Learning Preference Distributions From Pairwise Comparisons

Gokcan Tatli, Yi Chen and Ramya Korlakai Vinayak

Abstract

We introduce the problem of learning distribution of user preference over a population via pairwise comparison of a set of items. We consider the setting where each individual only answers one pairwise comparison. In this scenario, learning each user's individual preference is impossible. Hence the question of interest is what can we learn about the distribution of user preferences over the population? Since we are limited to binary answers to pairwise comparison queries, we focus on learning the mass of the underlying distribution in the regions defined by the intersecting midpoint hyperplanes between the pairs of items in the query set. We investigate this fundamental question in both 1-dimensional and higher dimensional settings under noiseless response settings. We show that the problem is identifiable in a 1-dimensional setting and provide recovery guarantees. We also show that the problem is not identifiable for higher dimensional settings. We propose using a regularized recovery for higher dimensional settings and provide guarantees on the total variation distance between the true mass in each of the regions and the distribution learned via regularized constrained optimization problem. We validate our findings through simulations and experiments on real datasets.

1 Introduction

Learning user preferences via pairwise comparison queries of type "Do you prefer item a or b?" is widely used in various applications, such as political science, to model voters' political preferences and to predict their voting behavior, and in recommendation systems, to model users' preferences for products or services [28, 14, 1, 21, 26]. The ideal point model [10, 22, 13, 32, 36, 9] often used in preference learning aims to model an individual's preferences over a set of alternatives by representing them as a single "ideal point" in a multidimensional space. The basic idea is that the closer an alternative is to the ideal point, the more preferred it is by the individual. The ideal point is usually estimated using data collected from the individual, such as their choices, ratings, or rankings.

Formally, the ideal point model is described as follows. Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m \in \mathbb{R}^d$ denote the known feature representations of m items. Let $\boldsymbol{u} \in \mathbb{R}^d$ represents the preference point of an individual. The answer to the preference query $\mathrm{Q}(a, b)$ is 1 if the individual prefers the item a over b and -1 otherwise. The ideal point model assumes that $\mathrm{Q}(a, b) = 1$ if the individual's preference point $\boldsymbol{u}$ is closer to the representation of item a, $\boldsymbol{x}_a$ than item b, $\boldsymbol{x}_b$. That is, $||\boldsymbol{x}_a - \boldsymbol{u}||_2 < ||\boldsymbol{x}_b - \boldsymbol{u}||_2$. The ideal point model with pairwise comparison queries is also well studied in the ranking literature. Several algorithms have been proposed to rank the preferences of the items of an individual [17, 22, 35, 30, 20].

In this work, we consider the problem of learning from diverse populations in preference learning. Many works on preference learning have focused on universal model, where the data from everyone is pooled in together to learn a single preference point on average for the population [18, 23, 5]. However, different individuals can have different preferences. While one can focus on learning an individual's preference separately, it takes $\mathcal{O}(d)$ queries

in $d-$dimensions, which can be prohibitively large due to cost, cognitive overload or privacy concerns. In many applications, learning a prior distribution over the user preference is more useful than learning each individual preference point. For example, if an ice cream company wants to come up with new flavors, knowing which regions of flavor profiles have more mass is beneficial in the discovery of new ice cream flavors. In a broader sense, learning a prior distribution of user preferences is advantageous in many downstream tasks, including, but not limited to, comparing preferences between distinct populations and addressing the cold start problem in recommendation systems. Information about a prior distribution can serve as a valuable prior in future inference tasks.

We investigate the problem of learning the distribution of user preferences using pairwise comparisons. In particular, our goal is to understand what we can learn about the distribution of user preferences with only one query per individual. Since we are limited to binary answers to pairwise comparison of m items, our goal is to learn the mass of the underlying distribution in the regions (polytopes) defined by the intersecting midpoint hyperplanes between the $\binom{m}{2}$ pairs of items. If we can query $\mathcal{O}(d \log d)$ or $\mathcal{O}(m^2)$ pairs, we can localize the user preference point to one of the regions. So, if we sample a large number of individuals from the population and if we can query each of them with a sufficiently large number of queries, we can build a histogram of the underlying distribution in these regions. However, due to limitations discussed before, querying large number of comparison pairs per individual can be prohibitive due privacy issues, limited interaction of individual with the platform, cognitive overload and cost.

Our contribution: We introduce the novel problem of learning the distribution of user preferences over the population via pairwise comparison queries with only one response per individual and investigate the fundamental questions of identifiability and recovery guarantees leading to the following contributions:

- We show that the problem is identifiable in 1D setting and is not identifiable in higher dimensional setting.
- For the 1D setting, we provide recovery guarantees for the mass in the regions defined by the intersection of hyperplanes at the mid-point of pairs of items used for queries.
- For the higher dimensional setting, we propose to use regularized recovery and provide guarantees on the total variation distance between the true mass in each of the regions and the estimated mass in terms of the regularization parameter and the interplay between the true mass and regularization.
- We provide experiments on synthetic datasets and real datasets that validate our results and observations.

Our work leads to several interesting open questions regarding learning from diverse populations in preference learning.

Related works: Learning a user preference point up to ϵ error with pairwise comparisons requires $\mathcal{O}(d/\epsilon \log d)$ queries under mild assumptions of coverage of query items on the space of preferences [25]. Recently, some works have considered the problem of simultaneously learning an unknown metric and a user preference point. Work in [36] proposed an alternating minimization algorithm for this problem. In [8], a more general problem setting with multiple individuals with different user preference points is considered. They propose a convex optimization based algorithm to learn the common unknown metric and the unknown different preferences of multiple users simultaneously with a sample complexity $\mathcal{O}(d)$ per person if the number of users is of the order of d. Basically, each of these requires at least $\mathcal{O}(d)$ queries per user as the goal in these works is to learn individual preference points.

