
Learning Low-Dimensional Metrics

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

Blake Mason *
University of Wisconsin
Madison, WI 53706
bmason3@wisc.edu

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

Abstract

This paper investigates the theoretical foundations of metric learning, focused on three key questions that are not fully addressed in prior work: 1) we consider learning general low-dimensional (low-rank) metrics as well as sparse metrics; 2) we develop upper and lower (minimax) bounds on the generalization error; 3) we quantify the sample complexity of metric learning in terms of the dimension of the feature space and the dimension/rank of the underlying metric; 4) we also bound the accuracy of the learned metric relative to the underlying true generative metric. All the results involve novel mathematical approaches to the metric learning problem, and also shed new light on the special case of ordinal embedding (aka non-metric multidimensional scaling).

1 Low-Dimensional Metric Learning

This paper studies the problem of learning a low-dimensional Euclidean metric from comparative judgments. Specifically, consider a set of n items with high-dimensional features $\mathbf{x}_i \in \mathbb{R}^p$ and suppose we are given a set of (possibly noisy) distance comparisons of the form

$$\text{sign}(\text{dist}(\mathbf{x}_i, \mathbf{x}_j) - \text{dist}(\mathbf{x}_i, \mathbf{x}_k)),$$

for a subset of all possible triplets of the items. Here we have in mind comparative judgments made by humans and the distance function implicitly defined according to human perceptions of similarities and differences. For example, the items could be images and the \mathbf{x}_i could be visual features automatically extracted by a machine. Accordingly, our goal is to learn a $p \times p$ symmetric positive semi-definite (psd) matrix \mathbf{K} such that the metric $d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_j) := (\mathbf{x}_i - \mathbf{x}_j)^T \mathbf{K} (\mathbf{x}_i - \mathbf{x}_j)$, where $d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_j)$ denotes the squared distance between items i and j with respect to a matrix \mathbf{K} , predicts the given distance comparisons as well as possible. Furthermore, it is often desired that the metric is *low-dimensional* relative to the original high-dimensional feature representation (i.e., $\text{rank}(\mathbf{K}) \leq d < p$). There are several motivations for this:

- Learning a high-dimensional metric may be infeasible from a limited number of comparative judgments, and encouraging a low-dimensional solution is a natural regularization.
- Cognitive scientists are often interested in visualizing human perceptual judgments (e.g., in a two-dimensional representation) and determining which features most strongly influence human perceptions. For example, educational psychologists in [1] collected comparisons between visual representations of chemical molecules in order to identify a small set of visual features that most significantly influence the judgments of beginning chemistry students.
- It is sometimes reasonable to hypothesize that a small subset of the high-dimensional features dominate the underlying metric (i.e., many irrelevant features).
- Downstream applications of the learned metric (e.g., for classification purposes) may benefit from robust, low-dimensional metrics.

* Authors contributed equally to this paper and are listed alphabetically.

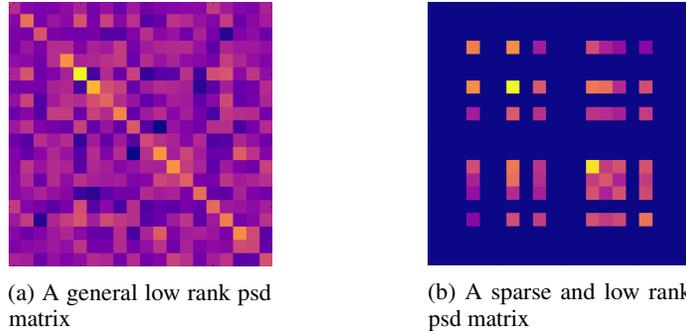


Figure 1: Examples of \mathbf{K} for $p = 20$ and $d = 7$. The sparse case depicts a situation in which only some of the features are relevant to the metric.

With this in mind, several authors have proposed nuclear norm and $\ell_{1,2}$ group lasso norm regularization to encourage low-dimensional and sparse metrics as in Fig. 1b (see [2] for a review). Relative to such prior work, the contributions of this paper are three-fold:

1. We develop novel upper bounds on the generalization error and sample complexity of learning low-dimensional metrics from triplet distance comparisons. Notably, unlike previous generalization bounds, our bounds allow one to easily quantify how the feature space dimension p and rank or sparsity $d < p$ of the underlying metric impacts the sample complexity.
2. We establish minimax lower bounds for learning low-rank and sparse metrics that match the upper bounds up to polylogarithmic factors, demonstrating the optimality of learning algorithms for the first time. Moreover, the upper and lower bounds demonstrate that learning sparse (and low-rank) metrics is essentially as difficult as learning a general low-rank metric. This suggests that nuclear norm regularization may be preferable in practice, since it places less restrictive assumptions on the problem.
3. We use the generalization error bounds to obtain model identification error bounds that quantify the accuracy of the learned \mathbf{K} matrix. This problem has received very little, if any, attention in the past and is crucial for interpreting the learned metrics (e.g., in cognitive science applications). This is a bit surprising, since the term “metric learning” strongly suggests accurately determining a metric, not simply learning a predictor that is parameterized by a metric.

1.1 Comparison with Previous Work

There is a fairly large body of work on metric learning which is nicely reviewed and summarized in the monograph [2], and we refer the reader to it for a comprehensive summary of the field. Here we discuss a few recent works most closely connected to this paper. Several authors have developed generalization error bounds for metric learning, as well as bounds for downstream applications, such as classification, based on learned metrics. To use the terminology of [2], most of the focus has been on must-link/cannot-link constraints and less on relative constraints (i.e., triplet constraints as considered in this paper). Generalization bounds based on algorithmic robustness are studied in [3], but the generality of this framework makes it difficult to quantify the sample complexity of specific cases, such as low-rank or sparse metric learning. Rademacher complexities are used to establish generalization error bounds in the must-link/cannot-link situation in [4, 5, 6], but do not consider the case of relative/triplet constraints. The sparse compositional metric learning framework of [7] does focus on relative/triplet constraints and provides generalization error bounds in terms of covering numbers. However, this work does not provide bounds on the covering numbers, making it difficult to quantify the sample complexity. To sum up, prior work does not quantify the sample complexity of metric learning based on relative/triplet constraints in terms of the intrinsic problem dimensions (i.e., dimension p of the high-dimensional feature space and the dimension d of the underlying metric), there is no prior work on lower bounds, and no prior work quantifying the accuracy of learned metrics themselves (i.e., only bounds on prediction errors, not model identification errors). Finally we mention that Fazel et al [8] also consider the recovery of sparse and low rank matrices from linear observations. Our situation is very different, our matrices are low rank because they are sparse - not sparse and simultaneously low rank as in their case.

2 The Metric Learning Problem

Consider n known points $\mathbf{X} := [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{p \times n}$. We are interested in learning a symmetric positive semidefinite matrix \mathbf{K} that specifies a metric on \mathbb{R}^p given ordinal constraints on distances between the known points. Let \mathcal{S} denote a set of triplets, where each $t = (i, j, k) \in \mathcal{S}$ is drawn uniformly at random from the full set of $n \binom{n-1}{2}$ triplets $\mathcal{T} := \{(i, j, k) : 1 \leq i \neq j \neq k \leq n, j < k\}$. For each triplet, we observe a $y_t \in \{\pm 1\}$ which is a noisy indication of the triplet constraint $d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_j) < d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_k)$. Specifically we assume that each t has an associated probability q_t of $y_t = -1$, and all y_t are statistically independent.

Objective 1: Compute an estimate $\widehat{\mathbf{K}}$ from \mathcal{S} that predicts triplets as well as possible.

In many instances, our triplet measurements are noisy observations of triplets from a true positive semi-definite matrix \mathbf{K}^* . In particular we assume

$$q_t > 1/2 \iff d_{\mathbf{K}^*}(\mathbf{x}_i, \mathbf{x}_j) < d_{\mathbf{K}^*}(\mathbf{x}_i, \mathbf{x}_k).$$

We can also assume an explicit known *link function*, $f : \mathbb{R} \rightarrow [0, 1]$, so that $q_t = f(d_{\mathbf{K}^*}(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{K}^*}(\mathbf{x}_i, \mathbf{x}_k))$.

Objective 2: Assuming an explicit known link function f estimate \mathbf{K}^* from \mathcal{S} .

