 3$. The result is true in the case when $n > d + 2$ but requires a more detailed analysis. This includes the construction of a vector \mathbf{x} such that $\mathbf{D}\mathbf{x} = \mathbf{1}$ and $\mathbf{1}^T \mathbf{x} \geq 0$ for any distance matrix a result in [17].

Proof. To prove Theorem 4 we need the following lemma, proved in the supplementary materials.

Lemma 4. Let \mathbf{D} be a Euclidean distance matrix on n points. Then \mathbf{D} is negative semidefinite on the subspace

$$\mathbf{1}^\perp := \{\mathbf{x} \in \mathbb{R}^n | \mathbf{1}^T \mathbf{x} = 0\}.$$

Furthermore, $\ker(\mathbf{D}) \subset \mathbf{1}^\perp$.

For any matrix \mathbf{M} , let $\lambda_i(\mathbf{M})$ denote its i th largest eigenvalue. Under the conditions of the theorem, we show that for $\sigma > 0$, $\lambda_2(\mathbf{D} - \sigma\mathbf{J}) = \sigma$. Since $\mathbf{C} = \mathbf{D} - \sigma\mathbf{J}$, this proves the theorem.

Note that, $\lambda_i(\mathbf{D} - \sigma\mathbf{J}) = \lambda_i(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T) + \sigma$ for $1 \leq i \leq n$ and σ arbitrary. So it suffices to show that $\lambda_2(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T) = 0$.

By Weyl's Theorem

$$\lambda_2(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T) \leq \lambda_2(\mathbf{D}) + \lambda_1(-\sigma\mathbf{1}\mathbf{1}^T).$$

Since $\lambda_1(-\sigma\mathbf{1}\mathbf{1}^T) = 0$, we have $\lambda_2(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T) \leq \lambda_2(\mathbf{D}) = 0$. By the Courant-Fischer Theorem

$$\lambda_2(\mathbf{D}) = \min_{U: \dim(U)=n-1} \max_{\mathbf{x} \in U, \mathbf{x} \neq 0} \frac{\mathbf{x}^T \mathbf{D} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \min_{U=\mathbf{1}^\perp} \max_{\mathbf{x} \in U, \mathbf{x} \neq 0} \frac{\mathbf{x}^T \mathbf{D} \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq 0$$

since \mathbf{D} negative semidefinite on $\mathbf{1}^\perp$. Now let \mathbf{v}_i denote the i th eigenvector of \mathbf{D} with eigenvalue $\lambda_i = 0$. Then

$$(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T)\mathbf{v}_i = \mathbf{D}\mathbf{v}_i = 0,$$

since $\mathbf{v}_i^T \mathbf{1} = 0$ by 4. So $\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T$ has at least $n - d - 2$ zero eigenvalues, since $\text{rank } \mathbf{D} \leq d + 2$. In particular, if $n > d + 3$, then $\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T$ must have at least two eigenvalues equal to 0. Therefore, $\lambda_2(\mathbf{D} - \sigma\mathbf{1}\mathbf{1}^T) = 0$. \square

The previous theorem along with Theorem 3 guarantees that we can recover \mathbf{G}^* as we increase the number of triplets sampled. The final theorem, which follows directly from Theorems 3 and 4, summarizes this.

Theorem 5. Assume $n > d + 2$ and consider the setting of Theorems 1 and 2. As $|\mathcal{S}| \rightarrow \infty$, $\widehat{\mathbf{D}} \rightarrow \mathbf{D}^*$ where $\widehat{\mathbf{D}}$ is the distance matrix corresponding to $\widehat{\mathbf{G}}$ (the solution to 8).

Proof. Recall $\widehat{\mathbf{D}} = \widehat{\mathbf{C}} + \lambda_2(\widehat{\mathbf{C}})\mathbf{J}$, so as $\widehat{\mathbf{C}} \rightarrow \mathbf{C}^*$, $\widehat{\mathbf{D}} \rightarrow \mathbf{D}^*$. \square

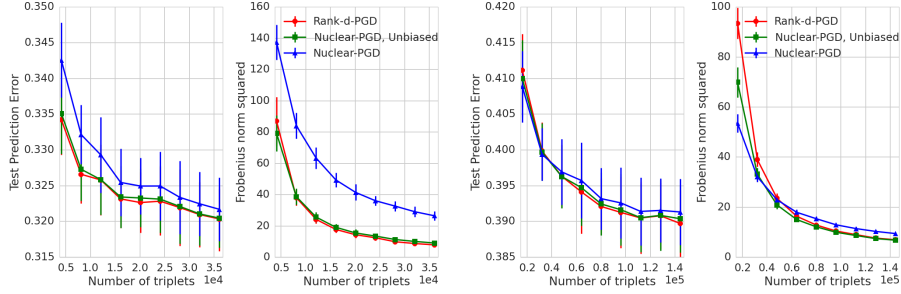


Figure 1: \mathbf{G}^* generated with $n = 64$ points in $d = 2$ and $d = 8$ dimensions on the left and right.