However, querying multiple pairs per user is an hurdle in various situations, e.g. privacy concerns, as it requires tracking a user over time, cognitive overload, and cost in obtaining answers to multiple queries, especially in larger dimensional spaces.

In a recent work [33], the authors introduced the problem of learning distribution of preferences over the population using distance queries rather than pairwise comparisons. When the distance measurements are available, they show that the problem of learning the underlying distribution of preferences is identifiable and provide recovery guarantees for 1-dimensional setting. They also provide numerical simulations that suggest similar results might hold for higher dimensional settings and network structures.

2 Problem Setting

Let m denote the number of items and $S_m := \{\boldsymbol{x}_1, \boldsymbol{x}_2, ..., \boldsymbol{x}_m\}$ where $\boldsymbol{x}_i \in \mathbb{R}^d$ denotes the known feature representations of the items. Under the ideal point model, each individual preference is also modeled as a point in the same space. Let $P^\star$ denote the unknown underlying distribution of user preferences. Each individual l has an unknown preference $\boldsymbol{u}_l \in \mathbb{R}^d$. We assume that $\boldsymbol{u}_l \overset{i.i.d.}{\sim} P^\star$. We further assume that the answer to the pairwise query to an individual l is i if item $\boldsymbol{x}_i$ is closer to their preference point $\boldsymbol{u}_l$ than $\boldsymbol{x}_j$. That is, $y_{ij}^{(l)} = 1$ if $\text{dist}(\boldsymbol{x}_i, \boldsymbol{u}_l) < \text{dist}(\boldsymbol{x}_j, \boldsymbol{u}_l)$ and $y_{ij}^{(l)} = -1$ otherwise.

Given answers to pairwise queries of the form "do you prefer item i or item j?", our goal is to understand what we can learn about $P^\star$. Note that given a set of m items, there are $\binom{m}{2}$ possible pairs for comparison. Each pair of items creates a hyperplane perpendicular at the midpoint joining the two items. Let h_{ij} denote the hyperplane perpendicular to the midpoint of the pairs (i, j) with $i < j$. These hyperplanes carve out regions in $\mathbb{R}^d$ that are polytopes. Let $\mathcal{H}(S_m)$ denote the set of partitions of $\mathbb{R}^d$ that is created by the set of hyperplanes arising from items in S_m and $|\mathcal{H}(S_m)|$ denotes the number of these regions. Note that for m items in $\mathbb{R}^d$, $|\mathcal{H}(S_m)| = \mathcal{O}(m^{2d})$ [6].

Let $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ denote the true mass in each of the regions arising due to $P^\star$ which is the true underlying distribution of user preferences. Let $q^\star_{ij}$ denote the mass of $P^\star$ to the side of $\boldsymbol{x}_i$ of the hyperplane h_{ij} and $q^\star_{ji} = 1 - q^\star_{ij}$ is the mass to the other side of the same hyperplane h_{ij}. Let $\mathbf{q}^\star \in \mathbb{R}^{2\binom{m}{2}}$ denote the vector that stacks q_{ij}'s for the ordered pairs, and then the q_{ji}'s. We note that the mass on either side of the hyperplanes, $\mathbf{q}^\star$, can be written as a linear combination of the mass in the regions $\mathbf{p}^\star_{\mathcal{H}(S_m)}$. That is, we can construct the following linear system of equations

$$\mathbf{M}\ \mathbf{p}^\star_{\mathcal{H}(S_m)} = \mathbf{q}^\star, \tag{1}$$

where $\mathbf{M}$ is a $2\binom{m}{2} \times |\mathcal{H}(S_m)|$ binary matrix where in each row, the 1's indicate the regions that contribute to the side of the hyperplane. Each column of $\mathbf{M}$ corresponds to a region (polytope) created by the intersection of the hyperplanes. Figure 1(a) shows an example of partition of $\mathbb{R}^2$ with 3 hyperplanes $h_{1,2}, h_{3,4}$ and $h_{5,6}$. With the enumeration of the regions shown in Figure 1(a), we can construct the binary matrix $\mathbf{M}$ (see Figure 1(b)), where the first 3 rows represent regions corresponding to $h_{1,2}$ towards the side of item 1, $h_{3,4}$ towards the side of item 3 and $h_{5,6}$ towards the side of item 5 respectively. Similarly, the last rows 3 represent regions corresponding to the other side of each of the hyperplanes. We also note that each column gives positions of the corresponding region $\mathbf{p}_i$ in terms of hyperplanes $h_{1,2}, h_{3,4}$ and $h_{5,6}$.

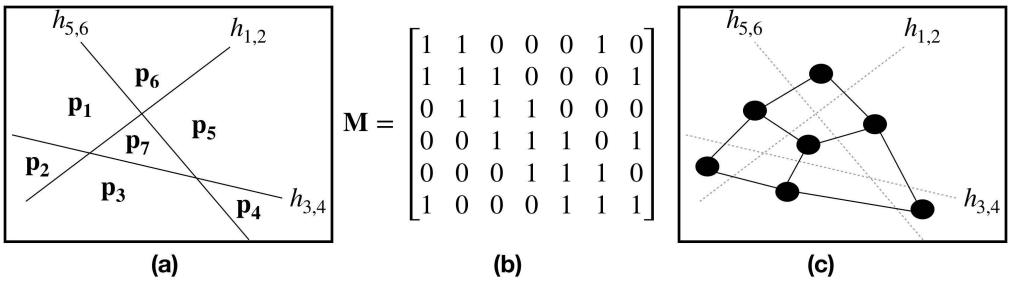

Figure 1: (a) Example of regions (polytopes) formed by intersection of three hyperplanes. (b) The corresponding matrix $\mathbf{M}$. (c) The corresponding graph where the regions are the nodes.