2.1 Definitions and Notation

Our triplet observations are nonlinear transformations of a linear function of the Gram matrix $\mathbf{G} := \mathbf{X}^T \mathbf{K} \mathbf{X}$. Indeed for any triple $t = (i, j, k)$, define

$$\begin{aligned} \mathbf{M}_t(\mathbf{K}) &:= d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{K}}(\mathbf{x}_i, \mathbf{x}_k) \\ &= \mathbf{x}_i^T \mathbf{K} \mathbf{x}_k + \mathbf{x}_k^T \mathbf{K} \mathbf{x}_i - \mathbf{x}_i^T \mathbf{K} \mathbf{x}_j - \mathbf{x}_j^T \mathbf{K} \mathbf{x}_i + \mathbf{x}_j^T \mathbf{K} \mathbf{x}_j - \mathbf{x}_k^T \mathbf{K} \mathbf{x}_k. \end{aligned}$$

So for every $t \in \mathcal{S}$, y_t is a noisy measurement of $\text{sign}(\mathbf{M}_t(\mathbf{K}))$. This linear operator may also be expressed as a matrix

$$\mathbf{M}_t := \mathbf{x}_i \mathbf{x}_k^T + \mathbf{x}_k \mathbf{x}_i^T - \mathbf{x}_i \mathbf{x}_j^T - \mathbf{x}_j \mathbf{x}_i^T + \mathbf{x}_j \mathbf{x}_j^T - \mathbf{x}_k \mathbf{x}_k^T,$$

so that $\mathbf{M}_t(\mathbf{K}) = \langle \mathbf{M}_t, \mathbf{K} \rangle = \text{Trace}(\mathbf{M}_t^T \mathbf{K})$. We will use \mathbf{M}_t to denote the operator and associated matrix interchangeably. Ordering the elements of \mathcal{T} lexicographically, we let \mathcal{M} denote the linear map,

$$\mathcal{M}(\mathbf{K}) = (\mathbf{M}_t(\mathbf{K}) \mid \text{for } t \in \mathcal{T}) \in \mathbb{R}^{n \binom{n-1}{2}}$$

Given a PSD matrix \mathbf{K} and a sample, $t \in \mathcal{S}$, we let $\ell(y_t \langle \mathbf{M}_t, \mathbf{K} \rangle)$ denote the loss of \mathbf{K} with respect to t ; e.g., the 0-1 loss $\mathbb{1}_{\{\text{sign}(y_t \langle \mathbf{M}_t, \mathbf{K} \rangle) \neq 1\}}$, the hinge-loss $\max\{0, 1 - y_t \langle \mathbf{M}_t, \mathbf{K} \rangle\}$, or the logistic loss $\log(1 + \exp(-y_t \langle \mathbf{M}_t, \mathbf{K} \rangle))$. Note that we insist that our losses be functions of our triplet differences $\langle \mathbf{M}_t, \mathbf{K} \rangle$. Further, note that this makes our losses invariant to rigid motions of the points \mathbf{x}_i . Other models proposed for metric learning use scale-invariant loss functions [9].

For a given loss ℓ , we then define the empirical risk with respect to our set of observations \mathcal{S} to be

$$\widehat{R}_{\mathcal{S}}(\mathbf{K}) := \frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} \ell(y_t \langle \mathbf{M}_t, \mathbf{K} \rangle).$$

This is an unbiased estimator of the true risk $R(\mathbf{K}) := \mathbb{E}[\ell(y_t \langle \mathbf{M}_t, \mathbf{K} \rangle)]$ where the expectation is taken with respect to a triplet t selected uniformly at random and the random value of y_t .

Finally, we let \mathbf{I}_n denote the identity matrix in $\mathbb{R}^{n \times n}$, $\mathbf{1}_n$ the n -dimensional vector of all ones and $\mathbf{V} := \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$ the *centering matrix*. In particular if $\mathbf{X} \in \mathbb{R}^{p \times n}$ is a set of points, $\mathbf{X} \mathbf{V}$ subtracts the mean of the columns of \mathbf{X} from each column. We say that \mathbf{X} is centered if $\mathbf{X} \mathbf{V} = 0$, or equivalently $\mathbf{X} \mathbf{1}_n = 0$. If \mathbf{G} is the Gram matrix of the set of points \mathbf{X} , i.e. $\mathbf{G} = \mathbf{X}^T \mathbf{X}$, then we say that \mathbf{G} is centered if \mathbf{X} is centered or if equivalently, $\mathbf{G} \mathbf{1}_n = 0$. Furthermore we use $\|\cdot\|_*$ to denote the nuclear norm, and $\|\cdot\|_{1,2}$ to denote the mixed $\ell_{1,2}$ norm of a matrix, the sum of the ℓ_2 norms of its rows. Unless otherwise specified, we take $\|\cdot\|$ to be the standard operator norm when applied to matrices and the standard Euclidean norm when applied to vectors. Finally we define the \mathbf{K} -norm of a vector as $\|\mathbf{x}\|_{\mathbf{K}}^2 := \mathbf{x}^T \mathbf{K} \mathbf{x}$.

2.2 Sample Complexity of Learning Metrics.

In most applications, we are interested in learning a matrix \mathbf{K} that is low-rank and positive-semidefinite. Furthermore as we will show in Theorem 2.1, such matrices can be learned using fewer samples than general psd matrices. As is common in machine learning applications, we relax the rank constraint to a nuclear norm constraint. In particular, let our constraint set be

$$\mathcal{K}_{\lambda, \gamma} = \{\mathbf{K} \in \mathbb{R}^{p \times p} \mid \mathbf{K} \text{ positive-semidefinite, } \|\mathbf{K}\|_* \leq \lambda, \max_{t \in \mathcal{T}} \langle \mathbf{M}_t, \mathbf{K} \rangle \leq \gamma\}.$$

Up to constants, a bound on $\langle \mathbf{M}_t, \mathbf{K} \rangle$ is a bound on $\mathbf{x}_i^T \mathbf{K} \mathbf{x}_i$. This bound along with assuming our loss function is Lipschitz, will lead to a tighter bound on the deviation of $\widehat{R}_S(\mathbf{K})$ from $R(\mathbf{K})$ crucial in our upper bound theorem.

Let $\mathbf{K}^* := \min_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} R(\mathbf{K})$ be the true risk minimizer in this class, and let $\widehat{\mathbf{K}} := \min_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} \widehat{R}_S(\mathbf{K})$ be the empirical risk minimizer. We achieve the following prediction error bounds for the empirical risk minimizer.

Theorem 2.1. *Fix $\lambda, \gamma, \delta > 0$. In addition assume that $\max_{1 \leq i \leq n} \|\mathbf{x}_i\|^2 = 1$. If the loss function ℓ is L -Lipschitz, then with probability at least $1 - \delta$*

$$R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*) \leq 4L \left(\sqrt{\frac{140\lambda^2 \frac{\|\mathbf{X}\mathbf{X}^T\|}{n} \log p}{|\mathcal{S}|}} + \frac{2 \log p}{|\mathcal{S}|} \right) + \sqrt{\frac{2L^2 \gamma^2 \log 2/\delta}{|\mathcal{S}|}}$$

Note that past generalization error bounds in the metric learning literature have failed to quantify the precise dependence on observation noise, dimension, rank, and our features \mathbf{X} . Consider the fact that a $p \times p$ matrix with rank d has $O(dp)$ degrees of freedom. With that in mind, one expects the sample complexity to be also roughly $O(dp)$. We next show that this intuition is correct if the original representation \mathbf{X} is isotropic (i.e., has no preferred direction).

The Isotropic Case. Suppose that $\mathbf{x}_1, \dots, \mathbf{x}_n, n > p$, are drawn independently from the isotropic Gaussian $\mathcal{N}(\mathbf{0}, \frac{1}{p}\mathbf{I})$. Furthermore, suppose that $\mathbf{K}^* = \frac{p}{\sqrt{d}}\mathbf{U}\mathbf{U}^T$ with $\mathbf{U} \in \mathbb{R}^{p \times d}$ is a generic (dense) orthogonal matrix with unit norm columns. The factor $\frac{p}{\sqrt{d}}$ is simply the scaling needed so that the average magnitude of the entries in \mathbf{K}^* is a constant, independent of the dimensions p and d . In this case, $\text{rank}(\mathbf{K}^*) = d$ and $\|\mathbf{K}^*\|_F = \text{trace}(\mathbf{U}^T\mathbf{U}) = p$. These two facts imply that the tightest bound on the nuclear norm of \mathbf{K}^* is $\|\mathbf{K}^*\|_* \leq p\sqrt{d}$. Thus, we take $\lambda = p\sqrt{d}$ for the nuclear norm constraint. Now let $\mathbf{z}_i = \sqrt{\frac{p}{\sqrt{d}}}\mathbf{U}^T \mathbf{x}_i \sim N(\mathbf{0}, \mathbf{I}_d)$ and note that $\|\mathbf{x}_i\|_{\mathbf{K}^*}^2 = \|\mathbf{z}_i\|^2 \sim \chi_d^2$.

Therefore, $\mathbb{E}\|\mathbf{x}_i\|_{\mathbf{K}^*}^2 = d$ and it follows from standard concentration bounds that with large probability $\max_i \|\mathbf{x}_i\|_{\mathbf{K}^*}^2 \leq 5d \log n =: \gamma$ see [10]. Also, because the $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \frac{1}{p}\mathbf{I})$ it follows that if $n > p \log p$, say, then with large probability $\|\mathbf{X}\mathbf{X}^T\| \leq 5n/p$. We now plug these calculations into Theorem 2.1 to obtain the following corollary.