4 Experimental Study

The section empirically studies the properties of estimators suggested by our theory. It is *not* an attempt to perform an exhaustive empirical evaluation of different embedding techniques; for that see [18, 4, 6, 3]. In what follows each of the n points is generated randomly: $\mathbf{x}_i \sim \mathcal{N}(0, \frac{1}{2d}I_d) \in \mathbb{R}^d$, $i = 1, \dots, n$, motivated by the observation that

$$\mathbb{E}[\langle \mathcal{L}_t, \mathbf{G}^* \rangle] = \mathbb{E}[\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 - \|\mathbf{x}_i - \mathbf{x}_k\|_2^2] \leq \mathbb{E}[\|\mathbf{x}_i - \mathbf{x}_j\|_2^2] = 2\mathbb{E}[\|\mathbf{x}_i\|_2^2] = 1$$

for any triplet $t = (i, j, k)$. We report the prediction error on a holdout set of 10,000 triplets and the error in Frobenius norm of the estimated Gram matrix over 36 random trials. We minimize the logistic MLE objective $\hat{R}_S(\mathbf{G}) = \frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} \log(1 + \exp(-y_t \langle \mathcal{L}_t, \mathbf{G} \rangle))$.

For each algorithm considered, the domain of the objective variable \mathbf{G} is the space of symmetric positive semi-definite matrices. None of the methods impose the constraint $\max_{i,j} |G_{ij}| \leq \gamma$ (as done above), since this was used to simplify the analysis and does not have a large impact in practice. *Rank-d Projected Gradient Descent (PGD)* performs gradient descent on the objective $\hat{R}_S(\mathbf{G})$ with line search, projecting onto the subspace spanned by the top d eigenvalues at each step (i.e. setting the smallest $n - d$ eigenvalues to 0). *Nuclear Norm PGD* performs gradient descent on $\hat{R}_S(\mathbf{G})$ projecting onto the nuclear norm ball with radius $\|\mathbf{G}^*\|_*$, where \mathbf{G}^* is the Gram matrix of the latent embedding. The nuclear norm projection can have the undesirable effect of shrinking the non-zero eigenvalues toward the origin. To compensate for this potential bias, we employ *Nuclear Norm PGD Debaised*, which takes the biased output of *Nuclear Norm PGD*, decomposes it into $\mathbf{U}\mathbf{E}\mathbf{U}^T$ where $\mathbf{U} \in \mathbb{R}^{n \times d}$ are the top d eigenvectors, and outputs $\mathbf{U}\text{diag}(\hat{\mathbf{s}})\mathbf{U}^T$ where $\hat{\mathbf{s}} = \arg \min_{\mathbf{s} \in \mathbb{R}^d} \hat{R}_S(\mathbf{U}\text{diag}(\mathbf{s})\mathbf{U}^T)$. This last algorithm is motivated by the observation that methods for minimizing $\|\cdot\|_1$ or $\|\cdot\|_*$ are good at identifying the true support of a signal, but output biased magnitudes [19]. *Rank-d PGD* and *Nuclear Norm PGD Debaised* are novel ordinal embedding algorithms.