For each pair (i, j), by querying a random sample of people, we can estimate the amount of mass $P^\star$ has to the right and to the left of the hyperplane at the midpoint of (i, j), denoted by h_{ij}. Given these estimates, the goal is to estimate the mass of $P^\star$ in the regions of intersections of hyperplanes $\mathcal{H}(S_m)$.

3 Identifiability

In this section we discuss the identifiability results for the problem setting described in Section 2 for the 1-dimensional case and higher-dimensional cases. Note that $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ is *identifiable* if it is the unique probability vector of size $|\mathcal{H}(S_m)|$ that gives rise to $\mathbf{q}^\star$. So, $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ is *not identifiable* if there exist a $\mathbf{p} \neq \mathbf{p}^\star_{\mathcal{H}(S_m)}$ such that $\mathbf{Mp} = \mathbf{Mp}^\star_{\mathcal{H}(S_m)}$.

3.1 One Dimensional Setting

Consider the setting when the items are represented in a 1-dimensional Euclidean space. Given m items, there are $\binom{m}{2}$ pairs for comparison which divide the space into $\binom{m}{2} + 1$ partitions. Therefore, the linear system of equations in (1) has $\binom{m}{2} + 1$ columns.

As we describe in Section 2, we assume that the answer to the pairwise query to an individual l is i if item $\boldsymbol{x}_i$ is closer to the preference point $\boldsymbol{u}_l$ than $\boldsymbol{x}_j$. Basically, $y^{(l)}_{ij} = 1$ if $\text{dist}(\boldsymbol{x}_i, \boldsymbol{u}_l) < \text{dist}(\boldsymbol{x}_j, \boldsymbol{u}_l)$ in the noiseless case. Then, we consider the corresponding hyperplane h_{ij} and the mass to the left of h_{ij}, denoted by $q(h^<_{ij})$. Under the assumption that items and users are supported on a 1-dimensional Euclidean space, hyperplanes in $\mathcal{H}(S_m)$ are lines passing through midpoints $\frac{\boldsymbol{x}_i + \boldsymbol{x}_j}{2}$. Therefore, we can infer that $q(h^<_{ij}) = \sum_{k \in h^<_{ij}} p_k$, i.e., $q(h^<_{ij})$ becomes equal to the cumulative distribution function (CDF) of the distribution $\mathbf{p}^*$ evaluated at the point $\frac{\boldsymbol{x}_i + \boldsymbol{x}_j}{2}$. Therefore, the binary matrix $\mathbf{M}$ can be written as a concatenation of two triangular matrices, where one is a lower triangular matrix with all diagonals (except the last one) and elements below the main diagonal are 1 and the other is an upper triangular matrix with all diagonal elements (except the first one) and all elements above the main diagonal is 1. Since $\mathbf{M}$ is full column rank by construction, linear system of equations in (1) is identifiable in 1D and we can write the distribution $\mathbf{p}^\star_{\mathcal{H}(\mathcal{S}_m)}$, i.e., true mass in regions, as $\mathbf{p}^\star_{\mathcal{H}(\mathcal{S}_m)} = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\mathbf{q}^\star$ in terms of the true $\mathbf{q}^\star$.

3.2 Higher Dimensional Settings

In this section, we discuss the identifiability of the linear system of equations in (1) for dimensions $d \geq 2$, where items and users are supported on d-dimensional Euclidean space.

We show that the problem setting is not identifiable in higher dimensions and provide the following proposition.

Proposition 1. *For $d \geq 2$, the binary matrix M which of size $2\binom{m}{2} \times |\mathcal{H}(S_m)|$ has rank$(\mathbf{M}) = \binom{m}{2} + 1$ and the solution to the linear system of equations (1) is not unique and hence $\mathbf{p}^{\star}_{\mathcal{H}(\mathcal{S}_m)}$ is not identifiable. We note that, since rank$(\mathbf{M}) = \binom{m}{2} + 1$, for any $k > \binom{m}{2} + 1$, there exists at least another solution to the equation (1).*

As a result, we note that uniqueness cannot be guaranteed for all k-sparse distributions. Additionally, we make the following remark based on the fact that $\mathbf{M}$ is a column-regular matrix, i.e. each column of $\mathbf{M}$ has exactly the same number of 1's.

Remark 1. *Robust Null Space Property (RNSP) is proposed as a sufficient condition for basis pursuit approach in compressed sensing literature [16, 15]. Later, in [24], authors proposed sufficient conditions for a column-regular binary matrix to achieve RNSP, which are the best sufficient conditions for column-regular binary matrices to the best of our knowledge. According to Theorem 9 in [24], a column-regular binary matrix satisfies RNSP when $k< d_L/\rho$, where d_L is the number of 1's in each column and ρ is the maximum inner product among columns. Our binary matrix $\mathbf{M}$ is column-regular binary matrix with $\binom{m}{2}$ 1's in each column. Since there are neighboring regions, i.e., regions that has only one different coordinate, maximum inner product among columns is $\binom{m}{2}-1$. Therefore, RNSP is achieved when $k = 1$.*