Corollary 2.1.1 (Sample complexity for isotropic points). *Fix $\delta > 0$, set $\lambda = p\sqrt{d}$, and assume that $\|\mathbf{X}\mathbf{X}^T\| = O(n/p)$ and $\gamma := \max_i \|\mathbf{x}_i\|_{\mathbf{K}^*}^2 = O(d \log n)$. Then for a generic $\mathbf{K}^* \in \mathcal{K}_{\lambda, \gamma}$, as constructed above, with probability at least $1 - \delta$,*

$$R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*) = O \left(\sqrt{\frac{dp(\log p + \log^2 n)}{|\mathcal{S}|}} \right)$$

This bound agrees with the intuition that the sample complexity should grow roughly like dp , the degrees of freedom on \mathbf{K}^* . Moreover, our minimax lower bound in Theorem 2.3 below shows that, ignoring logarithmic factors, the general upper bound in Theorem 2.1 is unimprovable in general.

Beyond low rank metrics, in many applications it is reasonable to assume that only a few of the features are salient and should be given nonzero weight. Such a metric may be learned by insisting \mathbf{K} to be row sparse in addition to being low rank. Whereas learning a low rank \mathbf{K} assumes that distance is well represented in a low dimensional subspace, a row sparse (and hence low rank) \mathbf{K} defines a metric using only a subset of the features. Figure 1 gives a comparison of a low rank versus a low rank and sparse matrix \mathbf{K} .

Analogous to the convex relaxation of rank by the nuclear norm, it is common to relax row sparsity by using the mixed $\ell_{1,2}$ norm. In fact, the geometry of the $\ell_{1,2}$ and nuclear norm balls are tightly related as the following lemma shows.

Lemma 2.2. *For a symmetric positive semi-definite matrix $\mathbf{K} \in \mathbb{R}^{p \times p}$, $\|\mathbf{K}\|_* \leq \|\mathbf{K}\|_{1,2}$.*

$$\text{Proof. } \|\mathbf{K}\|_{1,2} = \sum_{i=1}^p \sqrt{\sum_{j=1}^p \mathbf{K}_{i,j}^2} \geq \sum_{i=1}^p \mathbf{K}_{i,i} = \text{Trace}(\mathbf{K}) = \sum_{i=1}^p \lambda_i(\mathbf{K}) = \|\mathbf{K}\|_* \quad \square$$

This implies that the $\ell_{1,2}$ ball of a given radius is contained inside the nuclear norm ball of the same radius. In particular, it is reasonable to assume that it is easier to learn a \mathbf{K} that is sparse in addition to being low rank. Surprisingly, however, the following minimax bound shows that this is not necessarily the case.

To make this more precise, we will consider optimization over the set

$$\mathcal{K}'_{\lambda,\gamma} = \{\mathbf{K} \in \mathbb{R}^{p \times p} \mid \mathbf{K} \text{ positive-semidefinite, } \|\mathbf{K}\|_{1,2} \leq \lambda, \max_{t \in \mathcal{T}} \langle \mathbf{M}_t, \mathbf{K} \rangle \leq \gamma\}.$$

Furthermore, we must specify the way in which our data could be generated from noisy triplet observations of a fixed \mathbf{K}^* . To this end, assume the existence of a *link function* $f : \mathbb{R} \rightarrow [0, 1]$ so that $q_t = \mathbb{P}(y_t = -1) = f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle)$ governs the observations. There is a natural associated logarithmic loss function ℓ_f corresponding to the log-likelihood, where the loss of an arbitrary \mathbf{K} is

$$\ell_f(y_t \langle \mathbf{M}_t, \mathbf{K} \rangle) = \mathbb{1}_{\{y_t = -1\}} \log \frac{1}{f(\langle \mathbf{M}_t, \mathbf{K} \rangle)} + \mathbb{1}_{\{y_t = 1\}} \log \frac{1}{1 - f(\langle \mathbf{M}_t, \mathbf{K} \rangle)}$$

Theorem 2.3. *Choose a link function f and let ℓ_f be the associated logarithmic loss. For p sufficiently large, then there exists a choice of γ , λ , \mathbf{X} , and $|\mathcal{S}|$ such that*

$$\inf_{\widehat{\mathbf{K}}} \sup_{\mathbf{K} \in \mathcal{K}'_{\lambda,\gamma}} \mathbb{E}[R(\widehat{\mathbf{K}})] - R(\mathbf{K}) \geq C \sqrt{\frac{C_1^3 \ln 4 \lambda^2 \|\mathbf{X}\mathbf{X}^T\|}{2 |\mathcal{S}| n}}$$

where $C = \frac{C_f^2}{32} \sqrt{\frac{\inf_{|x| \leq \gamma} f(x)(1-f(x))}{\sup_{|\nu| \leq \gamma} f'(\nu)^2}}$ with $C_f = \inf_{|x| \leq \gamma} f'(x)$, C_1 is an absolute constant, and the infimum is taken over all estimators $\widehat{\mathbf{K}}$ of \mathbf{K} from $|\mathcal{S}|$ samples.

Importantly, up to polylogarithmic factors and constants, our minimax lower bound over the $\ell_{1,2}$ ball matches the upper bound over the nuclear norm ball given in Theorem 2.1. *In particular, in the worst case, learning a sparse and low rank matrix \mathbf{K} is no easier than learning a \mathbf{K} that is simply low rank.* However in many realistic cases, a slight performance gain is seen from optimizing over the $\ell_{1,2}$ ball when \mathbf{K}^* is row sparse, while optimizing over the nuclear norm ball does better when \mathbf{K}^* is dense. We show examples of this in the Section 3. The proof is given in the supplementary materials.

Note that if γ is in a bounded range, then the constant C has little effect. For the case that f is the logistic function, $C_f \geq \frac{1}{4} e^{-y_t \langle \mathbf{M}_t, \mathbf{K} \rangle} \geq \frac{1}{4} e^{-\gamma}$. Likewise, the term under the root will be also be bounded for γ in a constant range. The terms in the constant C arise when translating from risk and a KL-divergence to squared distance and reflects the noise in the problem.

2.3 Sample Complexity Bounds for Identification

Under a general loss function and arbitrary \mathbf{K}^* , we can not hope to convert our prediction error bounds into a recovery statement. However in this section we will show that as long as \mathbf{K}^* is low rank, and if we choose the loss function to be the log loss ℓ_f of a given link function f as defined prior to the statement of Theorem 2.3, recovery is possible. Firstly, note that under these assumptions we have an explicit formula for the risk,

$$R(\mathbf{K}) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle) \log \frac{1}{f(\langle \mathbf{M}_t, \mathbf{K} \rangle)} + (1 - f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle)) \log \frac{1}{1 - f(\langle \mathbf{M}_t, \mathbf{K} \rangle)}$$

and

$$R(\mathbf{K}) - R(\mathbf{K}^*) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} KL(f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle) \| f(\langle \mathbf{M}_t, \mathbf{K} \rangle)).$$

The following theorem shows that if the excess risk is small, i.e. $R(\widehat{\mathbf{K}})$ approximates $R(\mathbf{K}^*)$ well, then $\mathcal{M}(\widehat{\mathbf{K}})$ approximates $\mathcal{M}(\mathbf{K}^*)$ well. The proof, given in the supplementary materials, uses standard Taylor series arguments to show the KL-divergence is bounded below by squared-distance.

Lemma 2.4. *Let $C_f = \inf_{|x| \leq \gamma} f'(x)$. Then for any $\mathbf{K} \in \mathbf{K}_{\lambda, \gamma}$,*

$$\frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{M}(\mathbf{K}) - \mathcal{M}(\mathbf{K}^*)\|^2 \leq R(\mathbf{K}) - R(\mathbf{K}^*).$$

The following may give us hope that recovering \mathbf{K}^* from $\mathcal{M}(\mathbf{K}^*)$ is trivial, but the linear operator \mathcal{M} is non-invertible in general, as we discuss next. To see why, we must consider a more general class of operators defined on Gram matrices. Given a symmetric matrix \mathbf{G} , define the operator \mathbf{L}_t by

$$\mathbf{L}_t(\mathbf{G}) = 2\mathbf{G}_{ik} - 2\mathbf{G}_{ij} + \mathbf{G}_{jj} - \mathbf{G}_{kk}$$

If $\mathbf{G} = \mathbf{X}^T \mathbf{K} \mathbf{X}$ then $\mathbf{L}_t(\mathbf{G}) = \mathbf{M}_t(\mathbf{K})$, and more so $\mathbf{M}_t = \mathbf{X} \mathbf{L}_t \mathbf{X}^T$. Analogous to \mathcal{M} , we will combine the \mathbf{L}_t operators into a single operator \mathcal{L} ,

$$\mathcal{L}(\mathbf{G}) = (\mathbf{L}_t(\mathbf{G}) \mid \text{for } t \in \mathcal{T}) \in \mathbb{R}^{n \binom{n-1}{2}}.$$

Lemma 2.5. *The null space of \mathcal{L} is one dimensional, spanned by $\mathbf{V} = \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$.*

The proof is contained in the supplementary materials. In particular we see that \mathcal{M} is not invertible in general, adding a serious complication to our argument. However \mathcal{L} is still invertible on the subset of centered symmetric matrices orthogonal to \mathbf{V} , a fact that we will now exploit. We can decompose \mathbf{G} into \mathbf{V} and a component orthogonal to \mathbf{V} denoted \mathbf{H} ,

$$\mathbf{G} = \mathbf{H} + \sigma_{\mathbf{G}} \mathbf{V}$$

where $\sigma_{\mathbf{G}} := \frac{\langle \mathbf{G}, \mathbf{V} \rangle}{\|\mathbf{V}\|_F^2}$, and under the assumption that \mathbf{G} is centered, $\sigma_{\mathbf{G}} = \frac{\|\mathbf{G}\|_*}{n-1}$. Remarkably, the following lemma tells us that a *non-linear* function of \mathbf{H} uniquely determines \mathbf{G} .

