
One for All: Simultaneous Metric and Preference Learning over Multiple Users

Gregory Canal

University of Wisconsin-Madison
Madison, WI
gcanal@wisc.edu

Blake Mason

Rice University
Houston, TX
bm63@rice.edu

Ramya Korlakai Vinayak

University of Wisconsin-Madison
Madison, WI
ramya@ece.wisc.edu

Robert Nowak

University of Wisconsin-Madison
Madison, WI
rdnowak@wisc.edu

Abstract

This paper investigates simultaneous preference and metric learning from a crowd of respondents. A set of items represented by d -dimensional feature vectors and paired comparisons of the form “item i is preferable to item j ” made by each user is given. Our model jointly learns a distance metric that characterizes the crowd’s general measure of item similarities along with a latent ideal point for each user reflecting their individual preferences. This model has the flexibility to capture individual preferences, while enjoying a metric learning sample cost that is amortized over the crowd. We first study this problem in a noiseless, continuous response setting (i.e., responses equal to differences of item distances) to understand the fundamental limits of learning. Next, we establish prediction error guarantees for noisy, binary measurements such as may be collected from human respondents, and show how the sample complexity improves when the underlying metric is low-rank. Finally, we establish recovery guarantees under assumptions on the response distribution. We demonstrate the performance of our model on both simulated data and on a dataset of color preference judgments across a large number of users.

1 Introduction

In many data-driven recommender systems (e.g., streaming services, online retail), multiple users interact with a set of items (e.g., movies, products) that are common to all users. While each user has their individual preferences over these items, there may exist shared structure in how users *perceive* items when making preference judgments. This is a reasonable assumption, since collections of users typically have shared perceptions of similarity between items regardless of their individual item preferences [1–3]. In this work we develop and analyze models and algorithms for simultaneously learning individual preferences and the common metric by which users make preference judgments.

Specifically, suppose there exists a known, fixed set \mathcal{X} of n items, where each item $i \in 1, \dots, n$ is parameterized by a feature vector $\mathbf{x}_i \in \mathbb{R}^d$. We model the crowd’s preference judgments between items as corresponding to a common Mahalanobis distance metric $d_{\mathbf{M}}(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_{\mathbf{M}}$, where $\|\mathbf{x}\|_{\mathbf{M}} := \sqrt{\mathbf{x}^T \mathbf{M} \mathbf{x}}$ and \mathbf{M} is a $d \times d$ positive semidefinite matrix to be learned. Measuring distances with $d_{\mathbf{M}}$ has the effect of reweighting individual features as well as capturing pairwise interactions between features. To capture individual preferences amongst the items, we associate with each of K users an *ideal point* $\mathbf{u}_k \in \mathbb{R}^d$ for $k \in 1, \dots, K$ such that user k prefers items that are closer to \mathbf{u}_k than those items that are farther away, as measured by the common metric $d_{\mathbf{M}}$. The ideal point

model is attractive since it can capture nonlinear notions of preference, and preference rankings are determined simply by sorting item distances to each user point and can therefore be easily generalized to items outside of \mathcal{X} with known embedding features [4–7]. Furthermore, once a user’s point \mathbf{u}_k is estimated, in some generative modeling applications it can then be used to synthesize an “ideal” item for the user located exactly at \mathbf{u}_k , which by definition would be their most preferred item if it existed.

In order to learn the metric and ideal points, we issue a series of *paired comparison* queries to each user in the form “do you prefer item i or item j ?” Since such preferences directly correspond to distance rankings in \mathbb{R}^d , these comparisons provide a signal from which the user points $\{\mathbf{u}_k\}_{k=1}^K$ and common metric M can be estimated. *The main contribution of this work is a series of identifiability, prediction, and recovery guarantees to establish the first theoretical analysis of simultaneous preference and metric learning from paired comparisons over multiple users.* Our key observation is that by modeling a shared metric between all users rather than learning separate metrics for each user, the sample complexity is reduced from $O(d^2)$ paired comparisons per user to only $O(d)$, which is the sample cost otherwise required to learn each ideal point; in essence, when amortizing metric and preference learning over multiple users, the metric comes for free. Our specific contributions include:

- Necessary and sufficient conditions on the number of items and paired comparisons required for exact preference and metric estimation over generic items, when noiseless differences of item distances are known exactly. These results characterize the fundamental limits of our problem in an idealized setting, and demonstrate the benefit of amortized learning over multiple users. Furthermore, when specialized to $K = 1$ our results significantly advance the existing theory of identifiability for single-user simultaneous metric and preference learning [7].
- Prediction guarantees when learning from noisy, one-bit paired comparisons (rather than exact distance comparisons). We present prediction error bounds for two convex algorithms that learn full-rank and low-rank metrics respectively, and again illustrate the sample cost benefits of amortization.
- Recovery guarantees on the metric and ideal points when learning from noisy, binary labels under assumptions on the response distribution. Furthermore, we empirically validate the recovery performance of our multi-user learning algorithms on both synthetic datasets as well as on real psychometrics data studying individual and collective color preferences and perception.

Summary of related work: Metric and preference learning are both extensively studied problems (see [8] and [9] for surveys of each). A common paradigm in metric learning is that by observing distance comparisons, one can learn a linear [10–12], kernelized [13, 14], or deep metric [15, 16] and use it for downstream tasks such as classification. Similarly, it is common in preference learning to use comparisons to learn a ranking or to identify a most preferred item [5, 6, 17–19]. An important family of these algorithms reduces preference learning to identifying an ideal point for a fixed metric [5, 20]. The closest work to ours is [7], who perform metric and preference learning simultaneously from paired comparisons in the single-user case and propose an alternating minimization algorithm that achieves empirical success. However, that work leaves open the question of theoretical guarantees for the simultaneous learning problem, which we address here. A core challenge when establishing such guarantees is that the data are a function of multiple latent parameters (i.e., unknown metric and ideal point(s)) that interact with each other in a nonlinear manner, which complicates standard generalization and identifiability arguments. To this end, we introduce new theoretical tools and advance the techniques of [12] who showed theoretical guarantees for learning a Mahalanobis metric from triplet queries. We survey additional related work more extensively in Appendix B.

Notation: Let $[K] := 1 \dots K$. Unless specified otherwise, $\|\cdot\|$ denotes the ℓ_2 norm when acting on a vector, and the operator norm induced by the ℓ_2 norm when acting on a matrix. Let \mathbf{e}_i denote the i th standard basis vector, $\mathbf{1}$ the vector of all ones, $\mathbf{0}_{a,b}$ the $a \times b$ matrix of all zeros (or $\mathbf{0}$ if the dimensions are clear), and \mathbf{I} the identity matrix, where the dimensionality is inferred from context. For a symmetric $d \times d$ matrix \mathbf{A} , let $\text{vec}^*(\mathbf{A}) := [\mathbf{A}_{1,1}, \mathbf{A}_{1,2}, \dots, \mathbf{A}_{1,d}, \mathbf{A}_{2,2}, \mathbf{A}_{2,3}, \dots, \mathbf{A}_{2,d}, \dots, \mathbf{A}_{d,d}]^T$ denote the vectorized upper triangular portion of \mathbf{A} , which is a D -length vector where $D := d(d+1)/2$. Let $\mathbf{u} \otimes_S \mathbf{v} := \text{vec}^*(\mathbf{u}\mathbf{v}^T)$ denote the unique entries of the Kronecker product between vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$, and let \odot denote the Hadamard (or element-wise) product between two matrices.

2 Identifiability from unquantized measurements

In this section, we characterize the fundamental limits on the number of items and paired comparisons per user required to identify M and $\{\mathbf{u}_k\}_{k=1}^K$ exactly. In order to understand the fundamental

hardness of this problem, we begin by presenting identifiability guarantees under the idealized case where we receive *exact, noiseless* difference of distance measurements¹, before deriving similar results in the case of *noisy* realizations of the *sign* of these differences in the following sections.