3.2.1 Graph Regularization

In the face of non-identifiability, additional structural assumptions are needed for learning the mass in the regions, i.e., polytopes, $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$. We note that while $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$ is a $\mathcal{O}(m^{2d})$-dimensional probability vector, the entries corresponding to mass in regions have a geometry in the space $\mathcal{X} \subseteq \mathbb{R}^d$ (recall Figure 1(a)) that gives a notion of near-by and far-away regions. We construct a connected undirected graph with the polytopes as the nodes and two nodes are connected by an edge if they share a $(d-1)$-dimensional face between them (see Figure 1(c)). We propose using a graph regularizer (normalized by volume to account for differences in the sizes of the regions) to recover $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$. Intuitively, this means that we expect preferences to accumulate in spatially nearby regions (Figure 1(c)). Several works in signal recovery have used graph regularization to exploit local invariance in data as a side information and find a locally invariant representation of the data [4, 7, 19].

We note that this proposed graph structure can be constructed using the matrix $\mathbf{M}$. Recall that the rows of $\mathbf{M}$ correspond to hyperplanes and the columns correspond to the regions (polytopes) in $\mathcal{H}(S_m)$ providing a binary encoding for them by construction. That is, each entry of a given column of $\mathbf{M}$ determines which side of a hyperplane the corresponding region is located on. Therefore, there exists an edge between nodes corresponding to the regions that has only two different entries in their hyperplane coordinates, i.e., only if one pairwise comparison yields opposite results. Accordingly, neighboring regions have common $(d-1)$-dimensional faces in between. We define the weight matrix $\mathbf{W}$ for the graph regularization in the following way,

$$\mathbf{W}_{i,j} = \frac{\|\mathbf{M}_{:,i} - \mathbf{M}_{:,j}\|_1^{-1}}{\alpha_i \alpha_j}, \tag{2}$$

where $\alpha = [\alpha_1, \alpha_2, \ldots, \alpha_{|\mathcal{H}(S_m)|}]^T$ represent volumes of regions with corresponding mass $\mathbf{p} = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_{|\mathcal{H}(S_m)|}]^T$ respectively. We can define different weight matrices as long as

entries are inversely proportional to the distances between nodes. Heat kernel weighting [4], 0-1 weighting [7] are some of the widely used ones in the literature. We use $\mathbf{W}$ defined in equation (2) and form following graph Laplacian regularizer:

$$\frac{1}{2}\sum_{i=1}^{|\mathcal{H}(S_m)|}\sum_{j=1}^{|\mathcal{H}(S_m)|}|\mathbf{p}_i-\mathbf{p}_j|^2\mathbf{W}_{i,j} =: \mathbf{p}^T\mathbf{D}\mathbf{p}-\mathbf{p}^T\mathbf{W}\mathbf{p} =: \mathbf{p}^T\mathbf{L}\mathbf{p}, \tag{3}$$

where $\mathbf{D}_{i,i}=\sum_{j=1}^{|\mathcal{H}(S_m)|}\mathbf{W}_{i,j}$, $\mathbf{D}_{i,j}=0$ when $i\neq j$ and $\mathbf{L}=\mathbf{D}-\mathbf{W}$.

Essentially, we enforce the changes in neighboring regions to be smooth, which is similar to the local invariance property considered in [4, 7, 19]. We note that the eigenvectors of $\mathbf{L}$ are mutually orthogonal by spectral theory. Then, we conclude that orthogonal eigenvectors of nonzero eigenvalues force the candidate of the solution $\mathbf{p}$ to be close to the distribution $\overline{\alpha}$ by diminishing possible directions other than α, where $\overline{\alpha}$ is the normalized α. Note that regions in $\mathcal{H}(S_m)$ are not similar to an equally spaced grid. Therefore, we use a weighted version of the regularizer in (3) with respect to the volumes of the regions in $\mathcal{H}(S_m)$. We can write the whole objective function as follows:

$$f(\mathbf{p})=\|\mathbf{M}\mathbf{p}-\mathbf{q}\|_2^2+\lambda\mathbf{p}^T\mathbf{L}\mathbf{p}, \tag{4}$$

which induces the following optimization problem:

$$\begin{aligned}
\underset{\mathbf{p}\,\in\,\mathbb{R}^{|\mathcal{H}(S_m)|}}{\text{minimize}} \quad & \frac{1}{2}||\mathbf{M}\mathbf{p}-\mathbf{q}||_2^2+\frac{\lambda}{2}\mathbf{p}^T\mathbf{L}\mathbf{p} && \text{(5a)}\\
\text{subject to} \quad & \mathbf{1}^T\mathbf{p}=1, && \text{(5b)}\\
& p_i\geq 0, \quad i=1,\ldots.|\mathcal{H}(S_m)|. && \text{(5c)}
\end{aligned}$$

Proposition 2. *The convex optimization problem in (5) has a unique solution.*

We prove that the regularized problem is identifiable in Proposition 2 and show recovery guarantees in the next section.

4 Recovery Guarantees

In this section, we provide recovery guarantees for the 1-dimensional setting and higher-dimensional settings by bounding the total variation distance between $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ and the recovered mass in partitions $\mathcal{H}(S_m)$.