Figure 1 presents how the algorithms behave for $n = 64$ and $d = 2, 8$. We observe that the unbiased nuclear norm solution behaves near-identically to the rank- d solution and remark that this was observed in all of our experiments (see the supplementary materials for other values of n , d , and scalings of \mathbf{G}^*). A popular technique for recovering rank d embeddings is to perform (stochastic) gradient descent on $\hat{R}_S(\mathbf{U}^T\mathbf{U})$ with objective variable $\mathbf{U} \in \mathbb{R}^{n \times d}$ taken as the embedding [18, 4, 6]. In all of our experiments this method produced Gram matrices nearly identical to those produced by our *Rank-d-PGD* method, but *Rank-d-PGD* was an order of magnitude faster in our implementation. Also, in light of our isometry theorem, we can show that the Hessian of $\mathbb{E}[\hat{R}_S(\mathbf{G})]$ is nearly a scaled identity, leading us to hypothesize that a globally optimal linear convergence result for this non-convex optimization may be possible using the techniques of [20, 21]. Finally, we note that previous literature has reported that nuclear norm optimizations like *Nuclear Norm PGD* tend to produce less accurate embeddings than those of non-convex methods [4, 6]. The results imply that *Nuclear Norm PGD Debaised* appears to close the performance gap between the convex and non-convex solutions.

Acknowledgments This work was partially supported by the NSF grants CCF-1218189 and IIS-1447449, the NIH grant 1 U54 AI117924-01, the AFOSR grant FA9550-13-1-0138, and by ONR awards N00014-15-1-2620, and N00014-13-1-0129. We would also like to thank Amazon Web Services for providing the computational resources used for running our simulations.

References

- [1] Roger N Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function. i. *Psychometrika*, 27(2):125–140, 1962.
- [2] Joseph B Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29(2):115–129, 1964.
- [3] Sameer Agarwal, Josh Wills, Lawrence Cayton, Gert Lanckriet, David J Kriegman, and Serge Belongie. Generalized non-metric multidimensional scaling. In *International Conference on Artificial Intelligence and Statistics*, pages 11–18, 2007.
- [4] Omer Tamuz, Ce Liu, Ohad Shamir, Adam Kalai, and Serge J Belongie. Adaptively learning the crowd kernel. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 673–680, 2011.
- [5] Kevin G Jamieson and Robert D Nowak. Low-dimensional embedding using adaptively selected ordinal data. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 1077–1084. IEEE, 2011.
- [6] Laurens Van Der Maaten and Kilian Weinberger. Stochastic triplet embedding. In *Machine Learning for Signal Processing (MLSP), 2012 IEEE International Workshop on*, pages 1–6. IEEE, 2012.
- [7] Brian McFee and Gert Lanckriet. Learning multi-modal similarity. *The Journal of Machine Learning Research*, 12:491–523, 2011.
- [8] Matthäus Kleindessner and Ulrike von Luxburg. Uniqueness of ordinal embedding. In *COLT*, pages 40–67, 2014.
- [9] Yoshikazu Terada and Ulrike V Luxburg. Local ordinal embedding. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 847–855, 2014.
- [10] Ery Arias-Castro. Some theory for ordinal embedding. *arXiv preprint arXiv:1501.02861*, 2015.
- [11] Mark A Davenport, Yaniv Plan, Ewout van den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference*, 3(3), 2014.
- [12] Yu Lu and Sahand N Negahban. Individualized rank aggregation using nuclear norm regularization. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1473–1479. IEEE, 2015.
- [13] D. Park, J , Neeman, J. Zhang, S. Sanghavi, and I. Dhillon. Preference completion: Large-scale collaborative ranking from pairwise comparisons. *Proc. Int. Conf. Machine Learning (ICML)*, 2015.
- [14] Jon Dattorro. *Convex Optimization & Euclidean Distance Geometry*. Meboo Publishing USA, 2011.
- [15] Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics*, 9:323–375, 2005.
- [16] Joel A. Tropp. An introduction to matrix concentration inequalities, 2015.
- [17] Pablo Tarazaga and Juan E. Gallardo. Euclidean distance matrices: new characterization and boundary properties. *Linear and Multilinear Algebra*, 57(7):651–658, 2009.
- [18] Kevin G Jamieson, Lalit Jain, Chris Fernandez, Nicholas J Glattard, and Rob Nowak. Next: A system for real-world development, evaluation, and application of active learning. In *Advances in Neural Information Processing Systems*, pages 2638–2646, 2015.
- [19] Nikhil Rao, Parikshit Shah, and Stephen Wright. Conditional gradient with enhancement and truncation for atomic norm regularization. In *NIPS workshop on Greedy Algorithms*, 2013.
- [20] Samet Oymak, Benjamin Recht, and Mahdi Soltanolkotabi. Sharp time–data tradeoffs for linear inverse problems. *arXiv preprint arXiv:1507.04793*, 2015.
- [21] Jie Shen and Ping Li. A tight bound of hard thresholding. *arXiv preprint arXiv:1605.01656*, 2016.