We formally define our model as follows: if user k responds that they prefer item i to item j , then $\|\mathbf{x}_i - \mathbf{u}_k\|_M < \|\mathbf{x}_j - \mathbf{u}_k\|_M$. Equivalently, by defining

$$\delta_{i,j}^{(k)} := \|\mathbf{x}_i - \mathbf{u}_k\|_M^2 - \|\mathbf{x}_j - \mathbf{u}_k\|_M^2 = \mathbf{x}_i^T \mathbf{M} \mathbf{x}_i - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j - 2\mathbf{u}_k^T \mathbf{M} (\mathbf{x}_i - \mathbf{x}_j), \quad (1)$$

user k prefers item i over item j if $\delta_{i,j}^{(k)} < 0$ (otherwise j is preferred). In this section, we assume that $\delta_{i,j}^{(k)}$ is measured exactly, and refer to this measurement type as an *unquantized* paired comparison. Let m_k denote the number of unquantized paired comparisons answered by user k and let $m_T := \sum_{k=1}^K m_k$ denote the total number of comparisons made across all users.

It is not immediately clear if recovery of both \mathbf{M} and $\{\mathbf{u}_k\}_{k=1}^K$ is possible from such measurements, which depend quadratically on the item vectors. In particular, one can conceive of pathological examples where these parameters are not identifiable (i.e., there exists no unique solution). For instance, suppose $d = n$, $\mathbf{M} = \alpha \mathbf{I}$ for a scalar $\alpha > 0$, $\mathbf{x}_i = \mathbf{e}_i$ for $i \in [n]$, and for each user $\mathbf{u}_k = \beta_k \mathbf{1}$ for a scalar β_k . Then one can show that $\delta_{i,j}^{(k)} = 0$ for all i, j, k , and therefore $\alpha, \beta_1, \dots, \beta_K$ are unidentifiable from any set of paired comparisons over \mathcal{X} . In what follows, we derive necessary and sufficient conditions on the number and geometry of items, number of measurements per user, and interactions between measurements and users in order for the latent parameters to be identifiable.

Note that eq. (1) includes a nonlinear interaction between \mathbf{M} and \mathbf{u}_k ; however, by defining $\mathbf{v}_k := -2\mathbf{M}\mathbf{u}_k$ (which we refer to as user k 's ‘‘pseudo-ideal point,’’ reflecting the component of \mathbf{u}_k in the column space of \mathbf{M}) eq. (1) becomes linear in \mathbf{M} and \mathbf{v}_k :

$$\delta_{i,j}^{(k)} = \mathbf{x}_i^T \mathbf{M} \mathbf{x}_i - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j + (\mathbf{x}_i - \mathbf{x}_j)^T \mathbf{v}_k. \quad (2)$$

If \mathbf{M} is full-rank and the system of equations admits unique solutions for \mathbf{M} and $\{\mathbf{v}_k\}_{k=1}^K$, then \mathbf{u}_k can be recovered exactly from \mathbf{v}_k .² In other words, the non-convex eq. (1) can be solved in two stages by first solving a linear relaxation (2) in terms of \mathbf{M} and \mathbf{v}_k , and then solving for \mathbf{u}_k . Note that since \mathbf{M} is symmetric, we may write $\mathbf{x}_i^T \mathbf{M} \mathbf{x}_i = \langle \text{vec}^*(2\mathbf{M} - \mathbf{I} \odot \mathbf{M}), \mathbf{x}_i \otimes_S \mathbf{x}_i \rangle$. Defining $\boldsymbol{\nu}(\mathbf{M}) := \text{vec}^*(2\mathbf{M} - \mathbf{I} \odot \mathbf{M})$, from which \mathbf{M} can be determined, we have

$$\delta_{i,j}^{(k)} = [(\mathbf{x}_i \otimes_S \mathbf{x}_i - \mathbf{x}_j \otimes_S \mathbf{x}_j)^T \quad (\mathbf{x}_i - \mathbf{x}_j)^T] \begin{bmatrix} \boldsymbol{\nu}(\mathbf{M}) \\ \mathbf{v}_k \end{bmatrix}.$$

By concatenating all user measurements in a single linear system, we can directly show conditions for identifiability of \mathbf{M} and $\{\mathbf{v}_k\}_{k=1}^K$ by characterizing when the system admits a unique solution. To do so, we define a class of matrices that will encode the item indices in each pair queried to each user:

Definition 2.1. A $a \times b$ matrix \mathbf{S} is a *selection matrix* if for every $i \in [a]$, there exist distinct indices $p_i, q_i \in [b]$ such that $\mathbf{S}[i, p_i] = 1$, $\mathbf{S}[i, q_i] = -1$, and $\mathbf{S}[i, j] = 0$ for $j \in [b] \setminus \{p_i, q_i\}$.

In Appendix C, we characterize several theoretical properties of selection matrices, which will be useful in proving the results that follow.

For each user k , we represent their queried pairs by a $m_k \times n$ selection matrix denoted \mathbf{S}_k , where each row selects a pair of items corresponding to its nonzero entries. Letting $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, $\mathbf{X}_\otimes := [\mathbf{x}_1 \otimes_S \mathbf{x}_1, \dots, \mathbf{x}_n \otimes_S \mathbf{x}_n] \in \mathbb{R}^{D \times n}$, and $\boldsymbol{\delta}_k \in \mathbb{R}^{m_k}$ denote the vector of unquantized measurement values for user k , we can write the entire linear system over all users as a set of m_T equations with $D + dK$ variables to be recovered:

$$\boldsymbol{\Gamma} \begin{bmatrix} \boldsymbol{\nu}(\mathbf{M}) \\ \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_K \end{bmatrix} = \begin{bmatrix} \boldsymbol{\delta}_1 \\ \vdots \\ \boldsymbol{\delta}_K \end{bmatrix} \quad \text{where } \boldsymbol{\Gamma} := \begin{bmatrix} \mathbf{S}_1 \mathbf{X}_\otimes^T & \mathbf{S}_1 \mathbf{X}^T & \mathbf{0}_{m_1,d} & \cdots & \mathbf{0}_{m_1,d} \\ \mathbf{S}_2 \mathbf{X}_\otimes^T & \mathbf{0}_{m_2,d} & \mathbf{S}_2 \mathbf{X}^T & \cdots & \mathbf{0}_{m_2,d} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{S}_K \mathbf{X}_\otimes^T & \mathbf{0}_{m_K,d} & \mathbf{0}_{m_K,d} & \cdots & \mathbf{S}_K \mathbf{X}^T \end{bmatrix}. \quad (3)$$

¹We use the term ‘‘measurement’’ interchangeably with ‘‘paired comparison.’’

²If \mathbf{M} were rank deficient, only the component of \mathbf{u}_k in the row space of \mathbf{M} affects $\delta_{i,j}^{(k)}$. In this case, there is an equivalence class of user points that accurately model their responses. We then take \mathbf{u}_k to be the minimum norm solution, i.e., $\mathbf{u}_k = -\frac{1}{2} \mathbf{M}^\dagger \mathbf{v}_k$. This generalizes Proposition 1 of [7] for the multiple user case.

From this linear system, it is clear that $\nu(M)$ (and hence M) and $\{\mathbf{v}_k\}_{k=1}^K$ (and hence $\{\mathbf{u}_k\}_{k=1}^K$, if M is full-rank) can be recovered exactly if and only if Γ has full column rank. In the following sections, we present necessary and sufficient conditions for this to occur.