4.1 One Dimensional Setting

Recall from Section 3.1, that for 1-dimensional case, we have $\mathbf{p}^\star_{\mathcal{H}(S_m)}=(\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\mathbf{q}^\star$ with the true $\mathbf{q}^\star$. However, we usually only have finite samples and hence only have an estimate $\hat{\mathbf{q}}$ from observations. Therefore, we use the following constrained optimization to estimate $\mathbf{p}_{\mathcal{H}(S_m)}$ using $\hat{\mathbf{q}}$,

$$\begin{aligned}
\underset{\mathbf{p}\,\in\,\mathbb{R}^{|\mathcal{H}(S_m)|}}{\text{minimize}} \quad & \frac{1}{2}||\mathbf{M}\mathbf{p}-\hat{\mathbf{q}}||_2^2 && \text{(6a)}\\
\text{subject to} \quad & \mathbf{1}^T\mathbf{p}=1, && \text{(6b)}\\
& p_i\geq 0, \quad i=1,\ldots,|\mathcal{H}(S_m)|. && \text{(6c)}
\end{aligned}$$

Since $\mathbf{M}$ has full column rank in the 1-dimensional case, the objective function of the given optimization problem is strongly convex. Therefore, the above optimization problem in (6) is guaranteed to have a unique solution. We provide the following recovery guarantee for the noiseless setting in the 1-dimensional case.

Theorem 4.1. *When items and users are supported on a 1-dimensional Euclidean space, with probability at least 1−δ, the total variation distance between $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ and the recovered mass $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ in partitions $\mathcal{H}(S_m)$ is bounded as follows,*

$$TV(\mathbf{p}^\star_{\mathcal{H}(S_m)}, \hat{\mathbf{p}}_{\mathcal{H}(S_m)}) \leq \frac{\sqrt{|\mathcal{H}(S_m)|}\ cond(\mathbf{M}, 1)}{\sqrt{2\binom{m}{2}}} \max\left\{\sqrt{\frac{40}{n}}, \sqrt{\frac{25\log(3/\delta)}{n}}\right\}$$

where $cond(\mathbf{M}, 1)$ is the condition number of $\mathbf{M}$ with respect to l_1-norm and n is the total number of users.

4.2 Higher Dimensional Settings

Recall that in 2-dimensions and in higher-dimensional cases, the problem is not identifiable. In particular, the matrix $\mathbf{M}$ does not have a full column rank. Therefore, the optimization problem given in (6) corresponds to an underdetermined least squares problem with unit simplex constraint. In Section 3.2, we show that we cannot guarantee exact recovery of $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ even if $\mathbf{q}^\star$ is provided, since the linear system of equations in (1) is not identifiable, unlike th 1-dimensional setting. Therefore, in this section, we first provide the bound for the total variation distance between $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ and the recovered mass in partitions $\mathcal{H}(S_m)$ under the condition that $\mathbf{q}^\star$ is provided. Then, we extend that bound to the case where we only have estimated $\hat{\mathbf{q}}$. We propose using the optimization problem given in (5). This convex optimization problem can be rewritten as

$$\begin{aligned}
\underset{\mathbf{p} \in \mathbb{R}^{|\mathcal{H}(S_m)|}}{\text{minimize}} \quad & \frac{1}{2}\mathbf{p}^T(\mathbf{M}^T\mathbf{M} + \lambda\mathbf{L})\mathbf{p} - \mathbf{p}^T\mathbf{M}^T\mathbf{q}^\star && \text{(7a)}\\
\text{subject to} \quad & \mathbf{1}^T\mathbf{p} = 1, && \text{(7b)}\\
& p_i \geq 0, \quad i = 1, \ldots, |\mathcal{H}(S_m)|. && \text{(7c)}
\end{aligned}$$

Since $\mathbf{M}^T\mathbf{M} + \lambda\mathbf{L}$ is a positive definite matrix, there exists a unique real upper triangular matrix $\mathbf{R}$ with positive diagonal entries, where $\mathbf{M}^T\mathbf{M} + \lambda\mathbf{L} = \mathbf{R}^T\mathbf{R}$ by Cholesky decomposition. Therefore, we can write the problem in (7) as a constrained least squares optimization problem:

$$\begin{aligned}
\underset{\mathbf{p} \in \mathbb{R}^{|\mathcal{H}(S_m)|}}{\text{minimize}} \quad & \frac{1}{2}||\mathbf{R}\mathbf{p} - \mathbf{b}||_2^2 && \text{(8a)}\\
\text{subject to} \quad & \mathbf{1}^T\mathbf{p} = 1, && \text{(8b)}\\
& p_i \geq 0, \quad i = 1, \ldots, |\mathcal{H}(S_m)|, && \text{(8c)}
\end{aligned}$$

where $\mathbf{b} = \mathbf{R}^{-T}\mathbf{M}^T\mathbf{q}^\star$. Then, we provide the following recovery guarantees for the given problem.

Theorem 4.2. *Suppose items and users are supported on 2- or higher-dimensional Euclidean space and $\mathbf{q}^\star$ is provided. Then, with probability at least 1−δ, the total variation distance between $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ and the recovered mass in restricted regions, i.e. $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$, is bounded as follows,*

$$TV(\mathbf{p}^\star, \hat{\mathbf{p}}_{\mathcal{H}(S_m)}) \leq \frac{\lambda}{2}\sqrt{|\mathcal{H}(S_m)|}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2.$$

However, we usually do not have access to true $\mathbf{q}^\star$. We use an estimated value from the sampled data, $\hat{\mathbf{q}}$, and use the least squares optimization setting in (8) with $\mathbf{b} = \mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}}$. Similarly, we write the following theorem to provide recovery guarantees for the given least squares problem.