Lemma 2.6. *If $n > d + 1$, and \mathbf{G} is rank d and centered, then $-\sigma_{\mathbf{G}}$ is an eigenvalue of \mathbf{H} with multiplicity $n - d - 1$. In addition, given another Gram matrix \mathbf{G}' of rank d' , $\sigma_{\mathbf{G}'} - \sigma_{\mathbf{G}}$ is an eigenvalue of $\mathbf{H} - \mathbf{H}'$ with multiplicity at least $n - d - d' - 1$.*

Proof. Since \mathbf{G} is centered, $\mathbf{1}_n \in \ker \mathbf{G}$, and in particular $\dim(\mathbf{1}_n^\perp \cap \ker \mathbf{G}) = n - d - 1$. If $x \in \mathbf{1}_n^\perp \cap \ker \mathbf{G}$, then

$$\mathbf{G}x = \mathbf{H}x + \sigma_{\mathbf{G}} \mathbf{V}x \Rightarrow \mathbf{H}x = -\sigma_{\mathbf{G}} x.$$

For the second statement, notice that $\dim(\mathbf{1}_n^\perp \cap \ker \mathbf{G} - \mathbf{G}') \geq n - d - d' - 1$. A similar argument then applies. \square

If $n > 2d$, then the multiplicity of the eigenvalue $-\sigma_{\mathbf{G}}$ is at least $n/2$. So we can trivially identify it from the spectrum of \mathbf{H} . This gives us a *non-linear* way to recover \mathbf{G} from \mathbf{H} .

Now we can return to the task of recovering \mathbf{K}^* from $\mathcal{M}(\widehat{\mathbf{K}})$. Indeed the above lemma implies that \mathbf{G}^* (and hence \mathbf{K}^* if \mathbf{X} is full rank) can be recovered from \mathbf{H}^* by computing an eigenvalue of \mathbf{H}^* . However \mathbf{H}^* is recoverable from $\mathcal{L}(\mathbf{H}^*)$, which is itself well approximated by $\mathcal{L}(\widehat{\mathbf{H}}) = \mathcal{M}(\widehat{\mathbf{K}})$. The proof of the following theorem makes this argument precise.

Theorem 2.7. *Assume that \mathbf{K}^* is rank d , $\widehat{\mathbf{K}}$ is rank d' , $n > d + d' + 1$, \mathbf{X} is rank p and $\mathbf{X}^T \mathbf{K}^* \mathbf{X}$ and $\mathbf{X}^T \widehat{\mathbf{K}} \mathbf{X}$ are all centered. Let $C_{d,d'} = \left(1 + \frac{n-1}{(n-d-d'-1)}\right)$. Then with probability at least $1 - \delta$,*

$$\frac{n \sigma_{\min}(\mathbf{X} \mathbf{X}^T)^2}{|\mathcal{T}|} \|\widehat{\mathbf{K}} - \mathbf{K}^*\|_F^2 \leq \frac{2LC_{d,d'}}{C_f^2} \left[\left(\sqrt{\frac{140\lambda^2 \|\mathbf{X} \mathbf{X}^T\| \log p}{|\mathcal{S}|}} + \frac{2 \log p}{|\mathcal{S}|} \right) + \sqrt{\frac{2L^2 \gamma^2 \log \frac{2}{\delta}}{|\mathcal{S}|}} \right]$$

where $\sigma_{\min}(\mathbf{X} \mathbf{X}^T)$ is the smallest eigenvalue of $\mathbf{X} \mathbf{X}^T$.

The proof, given in the supplementary materials, relies on two key components, Lemma 2.6 and a type of *restricted isometry property* for \mathcal{M} on \mathbf{V}^\perp . Our proof technique is a streamlined and more general approach similar to that used in the special case of ordinal embedding. In fact, our new bound improves on the recovery bound given in [11] for ordinal embedding.

We have several remarks about the bound in the theorem. If \mathbf{X} is well conditioned, e.g. isotropic, then $\sigma_{\min}(\mathbf{X}\mathbf{X}^T) \approx \frac{n}{p}$. In that case $\frac{n\sigma_{\min}(\mathbf{X}\mathbf{X}^T)^2}{|\mathcal{T}|} \approx \frac{1}{p^2}$, so the left hand side is the average squared error of the recovery. In most applications the rank of the empirical risk minimizer $\widehat{\mathbf{K}}$ is approximately equal to the rank of \mathbf{K}^* , i.e. $d \approx d'$. Note that if $d + d' \leq \frac{1}{2}(n - 1)$ then $C_{d,d'} \leq 3$. Finally, the assumption that $\mathbf{X}^T \mathbf{K}^* \mathbf{X}$ are centered can be guaranteed by centering \mathbf{X} , which has no impact on the triplet differences $\langle \mathbf{M}_t, \mathbf{K}^* \rangle$, or insisting that \mathbf{K}^* is centered. As mentioned above C_f will have little effect assuming that our measurements $\langle \mathbf{M}_t, \mathbf{K} \rangle$ are bounded.

2.4 Applications to Ordinal Embedding

In the ordinal embedding setting, there are a set of items with unknown locations, $\mathbf{z}_1, \dots, \mathbf{z}_n \in \mathbb{R}^d$ and a set of triplet observations \mathcal{S} where as in the metric learning case observing $y_t = -1$, for a triplet $t = (i, j, k)$ is indicative of the $\|\mathbf{z}_i - \mathbf{z}_j\|^2 \leq \|\mathbf{z}_i - \mathbf{z}_k\|^2$, i.e. item i is closer to j than k . The goal is to recover the \mathbf{z}_i 's, up to rigid motions, by recovering their Gram matrix \mathbf{G}^* from these comparisons. Ordinal embedding case reduces to metric learning through the following observation. Consider the case when $n = p$ and $\mathbf{X} = \mathbf{I}_p$, i.e. the \mathbf{x}_i are standard basis vectors. Letting $\mathbf{K}^* = \mathbf{G}^*$, we see that $\|\mathbf{x}_i - \mathbf{x}_j\|_{\mathbf{K}}^2 = \|\mathbf{z}_i - \mathbf{z}_j\|^2$. So in particular, $\mathbf{L}_t = \mathbf{M}_t$ for each triple t , and observations are exactly comparative distance judgements. Our results then apply, and extend previous work on sample complexity in the ordinal embedding setting given in [11]. In particular, though Theorem 5 in [11] provides a consistency guarantee that the empirical risk minimizer $\widehat{\mathbf{G}}$ will converge to \mathbf{G}^* , they do not provide a convergence rate. We resolve this issue now.

In their work, it is assumed that $\|\mathbf{z}_i\|^2 \leq \gamma$ and $\|\mathbf{G}\|_* \leq \sqrt{dn}\gamma$. In particular, sample complexity results of the form $O(dn\gamma \log n)$ are obtained. However, these results are trivial in the following sense, if $\|\mathbf{z}_i\|^2 \leq \gamma$ then $\|\mathbf{G}\|_* \leq \gamma n$, and their results (as well as our upper bound) implies that true sample complexity is significantly smaller, namely $O(\gamma n \log n)$ which is independent of the ambient dimension d . As before, assume an explicit link function f with Lipschitz constant L , so the samples are noisy observations governed by \mathbf{G}^* , and take the loss to be the logarithmic loss associated to f .

We obtain the following improved recovery bound in this case. The proof is immediate from Theorem 2.7.

Corollary 2.7.1. *Let \mathbf{G}^* be the Gram matrix of n centered points in d dimensions with $\|\mathbf{G}^*\|_F^2 = \frac{\gamma^2 n^2}{d}$. Let $\widehat{\mathbf{G}} = \min_{\|\mathbf{G}\|_* \leq \gamma n, \|\mathbf{G}\|_\infty \leq \gamma} R_{\mathcal{S}}(\mathbf{G})$ and assume that $\widehat{\mathbf{G}}$ is rank d , with $n > 2d + 1$. Then,*

$$\frac{\|\widehat{\mathbf{G}} - \mathbf{G}^*\|_F^2}{n^2} = O\left(\frac{LC_{d,d}}{C_f^2} \sqrt{\frac{\gamma n \log n}{|\mathcal{S}|}}\right)$$

3 Experiments

To validate our complexity and recovery guarantees, we ran the following simulations. We generate $\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \frac{1}{p}\mathbf{I})$, with $n = 200$, and $\mathbf{K}^* = \frac{p}{\sqrt{d}}\mathbf{U}\mathbf{U}^T$ for a random orthogonal matrix $\mathbf{U} \in \mathbb{R}^{p \times d}$ with unit norm columns. In Figure 2a, \mathbf{K}^* has d nonzero rows/columns. In Figure 2b, \mathbf{K}^* is a dense rank- d matrix. We compare the performance of nuclear norm and $\ell_{1,2}$ regularization in each setting against an unconstrained baseline where we only enforce that \mathbf{K} be psd. Given a fixed number of samples, each method is compared in terms of the relative excess risk, $\frac{R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*)}{R(\mathbf{K}^*)}$, and the relative squared recovery error, $\frac{\|\widehat{\mathbf{K}} - \mathbf{K}^*\|_F^2}{\|\mathbf{K}^*\|_F^2}$, averaged over 20 trials. The y-axes of both plots have been trimmed for readability.