2.1 Necessary conditions for identifiability

To build intuition, note that the metric M has D degrees of freedom and each of the K pseudo-ideal points \mathbf{v}_k has d degrees of freedom. Hence, there must be at least $m_T \geq D + Kd$ measurements in total (i.e., rows of Γ) to have any hope of identifying M and $\{\mathbf{v}_k\}_{k=1}^K$. When amortized over the K users, this corresponds to each user providing at least $d + D/K$ measurements on average. In general, d of these measurements are responsible for identifying each user's own pseudo-ideal point (since \mathbf{v}_k is purely a function of user k 's responses), while the remaining D/K contribute towards a *collective* set of D measurements needed to identify the common metric. While these D measurements must be linearly independent from each other and from those used to learn the ideal points, a degree of overlap is acceptable in the additional d measurements each user provides, as the \mathbf{v}_k 's are independent of one another. We formalize this intuition in the following proposition, where we let $\mathbf{S}_T := [\mathbf{S}_1^T, \dots, \mathbf{S}_K^T]^T$ denote the concatenation of all user selection matrices.

Proposition 2.1. *If Γ has full column rank, then $\sum_{k=1}^K m_k \geq D + dK$ and the following must hold:*

- (a) *for all $k \in [K]$, $\text{rank}(\mathbf{S}_k \mathbf{X}^T) = d$, and therefore $\text{rank}(\mathbf{S}_k) \geq d$ and $m_k \geq d$*
- (b) *$\sum_{k=1}^K \text{rank}(\mathbf{S}_k [\mathbf{X}_{\otimes}^T \quad \mathbf{X}^T]) \geq D + dK$, and therefore $\sum_{k=1}^K \text{rank}(\mathbf{S}_k) \geq D + dK$*
- (c) *$\text{rank}(\mathbf{S}_T [\mathbf{X}_{\otimes}^T \quad \mathbf{X}^T]) = D + d$, and therefore $\text{rank}(\mathbf{S}_T) \geq D + d$, $\text{rank}([\mathbf{X}_{\otimes}^T \quad \mathbf{X}^T]) = D + d$, and $n \geq D + d + 1$*

If $\sum_{k=1}^K m_k = D + dK$ exactly, then (a) and (b) are equivalent to $m_k \geq d \forall k$ and each user's selection matrix having full row rank. (c) implies that the number of required items n scales as $\Omega(d^2)$; in higher dimensional feature spaces, this scaling could present a challenge since it might be difficult in practice to collect such a large number of items for querying. Finally, note that the conditions in Proposition 2.1 are *not* sufficient for identifiability: in Appendix C.6, we present a counterexample where these necessary properties are fulfilled, yet the system is not invertible.

2.2 Sufficient condition for identifiability

Next, we present a class of pair selection schemes that are sufficient for parameter identifiability and match the item and measurement count lower bounds in Proposition 2.1. This result leverages the idea that as long as the d measurements each user provides to learn their ideal point do not "overlap" with the D measurements collectively provided to learn the metric, then the set of m_T total measurements is sufficiently rich to ensure a unique solution. First, we define a property of certain selection matrices where each pair introduces at least one new item that has not yet been selected:

Definition 2.2. An $m \times n$ selection matrix \mathbf{S} is *incremental* if for all $i \in [m]$, at least one of the following is true, where p_i and q_i are as defined in Definition 2.1: (a) for all $j < i$, $\mathbf{S}[j, p_i] = 0$; (b) for all $j < i$, $\mathbf{S}[j, q_i] = 0$.

We now present a class of invertible measurement schemes that builds on the definition of incrementality. For simplicity assume that $m_T = D + dK$ exactly, which is the lower bound from Proposition 2.1. Additionally, assume without loss of generality that each $m_k > d$; if instead there existed a user k^* such that $m_{k^*} = d$ exactly, one can show under the necessary conditions in Proposition 2.1 that the system would separate into two subproblems where first the metric would need to be learned from the other $K - 1$ users, and then \mathbf{v}_{k^*} is solved for directly from user k^* 's measurements.

Proposition 2.2. *Let $K \geq 1$, and suppose $m_k > d \forall k \in [K]$, $m_T = D + dK$, and $n \geq D + d + 1$. Suppose that for each $k \in [K]$, there exists a $d \times n$ selection matrix $\mathbf{S}_k^{(1)}$ and $m_k - d \times n$ selection matrix $\mathbf{S}_k^{(2)}$ such that $\mathbf{S}_k = [(\mathbf{S}_k^{(1)})^T \quad (\mathbf{S}_k^{(2)})^T]^T$, and that the following are true:*

- (a) *For all $k \in [K]$, $\text{rank}(\mathbf{S}_k^{(1)}) = d$*

(b) Defining the $D \times n$ selection matrix $\mathbf{S}^{(2)}$ as $\mathbf{S}^{(2)} := [(\mathbf{S}_1^{(2)})^T \dots (\mathbf{S}_K^{(2)})^T]^T$, there exists a $D \times D$ permutation \mathbf{P} such that for each $k \in [K]$, $\begin{bmatrix} \mathbf{S}_k^{(1)} \\ \mathbf{P}\mathbf{S}^{(2)} \end{bmatrix}$ is incremental

Additionally, suppose each item \mathbf{x}_i is sampled i.i.d. from a distribution p_X that is absolutely continuous with respect to the Lebesgue measure. Then with probability 1, $\mathbf{\Gamma}$ has full column rank.

Remark 2.3. In Appendix C.6 we construct a pair selection scheme that satisfies the conditions³ in Proposition 2.2 while only using the minimum number of measurements and items, with $m_k = d + D/K$ (and therefore $m_T = D + dK$) and $n = D + d + 1$. Importantly, this construction confirms that the lower bounds on the number of measurements and items in Proposition 2.1 are in fact tight. Since $D = O(d^2)$, if $K = \Omega(d)$ then only $m_k = O(d)$ measurements are required per user. This scaling demonstrates the benefit of amortizing metric learning across multiple users, since in the single user case $D + d = \Omega(d^2)$ measurements would be required.

2.3 Single user case

In the case of a single user ($K = 1$), it is straightforward to show that the necessary and sufficient selection conditions in Proposition 2.1 and Proposition 2.2 respectively are *equivalent*, and simplify to the condition that $\text{rank}(\mathbf{S}) \geq D + d$ (where we drop the subscript on \mathbf{S}_1). In a typical use case, a practitioner is unlikely to explicitly select pair indices that result in \mathbf{S} being full-rank, and instead would select pairs uniformly at random from the set of $\binom{n}{2}$ unique item pairs. By proving a tail bound on the number of random comparisons required for \mathbf{S} to be full-rank, we have with high probability that randomly selected pairs are sufficient for metric and preference identifiability in the single user case. We summarize these results in the following corollary:

Corollary 2.3.1. *When $K = 1$, if $\mathbf{\Gamma}$ is full column rank then $\text{rank}(\mathbf{S}) \geq D + d$. Conversely, for a fixed \mathbf{S} satisfying $\text{rank}(\mathbf{S}) \geq D + d$, if each \mathbf{x}_i is sampled i.i.d. according to a distribution p_X that is absolutely continuous with respect to the Lebesgue measure then $\mathbf{\Gamma}$ is full column rank with probability 1. If each pair is selected independently and uniformly at random with $n = \Omega(D + d)$ and $m_T = \Omega(D + d)$, then if \mathbf{x}_i is drawn i.i.d. from p_X , $\mathbf{\Gamma}$ has full column rank with high probability.*

Importantly, the required item and sample complexity for randomly selected pairs matches the lower bounds in Proposition 2.1 up to a constant. As we describe in Appendix C.7, we conjecture that a similar result holds for the multiuser case ($K > 1$), which is left to future work.