Theorem 4.3. *We suppose that items and users are supported on 2- or higher-dimensional Euclidean space. Then, with probability at least 1−δ, the total variation distance between* $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ *and recovered mass* $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ *in partitions* $\mathcal{H}(S_m)$ *is bounded as follows,*

$$\begin{aligned} TV(\mathbf{p}^\star_{\mathcal{H}(S_m)}, \hat{\mathbf{p}}_{\mathcal{H}(S_m)}) &\leq \frac{\lambda}{2}\sqrt{|\mathcal{H}(S_m)|}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2 \\ &+ \sqrt{\frac{|\mathcal{H}(S_m)|\binom{m}{2}}{2}}\|\mathbf{R}^{-1}\|_2\|\mathbf{M}^T\|_2 \max\left\{\sqrt{\frac{40}{n}}, \sqrt{\frac{25\log(3/\delta)}{n}}\right\} \end{aligned}$$

where n is the total number of users.

5 Bounds on the Mass

In this section, we show that we can obtain lower and upper bounds for each entry of $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ without requiring any additional assumptions. We first provide lower and upper bounds in Proposition 3 under the condition that $\mathbf{q}^\star$ is provided. Then, we extend these results to the scenario where we only have estimated $\hat{\mathbf{q}}$. We suppose that $\mathcal{K}_j$ is the position of rows of $\mathbf{M}$ whose j-th entry is 1.

Proposition 3. *Each entry of* $\mathbf{p}^\star_{\mathcal{H}(\mathcal{S}_m)}$ *can be bounded below and above as follows*

$$\max\left\{0, \max_{i\in\mathcal{K}_j}\mathbf{q}^\star_{a_ib_i} - \mathbf{M}^T_{i,:}\mathbf{Q}^j_0\right\} \leq \mathbf{p}^\star_{\mathcal{H}(S_m)_j} \leq \min_{i\in\mathcal{K}_j}\mathbf{q}^\star_{a_ib_i}, \tag{9}$$

where $\mathbf{M}_{i,:}$ *is the ith row of* $\mathbf{M}$ *and*

$$\mathbf{Q}^j_0 = [\min_{i\in\mathcal{K}_1}\mathbf{q}^\star_{a_ib_i}, \ldots \min_{i\in\mathcal{K}_{j-1}}\mathbf{q}^\star_{a_ib_i},\ 0,\ \min_{i\in\mathcal{K}_{j+1}}\mathbf{q}^\star_{a_ib_i}, \ldots \min_{i\in\mathcal{K}_l}\mathbf{q}^\star_{a_ib_i}]^T.$$

Note that in practice, we only have access to estimated $\hat{\mathbf{q}}$ from the observations. Therefore, we use the following lemma to obtain the lower and upper bounds for the entries of $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ in terms of the estimated $\hat{\mathbf{q}}$ and the confidence intervals.

Lemma 1. *(Lemma 3 in [11]) Let* $\mathbf{v} \in \mathbb{R}^z$ *be the probabilities corresponding to a multinomial random vector with support size z. Let* $\hat{\mathbf{v}}$ *denote the empirical estimate of these probability values from N i.i.d. samples drawn from* $\mathbf{p}$*. Then, for all* $\varepsilon \geq \sqrt{20z/N}$*,*

$$Pr(||\hat{\mathbf{v}} - \mathbf{v}||_1 > \varepsilon) \leq 3e^{-N\varepsilon^2/25}.$$

Lemma 1 together with Proposition 3 above yields the following proposition.

Proposition 4. *We suppose that each query is answered by the same number of people. Then, with probability at least* $1-\delta$*, each entry of* $\mathbf{p}^\star_{\mathcal{H}(\mathcal{S}_m)}$ *can be bounded below and above as follows:*

$$\max\left\{0, \max_{i\in\mathcal{K}_j}\hat{\mathbf{q}}_{a_ib_i} - \mathbf{M}^T_{i,:}\hat{\mathbf{Q}}^j_0 - (|\mathbf{M}_{i,:}|_1 + 1)\gamma\right\} \leq \mathbf{p}^\star_{\mathcal{H}(S_m)_j} \leq \min_{i\in\mathcal{K}_j}\hat{\mathbf{q}}_{a_ib_i} + \gamma \tag{10}$$

where $\gamma = \sqrt{\frac{\log\left(4\binom{m}{2}/\delta\right)}{2n_p}}$ *and* n_p *represents the number of people who answer each query.*

6 Numerical Results

We evaluate the proposed approach for both simulated and real data. We provide numerical simulations in one-dimensional and higher-dimensional settings and show the total variation distance between $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$ and the recovered mass in partitions $\mathcal{H}(S_m)$. Similarly, we evaluate our results on two different real datasets and provide recovered mass in partitions $\mathcal{H}(S_m)$ for the corresponding populations.

6.1 Simulations

We start with noiseless and noisy one-dimensional settings and use the following set of true distributions (see Figure 2). We observe the relationship between the number of items and

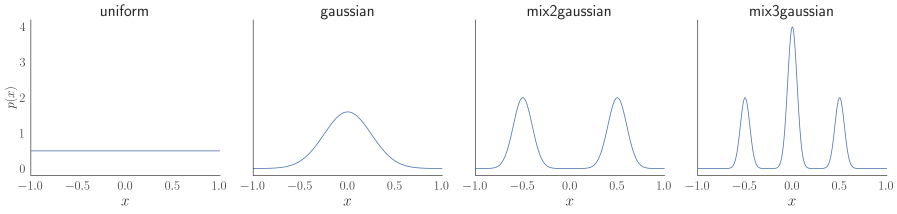

Figure 2: True distributions used in simulations.

the total variation error for the recovered mass in the partitions $\mathcal{H}(S_m)$ as well as between the number of people asked per query n_p and the error. We use CVXPY [12, 2] to solve the constrained least-squares problem given in (6) and its regularized version in (8). All simulations are run on Python 3.9.