In the case that \mathbf{K}^* is sparse, $\ell_{1,2}$ regularization outperforms nuclear norm regularization. However, in the case of dense low rank matrices, nuclear norm regularization is superior. Notably, as expected from our upper and lower bounds, the performances of the two approaches seem to be within constant

factors of each other. Therefore, unless there is strong reason to believe that the underlying \mathbf{K}^* is sparse, nuclear norm regularization achieves comparable performance with a less restrictive modeling assumption. Furthermore, in the two settings, both the nuclear norm and $\ell_{1,2}$ constrained methods outperform the unconstrained baseline, especially in the case where \mathbf{K}^* is low rank and sparse.

To empirically validate our sample complexity results, we compute the number of samples averaged over 20 runs to achieve a relative excess risk of less than 0.1 in Figure 3. First, we fix $p = 100$ and increment d from 1 to 10. Then we fix $d = 10$ and increment p from 10 to 100 to clearly show the linear dependence of the sample complexity on d and p as demonstrated in Corollary 2.1.1. To our knowledge, these are the first results quantifying the sample complexity in terms of the number of features, p , and the embedding dimension, d .

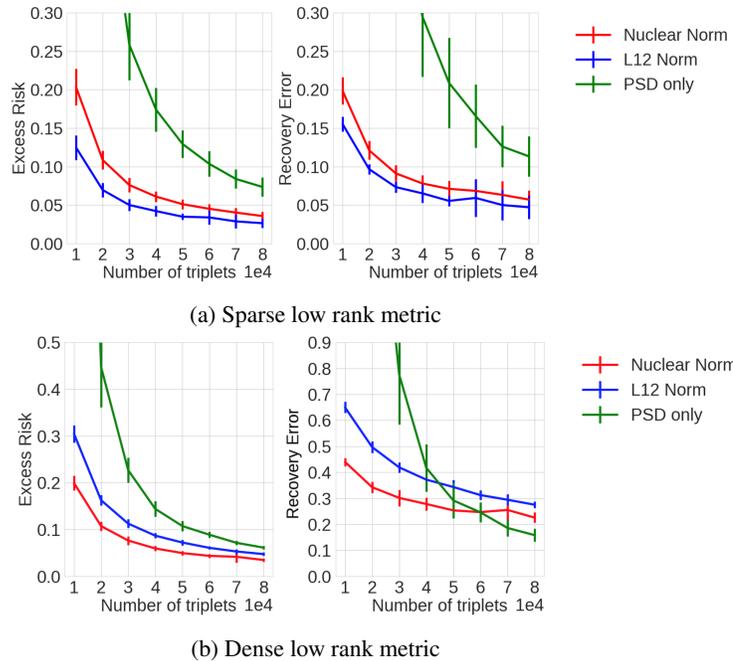


Figure 2: $\ell_{1,2}$ and nuclear norm regularization performance

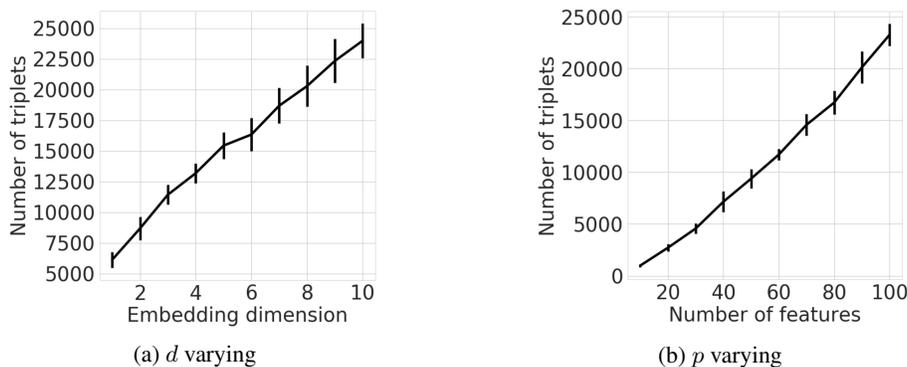


Figure 3: Number of samples to achieve relative excess risk < 0.1

Acknowledgments This work was partially supported by the NSF grants CCF-1218189 and IIS-1623605

References

- [1] Martina A Rau, Blake Mason, and Robert D Nowak. How to model implicit knowledge? similarity learning methods to assess perceptions of visual representations. In *Proceedings of the 9th International Conference on Educational Data Mining*, pages 199–206, 2016.
- [2] Aurélien Bellet, Amaury Habrard, and Marc Sebban. Metric learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 9(1):1–151, 2015.
- [3] Aurélien Bellet and Amaury Habrard. Robustness and generalization for metric learning. *Neurocomputing*, 151:259–267, 2015.
- [4] Zheng-Chu Guo and Yiming Ying. Guaranteed classification via regularized similarity learning. *Neural Computation*, 26(3):497–522, 2014.
- [5] Yiming Ying, Kaizhu Huang, and Colin Campbell. Sparse metric learning via smooth optimization. In *Advances in neural information processing systems*, pages 2214–2222, 2009.
- [6] Wei Bian and Dacheng Tao. Constrained empirical risk minimization framework for distance metric learning. *IEEE transactions on neural networks and learning systems*, 23(8):1194–1205, 2012.
- [7] Yuan Shi, Aurélien Bellet, and Fei Sha. Sparse compositional metric learning. *arXiv preprint arXiv:1404.4105*, 2014.
- [8] Samet Oymak, Amin Jalali, Maryam Fazel, Yonina C Eldar, and Babak Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Transactions on Information Theory*, 61(5):2886–2908, 2015.
- [9] Eric Heim, Matthew Berger, Lee Seversky, and Milos Hauskrecht. Active perceptual similarity modeling with auxiliary information. *arXiv preprint arXiv:1511.02254*, 2015.
- [10] Kenneth R Davidson and Stanislaw J Szarek. Local operator theory, random matrices and banach spaces. *Handbook of the geometry of Banach spaces*, 1(317-366):131, 2001.
- [11] Lalit Jain, Kevin G Jamieson, and Rob Nowak. Finite sample prediction and recovery bounds for ordinal embedding. In *Advances In Neural Information Processing Systems*, pages 2703–2711, 2016.
- [12] Mark A Davenport, Yaniv Plan, Ewout Van Den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3):189–223, 2014.
- [13] Joel A. Tropp. An introduction to matrix concentration inequalities, 2015.
- [14] Felix Abramovich and Vadim Grinshtein. Model selection and minimax estimation in generalized linear models. *IEEE Transactions on Information Theory*, 62(6):3721–3730, 2016.
- [15] Florentina Bunea, Alexandre B Tsybakov, Marten H Wegkamp, et al. Aggregation for gaussian regression. *The Annals of Statistics*, 35(4):1674–1697, 2007.
- [16] Philippe Rigollet and Alexandre Tsybakov. Exponential screening and optimal rates of sparse estimation. *The Annals of Statistics*, pages 731–771, 2011.
- [17] Jon Dattorro. *Convex Optimization & Euclidean Distance Geometry*. Meboo Publishing USA, 2011.

4 Supplementary Materials

5 Proof of Results

5.1 Proof of Theorem 2.1

Our argument follows standard statistical learning theory techniques used in the classification literature. This framework is also similar to that used in the one bit matrix completion literature, see [12]. The main ingredient in the proof is the use of a Matrix Bernstein to bound the Rademacher complexity of our class.