3 Prediction and generalization from binary labels

In practice, we do not have access to exact difference of distance measurements. Instead, paired comparisons are one-bit measurements (given by the user preferring one item over the other) that are sometimes noisy due to inconsistent user behavior or from model deviations. In this case, rather than simply solving a linear system, we must optimize a loss function that penalizes incorrect response predictions while enforcing the structure of our model. In this section, we apply a different set of tools from statistical learning theory to characterize the sample complexity of randomly selected paired comparisons under a general noise model, optimized under a general class of loss functions.

We assume that each pair p is sampled uniformly with replacement from the set of $\binom{n}{2}$ pairs, and the user k queried at each iteration is independently and uniformly sampled from the set of K users. For a pair $p = (i, j)$ given to user k , we observe a (possibly noisy) binary response $y_p^{(k)}$ where $y_p^{(k)} = -1$ indicates that user k prefers item i to j , and $y_p^{(k)} = 1$ indicates that j is preferred. Let $\mathcal{S} := \{(p, k, y_p^{(k)})\}_{p=(i,j)}$ be an i.i.d. joint dataset over pairs p , selected users k , and responses $y_p^{(k)}$, where $|\mathcal{S}|$ denotes the number of such data points. We wish to learn \mathbf{M} and vectors $\{\mathbf{u}_k\}_{k=1}^K$ that

³We note that these conditions are not exhaustive: in Appendix C.6 we construct an example where $\mathbf{\Gamma}$ is full column rank, yet the conditions in Proposition 2.2 are not met. A general set of matching necessary and sufficient identifiability conditions on $\{\mathbf{S}_k\}_{k=1}^K$ has remained elusive; towards this end, in Appendix C.7 we describe a more comprehensive set of conditions that we conjecture are sufficient for identifiability.

predict the responses in \mathcal{S} : given a convex, L -Lipschitz loss $\ell: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$,⁴ we wish to solve

$$\begin{aligned} \min_{\mathbf{M}, \{\mathbf{u}_k\}_{k=1}^K} & \frac{1}{|\mathcal{S}|} \sum_{\mathcal{S}} \ell \left(y_p^{(k)} \left(\|\mathbf{u}_k - \mathbf{x}_i\|_{\mathbf{M}}^2 - \|\mathbf{u}_k - \mathbf{x}_j\|_{\mathbf{M}}^2 \right) \right) \\ \text{s.t. } & \mathbf{M} \succeq 0, \|\mathbf{M}\|_F \leq \lambda_F, \|\mathbf{u}_k\|_2 \leq \lambda_u \forall k \in [K], |\delta_{i,j}^{(k)}| \leq \gamma \forall i, j, k \end{aligned} \quad (4)$$

where $\lambda_F, \lambda_u, \gamma > 0$ are hyperparameters and $\delta_p^{(k)}$ is defined as in eq. (1). The constraint $\mathbf{M} \succeq 0$ ensures that \mathbf{M} defines a metric, the Frobenius and ℓ_2 norm constraints prevent overfitting, and the constraint on $\delta_p^{(k)}$ is a technical point to avoid pathological cases stemming from coherent \mathbf{x} vectors.

The above optimization is non-convex due to the interaction between the \mathbf{M} and \mathbf{u} terms. Instead, as in Section 2 we define $\mathbf{v}_k := -2\mathbf{M}\mathbf{u}_k$ and solve the relaxation⁵

$$\begin{aligned} \min_{\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K} & \widehat{R}(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K) \text{ s.t. } \mathbf{M} \succeq 0, \|\mathbf{M}\|_F \leq \lambda_F, \|\mathbf{v}_k\|_2 \leq \lambda_v \forall k \in [K], |\delta_{i,j}^{(k)}| \leq \gamma \forall i, j, k \\ \text{where } \widehat{R}(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K) & := \frac{1}{|\mathcal{S}|} \sum_{\mathcal{S}} \ell \left(y_p^{(k)} \left(\mathbf{x}_i^T \mathbf{M} \mathbf{x}_i - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j + \mathbf{v}_k^T (\mathbf{x}_i - \mathbf{x}_j) \right) \right). \end{aligned} \quad (5)$$

The quantity $\widehat{R}(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K)$ is the *empirical risk*, given dataset \mathcal{S} . The empirical risk is an unbiased estimate of the true risk given by

$$R(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K) := \mathbb{E} \left[\ell \left(y_p^{(k)} \left(\mathbf{x}_i^T \mathbf{M} \mathbf{x}_i - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j + \mathbf{v}_k^T (\mathbf{x}_i - \mathbf{x}_j) \right) \right) \right],$$

where the expectation is with respect to a random draw of $p = (i, j), k$, and $y_p^{(k)}$ conditioned on the choice of p and k . Let $\widehat{\mathbf{M}}$ and $\{\widehat{\mathbf{v}}_k\}_{k=1}^K$ denote the minimizers of the empirical risk optimization in eq. (5), and let \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$ minimize the *true* risk, subject to the same constraints. The following theorem bounds the excess risk of the empirical optimum $R(\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K)$ relative to the optimal true risk $R(\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K)$.

Theorem 3.1. *Suppose $\|\mathbf{x}_i\|_2 \leq 1$ for all $i \in [n]$. With probability at least $1 - \delta$,*

$$\begin{aligned} R(\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K) - R(\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K) & \leq \sqrt{\frac{256L^2(\lambda_F^2 + K\lambda_v^2)}{|\mathcal{S}|} \log(d^2 + d + 1)} \\ & + \frac{\sqrt{128L^2(\lambda_F^2 + K\lambda_v^2)}}{3|\mathcal{S}|} \log(d^2 + d + 1) + \sqrt{\frac{8L^2\gamma^2 \log(\frac{2}{\delta})}{|\mathcal{S}|}}. \end{aligned} \quad (6)$$

Remark 3.2. To put this result in context, suppose $\|\mathbf{M}^*\|_F = d$ so that the average squared magnitude of each entry is a constant, in which case we can set $\lambda_F = d$. Similarly, if each entry of \mathbf{v}_k is dimensionless, then $\|\mathbf{v}_k\|_2 \propto \sqrt{d}$ and so we can set $\lambda_v = \sqrt{d}$. We then have that the excess risk in eq. (6) is $\widetilde{O} \left(\sqrt{\frac{d^2 + Kd}{|\mathcal{S}|}} \right)$ where \widetilde{O} suppresses logarithmic factors, implying a sample complexity of $d^2 + Kd$ measurements across all users, and therefore an average of $d + d^2/K$ measurements per user. If $K = \Omega(d)$, this is equivalent to $\widetilde{O}(d)$ measurements per user, which corresponds to the parametric rate required per user in order to estimate their pseudo-ideal point \mathbf{v}_k . Similar to the case of unquantized measurements, the $O(d^2)$ sample cost of estimating the metric from noisy one-bit comparisons has been amortized across all users, demonstrating the benefit of learning multiple user preferences simultaneously when the users share a common metric.

3.1 Low-rank modeling

In many settings, the metric \mathbf{M} may be low-rank with rank $r < d$ [8, 12]. In this case, \mathbf{M} only has dr degrees of freedom rather than d^2 degrees as in the full-rank case. Therefore if $K = \Omega(d)$, we intuitively expect the sample cost of learning the metric to be amortized to a cost of $O(r)$

⁴We restrict ourselves to the case where the loss is a function of $y_p^{(k)} \left(\|\mathbf{u}_k - \mathbf{x}_i\|_{\mathbf{M}}^2 - \|\mathbf{u}_k - \mathbf{x}_j\|_{\mathbf{M}}^2 \right)$.