1D Simulations: We fix the number of items, $m = 5$, and uniformly sample them from the interval of $[-1, 1]$. We vary the number of users per pair from 10^2 to 10^5 and repeat each setting 100 times. We employ a noise model where we flip a user's answer with probability 0.1. We provide the 1D noiseless and noisy simulation results in Figure 3.

As the number of people asked per query, n_p, increases, the TV distance between $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$ and $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ gets smaller as shown in Figure 3. In Figure 4, we fix the number of users

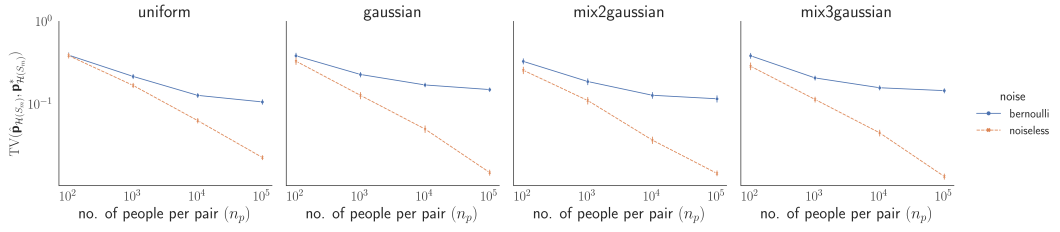

Figure 3: $\mathrm{TV}(\mathbf{p}^{\star}_{\mathcal{H}(S_m)}, \hat{\mathbf{p}}_{\mathcal{H}(S_m)})$ for 4 distributions in 1D

per pair to 1000 and provide the total variation distance between the recovered mass in partitions $\mathcal{H}(S_m)$ and the true underlying distributions with varying m for both 1D noisy and noiseless settings.

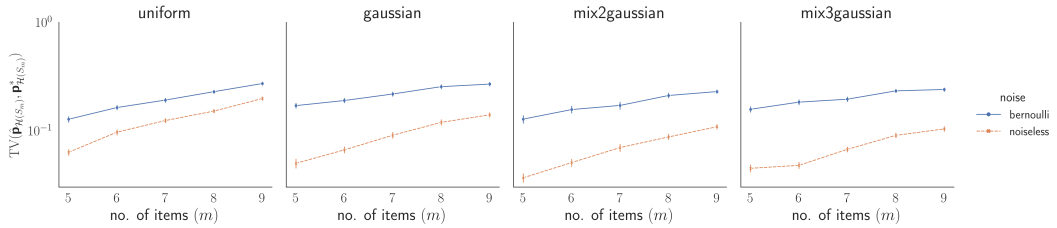

Figure 4: The total variation distance between the recovered distributions and the true distributions for the 4 distributions in noiseless and noisy 1D settings. Bernoulli line refers to coin-flipping noise.

2D Simulations: We use the same set of underlying distributions in 2D, where the items are drawn uniformly from $[-1, 1]^2$. Unlike in 1D, we cannot calculate the true mass of each region straightforwardly. Therefore, we employ the Monte Carlo method, where we sample 1 million points from the population distribution and use the fraction of points in each region as the corresponding true probability. In Figure 5, we provide the total variation error with a fixed m by varying the number of users per pair from 10^2 to 10^5 and repeating each setting 100 times.

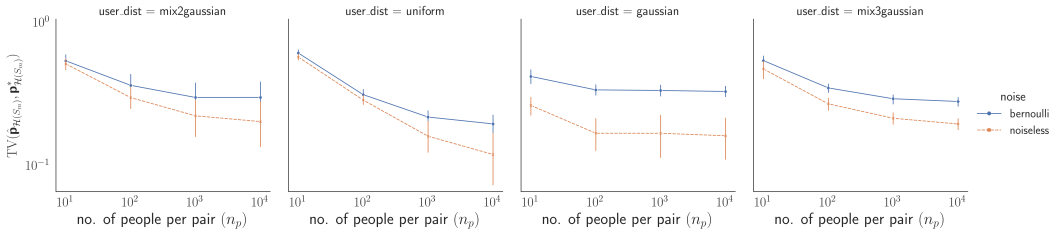

Figure 5: The total variation between the recovered distributions and the true distributions for the 4 distributions in noiseless and noisy settings 2D settings.

Similar to 1D, we fix the number of users per pair to 1000 and provide total variation error for both 2D noisy and noiseless settings with varying m in Figure 6.

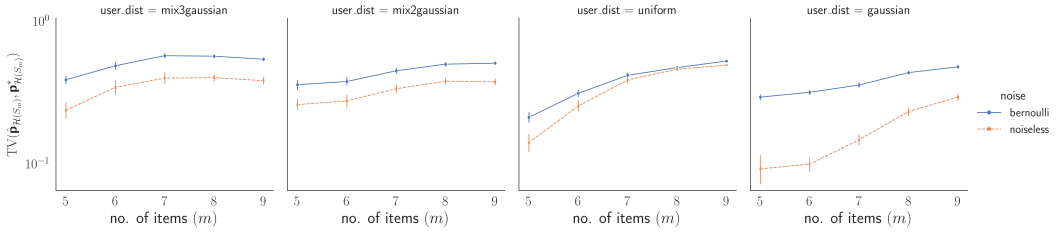

Figure 6: The total variation distance between the recovered distributions and the true distributions for the 4 distributions in noiseless and noisy 2D settings.

n**D Simulations**: To investigate the affects of dimensionality, we fix $m = 5$ and number of users per pair to 1000. Varying the dimension from 2 to 6, we plot the total variation error in Figure 7.