By the Bounded Difference inequality,

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

where $\beta = \sup_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} |\ell((y_t \langle \mathbf{M}_t, \mathbf{K} \rangle)) - \ell((y_{t'} \langle \mathbf{M}_{t'}, \mathbf{K} \rangle))| \leq L\gamma$ since $\langle \mathbf{M}_t, \mathbf{K} \rangle \leq \gamma$. Using standard symmetrization and contraction lemmas, we can introduce Rademacher random variables $\varepsilon_t \in \{-1, 1\}$ for all $t \in \mathcal{T}$ so that

$$\begin{aligned} \mathbb{E} \left[\sup_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} |\widehat{R}(\mathbf{K}) - R(\mathbf{K})| \right] &\leq \mathbb{E} \frac{2L}{|\mathcal{T}|} \sup_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} \left| \sum_{t \in \mathcal{S}} \varepsilon_t \langle \mathbf{M}_t, \mathbf{K} \rangle \right| \\ &\leq \mathbb{E} \frac{2L}{|\mathcal{S}|} \sup_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} \left\| \sum_{t \in \mathcal{S}} \varepsilon_t \mathbf{M}_t \right\| \|\mathbf{K}\|_* \\ &\leq \mathbb{E} \frac{2L\lambda}{|\mathcal{S}|} \sup_{\mathbf{K} \in \mathcal{K}_{\lambda, \gamma}} \left\| \sum_{t \in \mathcal{S}} \varepsilon_t \mathbf{M}_t \right\| \end{aligned}$$

We employ a matrix Bernstein bound, Theorem 6.6.1 in [13], to compute

$$\mathbb{E} \left\| \sum_{t \in \mathcal{S}} \varepsilon_t \mathbf{M}_t \right\| \leq \sqrt{140 \frac{\|\mathbf{X}\mathbf{X}^T\|}{n} |\mathcal{S}| \log p} + 2 \log p.$$

To see this, it suffices to bound $\left\| \sum_{t \in \mathcal{T}} \mathbf{M}_t^2 \right\|$ which is done in Lemma 5.1. Plugging this in above gives

$$\mathbb{E} \frac{2L\lambda}{|\mathcal{S}|} \left\| \sum_{t \in \mathcal{S}} \varepsilon_t \mathbf{M}_t \right\| \leq 2L \left(\sqrt{\frac{140\lambda^2 \|\mathbf{X}\mathbf{X}^T\| \log p}{|\mathcal{S}|}} + \frac{2 \log p}{|\mathcal{S}|} \right)$$

Lemma 5.1.

$$\frac{1}{n \binom{n-1}{2}} \left\| \sum_{t \in \mathcal{T}} \mathbf{M}_t^2 \right\| \leq 70 \frac{\|\mathbf{X}\mathbf{X}^T\|}{n}$$

Proof. Let \mathbf{e}_i be the i^{th} standard basis vector. For a triplet $t = (i, j, k)$, define

$$\mathbf{L}_t = \mathbf{e}_i \mathbf{e}_k^T + \mathbf{e}_k \mathbf{e}_i^T - \mathbf{e}_i \mathbf{e}_j^T - \mathbf{e}_j \mathbf{e}_i^T + \mathbf{e}_j \mathbf{e}_j^T - \mathbf{e}_k \mathbf{e}_k^T$$

(in particular \mathbf{L}_t is the matrix corresponding to the operator L_t given in Section 2.3). A computation shows that $\langle \mathbf{L}_t, \mathbf{X}^T \mathbf{K} \mathbf{X} \rangle = \langle \mathbf{M}_t, \mathbf{K} \rangle$ and moreover $\mathbf{M}_t = \mathbf{X} \mathbf{L}_t \mathbf{X}^T$. By definition,

$$\begin{aligned} \sum_{t \in \mathcal{T}} \mathbf{M}_t^2 &= \sum_{t \in \mathcal{T}} \mathbf{X} \mathbf{L}_t \mathbf{X}^T \mathbf{X} \mathbf{L}_t \mathbf{X} \\ &= \mathbf{X} \left(\sum_{t \in \mathcal{T}} \mathbf{L}_t \mathbf{X}^T \mathbf{X} \mathbf{L}_t \right) \mathbf{X}^T \end{aligned}$$

We now focus our attention on simplifying the middle term. Firstly, note that we can assume that the \mathbf{X} 's are centered, i.e. $\mathbf{X}\mathbf{1}_n = \mathbf{0}$. To see this, note that the \mathbf{L}_t 's are centered so in particular, $\mathbf{L}_t\mathbf{V} = \mathbf{L}_t$. Then

$$\mathbf{L}_t\mathbf{X}^T\mathbf{X}\mathbf{L}_t = \mathbf{L}_t\mathbf{V}^T\mathbf{X}^T\mathbf{X}\mathbf{V}\mathbf{L}_t = \mathbf{L}_t(\mathbf{X}\mathbf{V})^T(\mathbf{X}\mathbf{V})\mathbf{L}_t$$

so we can replace \mathbf{X} with $\mathbf{X}\mathbf{V}$, i.e. we can center \mathbf{X} . Also note that centering \mathbf{X} only diminishes the operator norm $\mathbf{X}\mathbf{X}^T$, so centering does not affect the statement of the bound, and furthermore a tighter statement is certainly possible by assuming that \mathbf{X} is centered.

Using the reduction to a centered \mathbf{X} , a computation (omitted due to length) shows that

$$\left(\sum_{t \in \mathcal{T}} \mathbf{L}_t\mathbf{X}^T\mathbf{X}\mathbf{L}_t\right)_{i,j} = \begin{cases} (2n-3)\|\mathbf{X}^T\mathbf{X}\|_* + (n^2-3n)\|\mathbf{x}_i\|^2 & i=j \\ (n-4)\langle \mathbf{x}_i, \mathbf{x}_j \rangle - (n-2)\|\mathbf{x}_j\|^2 - (n-2)\|\mathbf{x}_i\|^2 - \|\mathbf{X}^T\mathbf{X}\|_* & i \neq j \end{cases}$$

To bound $\|\sum_{t \in \mathcal{T}} \mathbf{L}_t\mathbf{X}^T\mathbf{X}\mathbf{L}_t\| \leq 7n^2$, by Gershgorin's Circle Theorem we just have to bound the sums of the absolute values of the entries in each row. This ends up being,

$$\begin{aligned} & (2n-3+n-1)\|\mathbf{X}^T\mathbf{X}\|_* + (n^2-3n+(n-1)(n-2))\|\mathbf{x}_i\|^2 + (n-2)\sum_{i \neq j} \|\mathbf{x}_j\|^2 \\ & \quad + (n-4)\sum_{i \neq j} |\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \\ & \leq (2n-3+n-1+n-2)\|\mathbf{X}^T\mathbf{X}\|_* + (n^2-3n+(n-1)(n-2)-2)\|\mathbf{x}_i\|^2 \\ & \quad + (n-4)\sum_j \|\mathbf{x}_i\|\|\mathbf{x}_j\| \\ & \leq (4n-6)\|\mathbf{X}^T\mathbf{X}\|_* + (2n^2-6n)\|\mathbf{x}_i\|^2 + n(n-4)\max_j \|\mathbf{x}_j\|^2 \\ & \leq (4n-6)n\max_j \|\mathbf{x}_j\|^2 + (2n^2-6n)\max_j \|\mathbf{x}_j\|^2 + (n^2-4n)\max_j \|\mathbf{x}_j\|^2 \\ & \leq 7n^2\max_j \|\mathbf{x}_j\|^2 \end{aligned}$$

So $\|\sum_{t \in \mathcal{T}} \mathbf{L}_t\mathbf{X}^T\mathbf{X}\mathbf{L}_t\| \leq 7n^2$ and

$$\frac{1}{n\binom{n-1}{2}} \left\| \sum_{t \in \mathcal{T}} \mathbf{X}\mathbf{L}_t\mathbf{X}^T\mathbf{X}\mathbf{L}_t\mathbf{X} \right\| \leq 70 \frac{\|\mathbf{X}\mathbf{X}^T\|}{n}$$

using the fact that $\frac{2n^2}{(n-1)(n-2)} \leq 10$ for positive $n \geq 3$. □

5.2 Proof of Theorem 2.3

We will need the following lemma relating the KL-divergence to squared distance in this section and in the proof of Theorem 2.7.

Lemma 5.2. *Let $y, z \in (0, 1)$, then*

$$2(z-y)^2 \leq KL(z||y) \leq \frac{(z-y)^2/2}{\inf_{x \in (0,1)} x(1-x)}$$

Proof. For $y, z \in (0, 1)$ let $g(z) = z \log \frac{z}{y} + (1-z) \log \frac{1-z}{1-y}$. Then $g'(z) = \log \frac{z}{1-z} - \log \frac{y}{1-y}$ and $g''(z) = \frac{1}{z(1-z)}$. By Taylor's theorem, for some η in the interval between y and z , $g(z) = \frac{g''(\eta)}{2}(z-y)^2$. So for a lower bound,

$$g(z) \geq \frac{(z-y)^2/2}{\sup_{x \in (0,1)} x(1-x)} \geq 2(z-y)^2.$$

Similarly an upper bound is given by,

$$g(z) \leq \frac{(z-y)^2/2}{\inf_{x \in (0,1)} x(1-x)}$$

□

Now we resume the proof of Theorem 2.3. Fix $\mathbf{X} = \mathbf{I}$. Given triplet comparisons generated according to \mathbf{K} , we are interested in finding the minimax lower bound,

$$\inf_{\widehat{\mathbf{K}}} \sup_{\mathbf{K} \in \mathcal{K}'_{\lambda, \gamma}} \mathbb{E}[R(\widehat{\mathbf{K}})] - R(\mathbf{K})$$

Where as previously computed in Section 2.3

$$R(\widehat{\mathbf{K}}) - R(\mathbf{K}) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} f(\langle \mathbf{M}_t, \mathbf{K} \rangle) \log \frac{f(\langle \mathbf{M}_t, \mathbf{K} \rangle)}{f(\langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)} + (1 - f(\langle \mathbf{M}_t, \mathbf{K} \rangle)) \log \frac{1 - f(\langle \mathbf{M}_t, \mathbf{K} \rangle)}{1 - f(\langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)}$$