⁵Specifically we may choose $\lambda_v = 2\lambda_F\lambda_u$, resulting in a constraint set containing the solution to (4). To see this, let $\mathbf{M}^*, \{\mathbf{u}_k^*\}_{k=1}^K$ be the solution to (4) and define $\mathbf{v}_k^* = -2\mathbf{M}^*\mathbf{u}_k^*$. Then $\|\mathbf{v}_k^*\|_2 = 2\|\mathbf{M}^*\mathbf{u}_k^*\|_2 \leq 2\|\mathbf{M}^*\|_F\|\mathbf{u}_k^*\|_2 \leq 2\lambda_F\lambda_u$.

measurements per user. Furthermore, as each \mathbf{v}_k is contained in the r -dimensional column space of \mathbf{M} , we also expect a sample complexity of $O(r)$ to learn each user's pseudo-ideal point. Hence, we expect the amortized sample cost per user to be $O(r)$ in the low-rank setting, which can be a significant improvement over $O(d)$ in the full-rank setting when $r \ll d$.

Algorithmically, ideally one would constrain the $\widehat{\mathbf{M}}$ and $\{\widehat{\mathbf{v}}_k\}_{k=1}^K$ that minimize the empirical risk such that $\text{rank}(\mathbf{M}) = r$ and $\mathbf{v}_k \in \text{colsp}(\mathbf{M})$; unfortunately, such constraints are not convex. Towards a convex algorithm, note that since $\mathbf{v}_k \in \text{colsp}(\mathbf{M})$, $\text{rank}([\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]) = \text{rank}(\mathbf{M}) = r$. Thus, it is sufficient to constrain the rank of $[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]$. We relax this constraint to a convex constraint on the nuclear norm $\|[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]\|_*$, and solve a similar optimization problem to eq. (5):

$$\min_{\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K} \widehat{R}(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K) \text{ s.t. } \mathbf{M} \succeq 0, \|[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]\|_* \leq \lambda_*, |\delta_{i,j}^{(k)}| \leq \gamma \forall i, j, k. \quad (7)$$

We again let \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$ minimize the true risk $R(\mathbf{M}, \{\mathbf{v}_k\}_{k=1}^K)$, subject to the same constraints. The following theorem bounds the excess risk over this constraint set:

Theorem 3.3. *Suppose $\|\mathbf{x}_i\|_2 \leq 1$ for all $i \in [n]$. With probability at least $1 - \delta$,*

$$\begin{aligned} R(\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K) - R(\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K) \leq & 2L \sqrt{\frac{2\lambda_*^2 \log(2d+K)}{|\mathcal{S}|} \left[\left(8 + \frac{4 \min(d, n)}{K}\right) \frac{\|\mathbf{X}\|^2}{n} + \frac{16}{\sqrt{K}} \right]} \\ & + \frac{8L\lambda_*}{3|\mathcal{S}|} \log(2d+K) + \sqrt{\frac{8L^2\gamma^2 \log(2/\delta)}{|\mathcal{S}|}}. \end{aligned}$$

To put this result in context, suppose that the items \mathbf{x}_i and ideal points \mathbf{u}_k are sampled i.i.d. from $\mathcal{N}(\mathbf{0}, \frac{1}{d}\mathbf{I})$. With this item distribution it is straightforward to show that with high probability, $\|\mathbf{X}\|^2 = O(\frac{n}{d})$ (see [21]). For a given $r < d$ let $\mathbf{M} = \frac{d}{\sqrt{r}}\mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a $d \times r$ matrix with orthonormal columns sampled uniformly from the Grassmanian. With this choice of scaling we have $\|\mathbf{M}\|_F = d$, so that each element of \mathbf{M} is dimensionless on average. Furthermore, recalling that $\mathbf{v}_k = -2\mathbf{M}\mathbf{u}_k$, with this choice of scaling $\mathbb{E}[\|\mathbf{v}_k\|_2^2] \propto d$ and so each entry of \mathbf{v}_k on average is dimensionless. To choose a setting for λ^* recall that $[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]$ has rank r and therefore

$$\|[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]\|_* \leq \sqrt{r} \|[\mathbf{M}, \mathbf{v}_1, \dots, \mathbf{v}_K]\|_F \leq \sqrt{r(d^2 + K \max_{k \in [K]} \|\mathbf{v}_k\|_2^2)},$$

which one can show is $O(\sqrt{r(d^2 + dK \log K)})$ with high probability and so we set $\lambda_* = O(\sqrt{r(d^2 + dK \log K)})$. With these term scalings, we have the following corollary:

Corollary 3.3.1. *Let $n \geq d$, $\mathbf{x}_i, \mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \frac{1}{d}\mathbf{I})$ and $\mathbf{M} = \frac{d}{\sqrt{r}}\mathbf{L}\mathbf{L}^T$, where \mathbf{L} is a $d \times r$ matrix with orthonormal columns. If $K = \Omega(d^2)$, then in the same setting as Theorem 3.3 with high probability*

$$R(\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K) - R(\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K) = \tilde{O}\left(\sqrt{\frac{dr + Kr}{|\mathcal{S}|}}\right).$$

Remark 3.4. The scaling $|\mathcal{S}| = O(dr + Kr)$ matches our intuition that $O(dr)$ collective measurements should be made across all users to account for the dr degrees of freedom in \mathbf{M} , in addition to $O(r)$ measurements per user to resolve their own pseudo-ideal point's r degrees of freedom. If $K = \Omega(d)$, then each user answering $O(r)$ queries is sufficient to amortize the cost of learning the metric with the same order of measurements per user as is required for their ideal point. Although Corollary 3.3.1 requires the even stronger condition that $K = \Omega(d^2)$, we believe this is an artifact of our analysis and that $K = \Omega(d)$ should suffice.⁶ Even so, a $\Omega(d^2)$ user count scaling might be reasonable in practice since recommender systems typically operate over large populations of users.

⁶The required user scaling of $K = \Omega(d^2)$ implies that Corollary 3.3.1 only recovers an amortized scaling of $O(r)$ measurements per user if the user count is very large. Nevertheless, the original statement of Theorem 3.3 does apply for any user count, and in this case the only drawback to not invoking $K = \Omega(d^2)$ is that the amortized scaling is larger than $O(r + dr/K)$ measurements per user. We believe this scaling can in fact be tightened to $O(r + dr/K)$ measurements per user for all user counts K (not just $K = \Omega(d^2)$), which we leave to future work. See the proof of Corollary 3.3.1 for additional discussion.

4 Recovery guarantees

The results in the previous section give guarantees on the generalization error of a learned metric and ideal points when predicting pair responses over \mathcal{X} , but do not bound the recovery error of the learned parameters $\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K$ with respect to \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$. Yet, in some settings such as data generated from human responses [22,23] it may be reasonable to assume that a true \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$ do exist that generate the observed data (rather than serving only as a model) and that practitioners may wish to estimate and interpret these latent variables, in which case accurate recovery is critical. Unfortunately, for an arbitrary noise model and loss function, recovering \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$ exactly is generally impossible if the model is not identifiable. However, we now show that with a small amount of additional structure, one can ensure that $\widehat{\mathbf{M}}$ and $\{\widehat{\mathbf{v}}_k\}_{k=1}^K$ accurately approximate \mathbf{M}^* and $\{\mathbf{v}_k^*\}_{k=1}^K$ if a sufficient number of one-bit comparisons are collected.