Lemma 2.4 implies,

$$R(\widehat{\mathbf{K}}) - R(\mathbf{K}) \geq \frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{M}(\mathbf{K}) - \mathcal{M}(\widehat{\mathbf{K}})\|_2^2.$$

where $C_f = \inf_{|x| \leq \gamma} f'(x)$. We will construct a set $\kappa \subset \mathcal{K}'_{\lambda, \gamma}$ so that for any two $\mathbf{K}^1, \mathbf{K}^2 \in \kappa$, with $\mathbf{K}^1 \neq \mathbf{K}^2$,

- $\frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{M}(\mathbf{K}^1) - \mathcal{M}(\mathbf{K}^2)\|_F^2 \geq 4s_n^2$, for $\mathbf{K}^1 \neq \mathbf{K}^2$
- Let $P_{\mathbf{K}}^S$ denote the sample distribution of a set of $|\mathcal{S}|$ samples conditioned on it being drawn from $\mathbf{K} \in \kappa$. Then we also require $KL(P_{\mathbf{K}^1}^S \| P_{\mathbf{K}^2}^S) \leq \frac{1}{16} \ln |\kappa|$

Following an argument similar to the proof of Theorem 2 in [14], it will then follow from a variant of Fano's inequality, namely Lemma A.1 from [15], that

$$\inf_{\widehat{\mathbf{K}}} \sup_{\mathbf{K} \in \mathcal{K}'_{\lambda, \gamma}} \mathbb{E}[R(\widehat{\mathbf{K}})] - R(\mathbf{K}) \geq s_n^2.$$

By Lemma 8.3 of [16], there exists a subset $\kappa \subset \mathcal{K}'_{\lambda, \gamma}$, and an absolute constant $0 < C_1 < 1$ such that

- $\ln |\kappa| \geq C_1 d \ln \frac{p}{d}$
- Each element of κ has sparsity d , is 0 away from the diagonal, and on the diagonal the elements are either 0 or γ , for a value of $\gamma \geq 0$ we will choose later.
- For all $\mathbf{K}^i, \mathbf{K}^j \in \kappa$, $\|\mathbf{K}^i - \mathbf{K}^j\|_0 \geq C_1 d$.

Therefore, for $\mathbf{K}^1, \mathbf{K}^2 \in \kappa$, we need only to show $KL(\mathbf{K}^1 \| \mathbf{K}^2) \leq \frac{1}{16} \ln |\kappa|$. Using the fact that $\mathbf{X} = \mathbf{I}$,

$$\begin{aligned} \frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{M}(\mathbf{K}^1) - \mathcal{M}(\mathbf{K}^2)\|_2^2 &\geq \frac{2C_f^2}{|\mathcal{T}|} p \sum_{j < k} ((\mathbf{K}_{kk}^1 - \mathbf{K}_{kk}^2) - (\mathbf{K}_{jj}^1 - \mathbf{K}_{jj}^2))^2 \\ &\geq \frac{2C_f^2 C_1 p d (p - 2d) \gamma^2}{|\mathcal{T}|} \end{aligned}$$

To see the second to last inequality, note that there are at least $C_1 d (p - 2d)$ pairs of indices j, k where $\mathbf{K}_{kk}^1 \neq \mathbf{K}_{kk}^2$ but $\mathbf{K}_{jj}^1 = \mathbf{K}_{jj}^2$, because \mathbf{K}^1 and \mathbf{K}^2 share at least $p - 2d$ entries on their diagonal that are both 0. Each such entry contributes a γ^2 to the sum.

In particular choose,

$$s_n^2 = \frac{C_f^2 C_1 p d (p - 2d) \gamma^2}{2|\mathcal{T}|}.$$

We proceed by selecting γ such that $KL(P_{\mathbf{K}^1}^{\mathcal{S}} || P_{\mathbf{K}^2}^{\mathcal{S}}) \leq \frac{1}{16} \ln |\kappa|$. Assume our samples are $\mathcal{S} = \{(t, y_t)\}$. Then since the samples are i.i.d.

$$KL(P_{\mathbf{K}^1}^{\mathcal{S}} || P_{\mathbf{K}^2}^{\mathcal{S}}) = \sum_{t \in \mathcal{S}} KL(P_{\mathbf{K}^1}(t) || P_{\mathbf{K}^2}(t))$$

where $P_{\mathbf{K}^i}(t)$ is the distribution of y_t conditioned on \mathbf{K}^i , in particular the probability of $y_t = -1$ is $f(\langle \mathbf{M}_t, \mathbf{K}^i \rangle)$.

We can bound each term of the sum above using the upper bound from Lemma 5.2.

$$\begin{aligned} KL(P_{\mathbf{K}^1}(t) || P_{\mathbf{K}^2}(t)) &\leq \frac{(f(\langle \mathbf{M}_t, \mathbf{K}^1 \rangle) - f(\langle \mathbf{M}_t, \mathbf{K}^2 \rangle))^2}{2 \inf_t f(\langle \mathbf{M}_t, \mathbf{K}^2 \rangle)(1 - f(\langle \mathbf{M}_t, \mathbf{K}^2 \rangle))} \\ &\leq \frac{(\langle \mathbf{M}_t, \mathbf{K}^1 - \mathbf{K}^2 \rangle)^2 \sup_{|\nu| \leq \gamma} f'(\nu)^2}{2 \inf_{|x| \leq \gamma} f(x)(1 - f(x))} \\ &\leq \frac{\gamma^2 \sup_{|\nu| \leq \gamma} f'(\nu)^2}{2 \inf_{|x| \leq \gamma} f(x)(1 - f(x))} \end{aligned}$$

Summing over $t \in \mathcal{S}$, we require that

$$KL(P_{\mathbf{K}^1}^{\mathcal{S}} || P_{\mathbf{K}^2}^{\mathcal{S}}) \leq \frac{\gamma^2 |\mathcal{S}| \sup_{|\nu| \leq \gamma} f'(\nu)^2}{2 \inf_{|x| \leq \gamma} f(x)(1 - f(x))} \leq \frac{C_1}{16} d \ln \frac{p}{d} \leq \frac{1}{16} \ln |\kappa|,$$

so in particular, we will take

$$\frac{\gamma^2 \sup_{|\nu| \leq \gamma} f'(\nu)^2}{2 \inf_{|x| \leq \gamma} f(x)(1 - f(x))} = \frac{C_1}{16 |\mathcal{S}|} d \ln \frac{p}{d}$$

From this point on, let's take $\lambda = p$, and $d = \frac{p}{4}$. Now we have a few additional constraints on γ ,

- Since $\|\mathbf{K}^i\|_{1,2} \leq \lambda$ for each $\mathbf{K}^i \in \kappa$, we require $\gamma d \leq \lambda$, so in particular $\gamma \leq 4$.
- In addition, we are going to require $\gamma \geq 1$ since we will need $p\gamma \geq \lambda$ (used below).

Based on these conditions, we just take $\gamma = 2$ and after simplification choose,

$$|\mathcal{S}| := \frac{C_1 p \ln 4 \inf_{|x| \leq \gamma} f(x)(1 - f(x))}{32 \gamma^2 \sup_{|\nu| \leq \gamma} f'(\nu)^2}$$

Now we are finally in a position to use our choice of γ, d, λ and $|\mathcal{S}|$. We see that

$$\begin{aligned} s_n^2 &= \frac{C_f^2 C_1 p d (p - 2d) \gamma^2}{2 |\mathcal{T}|} = \frac{C_f^2 C_1 p^2 \gamma \lambda}{16 |\mathcal{T}|} && \text{(since } p\gamma \geq \lambda) \\ &\geq \frac{C_f^2 C_1 \gamma \lambda}{8p} \\ &\geq \frac{C_f^2 \sqrt{\inf_{|x| \leq \gamma} f(x)(1 - f(x))}}{8p \sqrt{\sup_{|\nu| \leq 2} f'(\nu)^2}} \sqrt{\frac{C_1^3 \ln 4}{32} \frac{p}{|\mathcal{S}|}} \\ &= \frac{C_f^2}{32} \sqrt{\frac{\inf_{|x| \leq \gamma} f(x)(1 - f(x))}{\sup_{|\nu| \leq \gamma} f'(\nu)^2}} \sqrt{\frac{C_1^3 \ln 4}{2} \frac{\lambda^2 \|\mathbf{X} \mathbf{X}^T\|}{n |\mathcal{S}|}} \end{aligned}$$

where the final equality follows from the fact that we have chosen $\mathbf{X} = \mathbf{I}_n$ so $n = p$. \square

5.3 Proof of Lemma 2.4

Proof of Lemma 2.4. As computed prior to the statement of Theorem 2.7.

$$R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*) = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} KL(f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle) || f(\langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle))$$

Now using Lemma 5.2 with $z = f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle)$ and $y = f(\langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)$ we see

$$KL(f(\langle \mathbf{M}_t, \mathbf{K}^* \rangle) \| f(\langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)) \geq 2C_f^2 (\langle \mathbf{M}_t, \mathbf{K}^* \rangle - \langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)^2$$

Summing over all $t \in \mathcal{T}$

$$\begin{aligned} R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*) &\geq \frac{2C_f^2}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\langle \mathbf{M}_t, \mathbf{K}^* \rangle - \langle \mathbf{M}_t, \widehat{\mathbf{K}} \rangle)^2 \\ &= \frac{2C_f^2}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\langle \mathbf{M}_t, \widehat{\mathbf{K}} - \mathbf{K}^* \rangle)^2 = \frac{2C_f^2}{|\mathcal{T}|} \|\mathcal{M}(\widehat{\mathbf{K}}) - \mathcal{M}(\mathbf{K}^*)\|_2^2. \end{aligned}$$

□

5.4 Proof of Theorem 2.7

Before launching into the proof of Theorem 2.7, we first prove an auxiliary set of results that depend on the classical correspondence between centered Gram matrices and Euclidean distance matrices. For a more in depth discussion of this correspondence, we refer interested readers to [17]. Let \mathbb{S}_h^n be the subspace of symmetric hollow matrices, i.e. symmetric matrices with zero diagonal, and let \mathbb{S}_c^n be the subspace of centered Gram matrices, i.e. positive semi-definite matrices with $\mathbf{1}_n$ in their kernel.