We assume a model akin to that of [12] for the case of triplet metric learning. Let $f: \mathbb{R} \rightarrow [0, 1]$ be a strictly monotonically increasing *link function* satisfying $f(x) = 1 - f(-x)$; for example, $f(x) = (1 + e^{-x})^{-1}$ is the logistic link and $f(x) = \Phi(x)$ is the probit link where $\Phi(\cdot)$ denotes the CDF of a standard normal distribution. Defining $\delta_p(\mathbf{M}, \mathbf{v}) := \mathbf{x}_i^T \mathbf{M} \mathbf{x}_i - \mathbf{x}_j^T \mathbf{M} \mathbf{x}_j + \mathbf{v}^T (\mathbf{x}_i - \mathbf{x}_j)$ for $p = (i, j)$, we assume that $\mathbb{P}(y_p^{(k)} = -1) = f(-\delta_p(\mathbf{M}^*, \mathbf{v}_k^*))$ for some $\mathbf{M}^* \succeq \mathbf{0}$ and $\mathbf{v}_k^* \in \text{colsp}(\mathbf{M}^*)$. This naturally reflects the idea that some queries are easier to answer (and thus less noisy) than others. For instance, if $\delta_{ij}^{(k)} \ll 0$ such as may occur when \mathbf{x}_i very nearly equals user k 's ideal point, we may assume that user k almost always prefers item i to j and so $f(-\delta_{ij}^{(k)}) \rightarrow 1$ (since f is monotonic). Furthermore, we assume that eq. (5) is optimized with the *negative log-likelihood* loss ℓ_f induced by f : $\ell_f(y_p, p; \mathbf{M}, \mathbf{v}) := -\log(f(y_p \delta_p(\mathbf{M}, \mathbf{v}_k)))$. In Appendix E, we show that we may lower bound the excess risk of $\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K$ by the squared error between the *unquantized* measurements corresponding to $\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K$ and $\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K$. We then utilize tools from Section 2 combined with the results in Section 3 to arrive at the following recovery guarantee.

Theorem 4.1. *Fix a strictly monotonic link function f satisfying $f(x) = 1 - f(-x)$. Suppose for a given item set \mathcal{X} with $n \geq D + d + 1$ and $\|\mathbf{x}_i\| \leq 1 \forall i \in [n]$ that the pairs and users in dataset \mathcal{S} are sampled independently and uniformly at random, and that user responses are sampled according to $\mathbb{P}(y_p^{(k)} = -1) = f(-\delta_p(\mathbf{M}^*, \mathbf{v}_k^*))$ where $\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K$ satisfy the constraints in (5). Let $\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K$ be the solution to (5) solved using loss ℓ_f . Then with probability at least $1 - \delta$,*

$$\frac{1}{n} \sigma_{\min}(\mathbf{J}[\mathbf{X}_{\otimes}^T, \mathbf{X}^T])^2 \left(\|\widehat{\mathbf{M}} - \mathbf{M}^*\|_F^2 + \frac{1}{K} \sum_{k=1}^K \|\widehat{\mathbf{v}}_k - \mathbf{v}_k^*\|^2 \right) \leq \frac{4}{C_f^2} \sqrt{\frac{L^2(\lambda_F^2 + K\lambda_v^2)}{|\mathcal{S}|} \log(d^2 + d + 1)} + \frac{\sqrt{8L^2(\lambda_F^2 + K\lambda_v^2)}}{3C_f^2 |\mathcal{S}|} \log(d^2 + d + 1) + \frac{1}{C_f^2} \sqrt{\frac{L^2 \gamma^2 \log(\frac{2}{\delta})}{2|\mathcal{S}|}},$$

where $C_f = \min_{z: |z| \leq \gamma} f'(z)$ and $\mathbf{J} := \mathbf{I}_n - \frac{1}{n} \mathbf{1}_n \mathbf{1}_n^T$ is the centering matrix. Furthermore, if \mathcal{X} is constructed by sampling each item i.i.d. from a distribution p_X with support on the unit ball that is absolutely continuous with respect to the Lebesgue measure, then with probability 1, $\sigma_{\min}(\mathbf{J}[\mathbf{X}_{\otimes}^T, \mathbf{X}^T]) > 0$.

Remark 4.2. The key conclusion from this result is that since $\sigma_{\min}(\mathbf{J}[\mathbf{X}_{\otimes}^T, \mathbf{X}^T]) > 0$ almost surely, the recovery error of $\widehat{\mathbf{M}}, \{\widehat{\mathbf{v}}_k\}_{k=1}^K$ with respect to $\mathbf{M}^*, \{\mathbf{v}_k^*\}_{k=1}^K$ is upper bounded by a decreasing function of the number of data points, $|\mathcal{S}|$. In other words, the metric and pseudo-ideal points are identifiable from one-bit paired comparisons under an assumed response distribution. As discussed in Section 2, the ideal points $\{\mathbf{u}_k^*\}_{k=1}^K$ are then also identifiable as long as \mathbf{M}^* is full-rank. We present an analogous result for the case of a low-rank metric in Appendix E, and leave to future work a study of the scaling of $\sigma_{\min}(\mathbf{J}[\mathbf{X}_{\otimes}^T, \mathbf{X}^T])$ with respect to d and n .

5 Experimental results

Below we analyze the performance of the empirical risk minimizers given in eqs. (5) and (7) on both simulated and real-world data,⁷ with further details deferred to Appendix F.

⁷Code available at <https://github.com/gregcanal/multiuser-metric-preference>

Simulated experiments: We first simulate data in a similar setting to Cor. 3.3.1 where $x_i, \mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \frac{1}{d}\mathbf{I})$ and $M^* = \frac{d}{\sqrt{r}}\mathbf{L}\mathbf{L}^T$ where $\mathbf{L} \in \mathbb{R}^{d \times r}$ is a random orthogonal matrix. To construct the training dataset, we query a fixed number of randomly selected pairs per user and evaluate prediction accuracy on a held-out test set, where all responses are generated according to a logistic link function (injecting response noise). We evaluate the prediction accuracy of the Frobenius norm regularized optimization in eq. (5) (referred to as **Frobenius metric**), designed for full-rank matrix recovery, as well as the nuclear norm regularized optimization in eq. (7) (referred to as **Nuclear full**), designed for low-rank metrics. We also compare against several ablation methods that modify the constraint sets in (5) and (7): **Nuclear metric**, where $\|M\|_*$ and $\|\mathbf{v}_k\|_2$ are constrained; **Nuclear split**, where $\|M\|_*$ and $\|[\mathbf{v}_1, \dots, \mathbf{v}_K]\|_*$ are constrained; and **PSD only**, where only $M \succeq \mathbf{0}$ is enforced. We also compare against **Nuclear full, single**, which is equivalent to **Nuclear full** when applied *separately* to each user (learning a unique metric and ideal point), where test accuracy is averaged over all users. To compare performance under a best-case hyperparameter setting, we tune each method’s respective constraints using oracle knowledge of M^* and $\{\mathbf{u}_k^*\}_{k=1}^K$. Finally, we also evaluate prediction accuracy when the ground-truth parameters are known exactly (i.e., $M = M^*, \mathbf{v}_k = -2M^*\mathbf{u}_k^*$), which we call **Oracle**. This accuracy gives the “best case” performance, and reflects the inherent response noise in the generated data.

To test a low-rank setting, we set $d = 10, r = 1, n = 100$, and $K = 10$. We observe that **Nuclear full** outperforms the baseline methods in terms of test accuracy, and is closely followed by **Nuclear split** (Fig. 1a). Interestingly **Nuclear metric**, which also enforces a nuclear norm constraint on M , does not perform as well, possibly because it does not encourage the pseudo-ideal points to lie in the same low-rank subspace. While **Nuclear metric** does demonstrate slightly improved metric recovery (Fig. 1b), **Nuclear full** and **Nuclear split** recover higher quality metrics for lower query counts (which is the typical operating regime for human-in-the-loop systems) and exhibit significantly better ideal point recovery (Fig. 1c), illustrating the importance of proper subspace alignment between the pseudo-ideal points. To this end, unlike **Nuclear split**, **Nuclear full** explicitly encourages the pseudo-ideal points to align with the column space of M , which may explain its slight advantage. Finally, we note that **Nuclear full, single** results in the worst prediction accuracy, demonstrating the benefit of a common metric model when the underlying metric is in fact shared. While the single user case is not the focus of this work, in Appendix F.4 we compare the performance of **Nuclear full, single** against the methods proposed in [7], which only considers the single user case.

Color dataset: We also study the performance of our model on a dataset of pairwise color preferences across multiple respondents ($K = 48$) [24]. In this setting, each color ($n = 37$) is represented as a 3-dimensional vector in CIELAB color space (lightness, red vs. green, blue vs. yellow), which was designed as a uniform space for how humans perceive color [25]. Each respondent was asked to order pairs of color by preference, as described in [26, Sec. 3.1]. Since all $2\binom{37}{2}$ possible pairs (including each pair reversal) were queried for each respondent, we may simulate random pair sampling exactly.

As there are only $d = 3$ features, we constrain the Frobenius norm of the metric and optimize eq. (5) using the hinge loss. Varying the number of pairs queried per user, we plot prediction accuracy on a held-out test set (**Learned M, crowd** in Fig. 1d). As CIELAB is designed to be perceptually uniform, we compare against a solution to eq. (5) that *fixes* $M = \mathbf{I}$ and only learns the points $\{\mathbf{v}_k\}_{k=1}^{48}$ (**Identity M** in Fig. 1d). This method leads to markedly lower prediction accuracy than simultaneously learning the metric and ideal points; this result suggests that although people’s *perception* of color is uniform in this space, their *preferences* are not. We also compare against a baseline that solves the same optimization as eq. (5) *separately* for each individual respondent (learning a unique metric and ideal point per user), with prediction accuracy averaged over all respondents (**Learned M, single** in Fig. 1d). Although learning individual metrics appears to result in better prediction after many queries, in the low-query regime (< 20 pairs per user) learning a common metric across all users results in slightly improved performance (see Appendix F for zoomed plot). As $d = 3$ is small relative to the number of queries given to each user, the success of individual metric learning is not unexpected; however, collecting $O(d^2)$ samples per user is generally infeasible for larger d unlike collective metric learning which benefits from crowd amortization. Finally, learning a single metric common to all users allows for insights into the crowd’s general measure of color similarity. As can be seen in Fig. 1e, the learned metric is dominated by the “lightness” feature, indicating that people’s preferences correspond most strongly to a color’s lightness. As an external validation, this is consistent with the findings of Fig. 1 of [24].

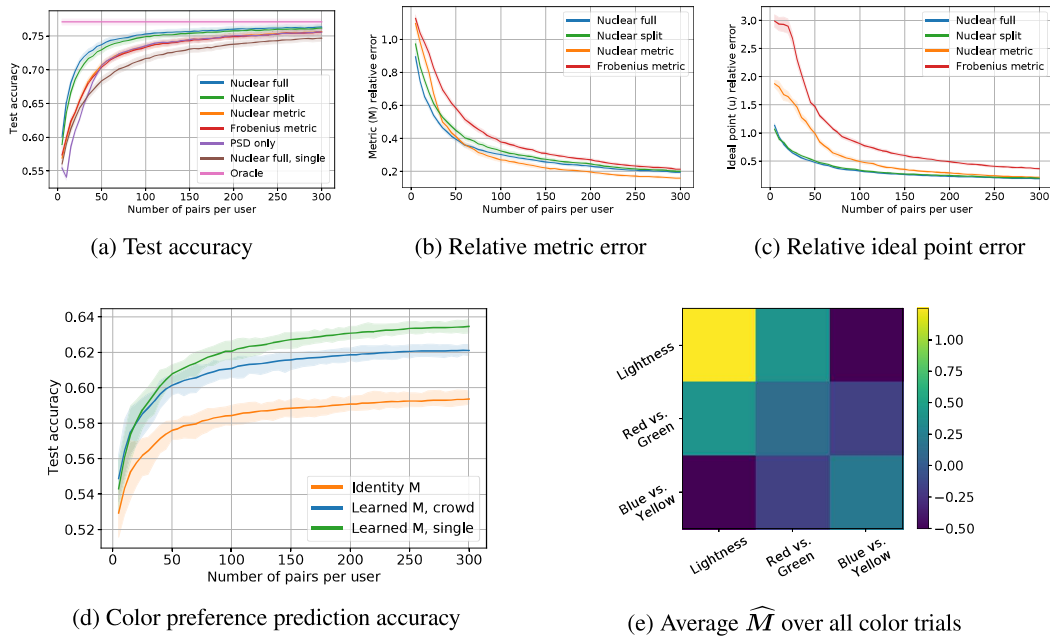


Figure 1: (a-c) Normally distributed items with $d = 10$, $r = 1$, $n = 100$ and $K = 10$. Error bars indicate ± 1 standard error about the sample mean. For visual clarity, **PSD only** and **Nuclear full, single** baselines are omitted from (b-c) due to poor performance. (d) Average color preference prediction accuracy, where error bars indicate 2.5% and 97.5% percentiles. (e) Estimated color preference metric. For (a-e), random train/test splitting was repeated over 30 trials.

6 Discussion

The main contribution of this work is a model for multi-user simultaneous metric and preference learning consisting of a shared metric that captures the crowd’s common perceptions between items, as well as user-specific ideal points that characterize individual preferences. Our core result is that when querying paired comparisons over a large number of users, the sample cost of metric learning is distributed over the crowd, with the total number of queries per user scaling with the rank of the underlying metric.

One interesting avenue for future work is to study more flexible metric models, since as with any Mahalanobis metric model it may not be the case that linear weightings of quadratic feature combinations provide enough flexibility to adequately model certain preference judgements. While it may be possible to generalize the linear metric results studied here to a more general Hilbert space, another avenue to increase model flexibility is to acquire a richer set of features in the item set, which would typically involve increasing the ambient dimension. In this setting, an important consideration would be to ameliorate the requirement that the number of items scales as $\Omega(d^2)$. While this condition is necessary for identifiability if the metric is full-rank, we believe that the item count should only need to scale as $\Omega(dr)$ if the metric has rank $r < d$ (see Appendix C.3 for additional discussion). Furthermore, in the case of metric learning from triplet queries it is known that only $\Omega(r)$ items are required for metric recovery in the low-rank case [12]. This suggests a potential strategy of supplementing paired comparisons with triplet queries to reduce the number of required items.

Acknowledgements and Disclosure of Funding

We thank the reviewers for their feedback, as well as Karen Schloss for providing the color preference data used in our experiments and for useful discussions and references about color preference learning. This work was partially supported by AFOSR/AFRL grant FA9550-18-1-0166 and ARMY MURI grant W911NF-15-1-0479.