Note that $\dim \mathbb{S}_h^n = \dim \mathbb{S}_c^n = \binom{n}{2}$. In fact these spaces are isomorphic with an explicit linear isomorphism given by the maps

$$\mathbb{S}_h^n \rightarrow \mathbb{S}_c^n : \mathbf{D} \rightarrow -\frac{1}{2} \mathbf{V} \mathbf{D} \mathbf{V}$$

with inverse

$$\mathbb{S}_c^n \rightarrow \mathbb{S}_h^n : \mathbf{G} \rightarrow \text{diag}(\mathbf{G}) \mathbf{1}_n^T - 2\mathbf{G} + \mathbf{1}_n \text{diag}(\mathbf{G})^T$$

where again, $\mathbf{V} = \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$.

Given a set of centered points $\mathbf{X} \in \mathbb{R}^p$, then under the isomorphism above, the associated Gram matrix $\mathbf{G} \in \mathbb{S}_c^n$ maps to the squared distance matrix $\mathbf{D} \in \mathbb{S}_h^n$. In particular, a matrix in \mathbb{S}_h^n is a valid Euclidean distance matrix if and only if $-\frac{1}{2} \mathbf{V} \mathbf{D} \mathbf{V}$ is a centered Gram matrix.

Given a triplet $t = (i, j, k) \in \mathcal{T}$, we can define an operator $\Delta_t(\mathbf{D}) := \mathbf{D}_{ij} - \mathbf{D}_{ik}$ and

$$\Delta(\mathbf{D}) := (\Delta_t(\mathbf{D}) \mid \text{for } t \in \mathcal{T})$$

analogous to \mathcal{L} and \mathcal{M} . In particular, for associated \mathbf{D} and \mathbf{G} , $\Delta_t(\mathbf{D}) = \mathcal{L}_t(\mathbf{G})$ for all t so $\Delta(\mathbf{D}) = \mathcal{L}(\mathbf{G})$. We can now prove the key lemmas used in the proof of 2.7.

Lemma 5.3. *The null space of \mathcal{L} is one dimensional, spanned by \mathbf{V} .*

Proof. Lemma 2 in [11] shows $\ker \Delta$ is one dimensional and is spanned by $\mathbf{J} = \mathbf{1}_n \mathbf{1}_n^T - \mathbf{I}_n$. A computation shows that $-\frac{1}{2} \mathbf{V} \mathbf{J} \mathbf{V} = \frac{1}{2} \mathbf{V}$. Since $\mathcal{L}(\mathbf{V}) = \Delta(\mathbf{J}) = \mathbf{0}$, \mathbf{V} spans $\ker \mathcal{L}$. □

We rely on an analogous statement for distance matrices given in Lemma 3 in [11].

Lemma 5.4. *Let $\mathbf{G} \in \mathbb{S}_c^n$ and \mathbf{H} the component of \mathbf{G} orthogonal \mathbf{V} then $\|\mathcal{L}(\mathbf{H})\|^2 \geq n \|\mathbf{H}\|_F^2$.*

Proof. Again, let \mathbf{D} be the symmetric hollow matrix corresponding to \mathbf{G} . We can take a decomposition of \mathbf{D} into a component perpendicular to $\ker \Delta$

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

Applying $-\frac{1}{2} \mathbf{V} \cdot \mathbf{V}$ to both sides we get,

$$\mathbf{G} = -\frac{1}{2} \mathbf{V} \mathbf{C} \mathbf{V} + \frac{\sigma_D}{2} \mathbf{V}.$$

We claim that $\mathbf{H} = -\frac{1}{2} \mathbf{V} \mathbf{C} \mathbf{V}$ and $\sigma_G = \sigma_D/2$. It suffices to prove that $\mathbf{V} \mathbf{C} \mathbf{V}$ is perpendicular to \mathbf{V} . To see this note that $\langle \mathbf{V} \mathbf{C} \mathbf{V}, \mathbf{V} \rangle = \langle \mathbf{C}, \mathbf{V} \rangle = 0$, since \mathbf{C} is hollow and perpendicular to \mathbf{J} .

We now apply Lemma 3 in [11] which shows that the minimal eigenvalue of Δ is n .

$$\begin{aligned}
\|\mathcal{L}(\mathbf{H})\|^2 &= \|\Delta(\mathbf{C})\|^2 \\
&\geq n\|\mathbf{C}\|_F^2 && \text{(since } \mathbf{C} \text{ is perpendicular to the kernel of } \Delta) \\
&\geq n\left\|\frac{1}{2}\mathbf{V}\mathbf{C}\mathbf{V}\right\|_F^2 && \text{(Since } \mathbf{V} \text{ is a projection.)} \\
&\geq n\|\mathbf{H}\|_F^2
\end{aligned}$$

□

Proof of Theorem 2.7. We begin by applying Lemma 2.6 in the specific case where $\mathbf{G}^* = \mathbf{X}^T \mathbf{K}^* \mathbf{X}$ and $\widehat{\mathbf{G}} = \mathbf{X}^T \widehat{\mathbf{K}} \mathbf{X}$ with \mathbf{H}^* and $\widehat{\mathbf{H}}$ defined analogously to above. Firstly, by definition

$$\widehat{\mathbf{G}} - \mathbf{G}^* = \widehat{\mathbf{H}} - \mathbf{H}^* + (\sigma_{\widehat{\mathbf{G}}} - \sigma_{\mathbf{G}^*})\mathbf{V}$$

By orthogonality

$$\begin{aligned}
\|\widehat{\mathbf{G}} - \mathbf{G}^*\|_F^2 &= \|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2 + (\sigma_{\widehat{\mathbf{G}}} - \sigma_{\mathbf{G}^*})^2 \|\mathbf{V}\|_F^2 \\
&= \|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2 + (n-1)(\sigma_{\widehat{\mathbf{G}}} - \sigma_{\mathbf{G}^*})^2 && \text{(Since } \|\mathbf{V}\|_F^2 = n-1) \\
&\leq \|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2 + \frac{n-1}{(n-d-d'-1)} \|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2 \\
&\text{(By Lemma 2.6 } \sigma_{\widehat{\mathbf{G}}} - \sigma_{\mathbf{G}^*} \text{ is a repeated eigenvalue with multiplicity } n-d-d'-1) \\
&= C_{d,d'} \|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2.
\end{aligned}$$

Now,

$$\begin{aligned}
\|\mathcal{M}(\widehat{\mathbf{K}}) - \mathcal{M}(\mathbf{K}^*)\|_2^2 &= \|\mathcal{L}(\mathbf{X}^T \widehat{\mathbf{K}} \mathbf{X}) - \mathcal{L}(\mathbf{X}^T \mathbf{K}^* \mathbf{X})\|^2 \\
&\geq n\|\widehat{\mathbf{H}} - \mathbf{H}^*\|_F^2 && \text{(Using Lemma 5.4)} \\
&\geq \|\widehat{\mathbf{G}} - \mathbf{G}^*\|_F^2 && \text{(From the above.)} \\
&= \frac{n}{C_{d,d'}} \|\mathbf{X}^T \widehat{\mathbf{K}} \mathbf{X} - \mathbf{X}^T \mathbf{K}^* \mathbf{X}\|_F^2 \\
&\geq \frac{n\sigma_{\min}(\mathbf{X}\mathbf{X}^T)^2}{C_{d,d'}} \|\widehat{\mathbf{K}} - \mathbf{K}^*\|_F^2
\end{aligned}$$

To see the last line, recall $\text{vec}(\mathbf{X}^T \mathbf{K} \mathbf{X}) = (\mathbf{X}^T \otimes \mathbf{X}^T) \text{vec}(\mathbf{K})$. Now, the minimal eigenvalue of $\mathbf{X}^T \otimes \mathbf{X}^T$ is $\sigma_{\min}(\mathbf{X}\mathbf{X}^T)$ which is nonzero since \mathbf{X} is rank p .

So we see from Lemma 2.4, that

$$\frac{n\sigma_{\min}(\mathbf{X}\mathbf{X}^T)^2}{|\mathcal{T}|} \|\mathbf{K} - \widehat{\mathbf{K}}\|_F^2 \leq \frac{C_{d,d'}}{C_f^2} (R(\widehat{\mathbf{K}}) - R(\mathbf{K}^*))$$

The result now follows from Theorem 2.1. □