References

- [1] R. N. Shepard, "The analysis of proximities: Multidimensional scaling with an unknown distance function. i." *Psychometrika*, vol. 27, no. 2, pp. 125–140, 1962.
- [2] C. H. Coombs, "A theory of data." 1964.
- [3] O. Tamuz, C. Liu, S. Belongie, O. Shamir, and A. T. Kalai, "Adaptively learning the crowd kernel," *arXiv preprint arXiv:1105.1033*, 2011.
- [4] G. S. Carpenter and K. Nakamoto, "Consumer preference formation and pioneering advantage," *Journal of Marketing research*, vol. 26, no. 3, pp. 285–298, 1989.
- [5] K. G. Jamieson and R. Nowak, "Active ranking using pairwise comparisons," in *Advances in Neural Information Processing Systems*, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, Eds., vol. 24. Curran Associates, Inc., 2011. [Online]. Available: <https://proceedings.neurips.cc/paper/2011/file/6c14da109e294d1e8155be8aa4b1ce8e-Paper.pdf>
- [6] G. Canal, A. Massimino, M. Davenport, and C. Rozell, "Active embedding search via noisy paired comparisons," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 902–911. [Online]. Available: <https://proceedings.mlr.press/v97/canal19a.html>
- [7] A. Xu and M. Davenport, "Simultaneous preference and metric learning from paired comparisons," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 454–465. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/0561bc7ecba98e39ca7994f93311ba23-Paper.pdf>
- [8] A. Bellet, A. Habrard, and M. Sebban, "Metric learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 9, no. 1, pp. 1–151, 2015.
- [9] J. Fürnkranz and E. Hüllermeier, "Preference learning and ranking by pairwise comparison," in *Preference learning*. Springer, 2010, pp. 65–82.
- [10] K. Q. Weinberger, J. Blitzer, and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," in *Advances in neural information processing systems*, 2006, pp. 1473–1480.
- [11] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon, "Information-theoretic metric learning," in *Proceedings of the 24th international conference on Machine learning*, 2007, pp. 209–216.
- [12] B. Mason, L. Jain, and R. Nowak, "Learning low-dimensional metrics," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/f12ee9734e1edf70ed02d9829018b3d9-Paper.pdf>
- [13] R. Chatpatanasiri, T. Korsrilabutr, P. Tangchanachaianan, and B. Kijsirikul, "A new kernelization framework for mahalanobis distance learning algorithms," *Neurocomputing*, vol. 73, no. 10-12, pp. 1570–1579, 2010.
- [14] M. Kleindessner and U. von Luxburg, "Kernel functions based on triplet comparisons," *arXiv preprint arXiv:1607.08456*, 2016.
- [15] M. Kaya and H. Ş. Bilge, "Deep metric learning: A survey," *Symmetry*, vol. 11, no. 9, p. 1066, 2019.
- [16] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *International workshop on similarity-based pattern recognition*. Springer, 2015, pp. 84–92.
- [17] K. G. Jamieson and R. D. Nowak, "Low-dimensional embedding using adaptively selected ordinal data," in *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2011, pp. 1077–1084.
- [18] K.-S. Jun, L. Jain, B. Mason, and H. Nassif, "Improved confidence bounds for the linear logistic model and applications to bandits," in *International Conference on Machine Learning*. PMLR, 2021, pp. 5148–5157.

- [19] G. H. Canal, M. R. O’Shaughnessy, C. J. Rozell, and M. A. Davenport, “Joint estimation of trajectory and dynamics from paired comparisons,” in *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2019, pp. 121–125.
- [20] A. K. Massimino and M. A. Davenport, “As you like it: Localization via paired comparisons,” *Journal of Machine Learning Research*, vol. 22, no. 186, pp. 1–39, 2021.
- [21] K. R. Davidson and S. J. Szarek, “Local operator theory, random matrices and banach spaces,” *Handbook of the geometry of Banach spaces*, vol. 1, no. 317-366, p. 131, 2001.
- [22] B. Mason, M. A. Rau, and R. Nowak, “Cognitive task analysis for implicit knowledge about visual representations with similarity learning methods,” *Cognitive Science*, vol. 43, no. 9, p. e12744, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cogs.12744>
- [23] M. A. Rau, B. Mason, and R. Nowak, “How to model implicit knowledge? similarity learning methods to assess perceptions of visual representations.” *International Educational Data Mining Society*, 2016.
- [24] S. E. Palmer and K. B. Schloss, “An ecological valence theory of human color preference,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 19, pp. 8877–8882, 2010.
- [25] K. B. Schloss, L. Lessard, C. Racey, and A. C. Hurlbert, “Modeling color preference using color space metrics,” *Vision Research*, vol. 151, pp. 99–116, 2018.
- [26] S. E. Palmer, K. B. Schloss, and J. Sammartino, “Visual aesthetics and human preference,” *Annual review of psychology*, vol. 64, pp. 77–107, 2013.
- [27] B. Kulis *et al.*, “Metric learning: A survey,” *Foundations and Trends® in Machine Learning*, vol. 5, no. 4, pp. 287–364, 2013.
- [28] H.-J. Ye, D.-C. Zhan, and Y. Jiang, “Fast generalization rates for distance metric learning,” *Machine Learning*, vol. 108, no. 2, pp. 267–295, 2019.
- [29] F. Liu, X. Huang, Y. Chen, and J. Suykens, “Fast learning in reproducing kernel krein spaces via signed measures,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 388–396.
- [30] M. Huai, H. Xue, C. Miao, L. Yao, L. Su, C. Chen, and A. Zhang, “Deep metric learning: The generalization analysis and an adaptive algorithm.” in *IJCAI*, 2019, pp. 2535–2541.
- [31] G. Canal, S. Fenu, and C. Rozell, “Active ordinal querying for tuplewise similarity learning,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, pp. 3332–3340, Apr. 2020. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/5734>
- [32] B. Mason, A. Tripathy, and R. Nowak, “Learning nearest neighbor graphs from noisy distance samples,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [33] L. Jain, K. G. Jamieson, and R. Nowak, “Finite sample prediction and recovery bounds for ordinal embedding,” in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/file/4e0d67e54ad6626e957d15b08ae128a6-Paper.pdf>
- [34] R. A. Bradley and M. E. Terry, “Rank analysis of incomplete block designs: I. the method of paired comparisons,” *Biometrika*, vol. 39, no. 3/4, pp. 324–345, 1952.
- [35] P. Rao and L. L. Kupper, “Ties in paired-comparison experiments: A generalization of the bradley-terry model,” *Journal of the American Statistical Association*, vol. 62, no. 317, pp. 194–204, 1967.
- [36] R. D. Luce, *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.
- [37] R. L. Plackett, “The analysis of permutations,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 24, no. 2, pp. 193–202, 1975.
- [38] L. L. Thurstone, “A law of comparative judgment.” *Psychological review*, vol. 34, no. 4, p. 273, 1927.
- [39] I. Melekhov, J. Kannala, and E. Rahtu, “Siamese network features for image matching,” in *2016 23rd international conference on pattern recognition (ICPR)*. IEEE, 2016, pp. 378–383.
- [40] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer, “An efficient boosting algorithm for combining preferences,” *Journal of machine learning research*, vol. 4, no. Nov, pp. 933–969, 2003.

- [41] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender, “Learning to rank using gradient descent,” in *Proceedings of the 22nd international conference on Machine learning*, 2005, pp. 89–96.
- [42] Z. Zheng, K. Chen, G. Sun, and H. Zha, “A regression framework for learning ranking functions using relative relevance judgments,” in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 287–294.
- [43] D. Chumbalov, L. Maystre, and M. Grossglauser, “Scalable and efficient comparison-based search without features,” in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 1995–2005. [Online]. Available: <https://proceedings.mlr.press/v119/chumbalov20a.html>
- [44] N. Houlsby, F. Huszar, Z. Ghahramani, and J. Hernández-lobato, “Collaborative gaussian processes for preference learning,” *Advances in neural information processing systems*, vol. 25, 2012.
- [45] B. Dubois, “Ideal point versus attribute models of brand preference: A comparison of predictive validity,” *ACR North American Advances*, 1975.
- [46] S. Janson, “Tail bounds for sums of geometric and exponential variables,” *Statistics & Probability Letters*, vol. 135, pp. 1–6, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167715217303711>
- [47] M. A. Davenport and J. Romberg, “An overview of low-rank matrix recovery from incomplete observations,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 608–622, 2016.
- [48] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [49] J. A. Tropp, “An introduction to matrix concentration inequalities,” *arXiv preprint arXiv:1501.01571*, 2015.
- [50] B. Laurent and P. Massart, “Adaptive estimation of a quadratic functional by model selection,” *The Annals of Statistics*, vol. 28, no. 5, pp. 1302–1338, 2000. [Online]. Available: <http://www.jstor.org/stable/2674095>

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes], in Section 6
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes], in Appendix A
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes], please see linked code repository as well as Sections 5 and F.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes], please see Section F
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] in Section F

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes], please see linked code repository.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]