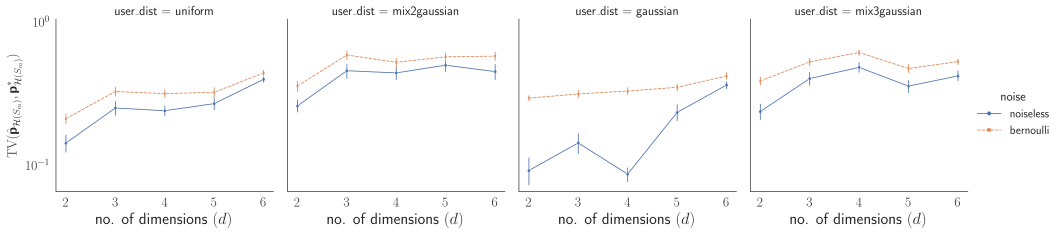

Figure 7: The total variation between the recovered distributions and the true distributions for the 4 distributions in noiseless and noisy 2D settings.

Bounds on the Mass: Referring to Section 5, we provide the following plot of the lower-upper bounds for $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$. We uniformly sample 5 items from $[-1, 1]^2$ and fix them so that we can repeat the algorithm 100 times.

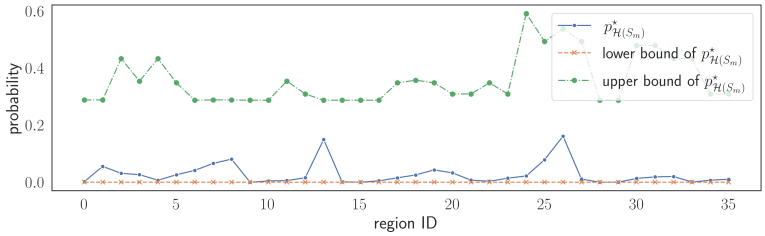

Figure 8: Upper and lower bounds for $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$.

6.2 Experiments

We use two different real datasets to validate our results.

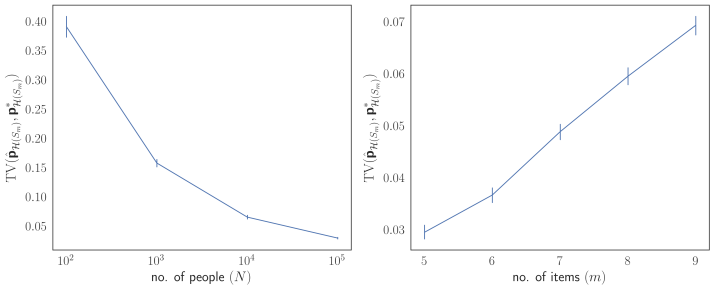

Figure 9: The total variation distance between recovered distributions and the true distributions for the colors dataset. Left: fixing number of users per pair. Right fixing m.

Colors dataset: Colors dataset is formed by answers to pairwise queries from 48 different users/respondents. Each person was asked to order 37 different colors, which enables us to have all possible pairwise queries for each person. In this dataset, each color is considered as a 3-dimensional vector in CIELAB color space (lightness, red vs. green, blue vs. yellow). For our experiment, we use the 1D user embedding of the color data set learned from [9]. Then, we project the CIELAB color space onto the 1D user embedding space.

We fix $m = 5$ to consider some subset of colors and uniformly sample $\{10, 100, 1000, 10000\}$ users from all 48 users for each pair to estimate $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$. We provide the total variation distance between recovered mass $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ and $\mathbf{p}^{\star}_{\mathcal{H}(S_m)}$ in Figure 9. Similarly, we fix the number of users per pair and show the change in the total variation error with varying numbers of colors, i.e. $m \in \{5, 6, 7, 8, 9\}$. To obtain a consistent result, we repeat each setting 100 times. Additionally, we provide Figure 10 to illustrate the true distribution of people in the 1D embedding space of the color dataset and the estimated distribution recovered by our method.

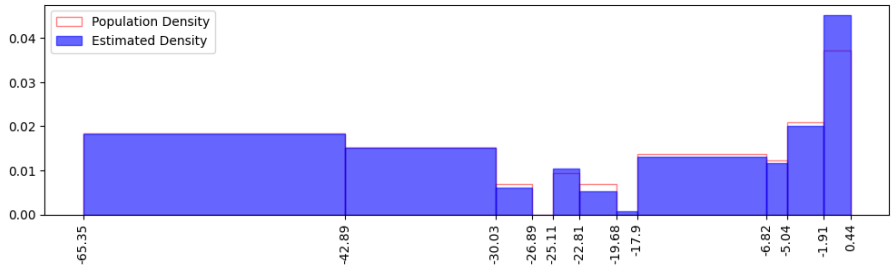

Figure 10: Mass recovered by our method vs. population mass

Zappos: UT Zappos50K dataset [37, 38] is a really large dataset with 50,025 catalog images from the website called Zappos.com. The entire data set contains images of shoes in different categories, such as shoes, sandals, slippers, and boots. We manually pick five shoes from this dataset (see Figure 11) and collect responses from 100 Amazon Mechanical Turk workers for each possible pairwise query to obtain a ground truth of the user distribution, and then for each pair of questions, we asked 100 people to estimate the user preference. We defer details of the setting to Appendix B. Figure 12 shows the ground truth and the estimated user preference distribution.

Figure 11: The 5 shoes we pick for pairwise comparison task on Amazon Mechanical Turk.

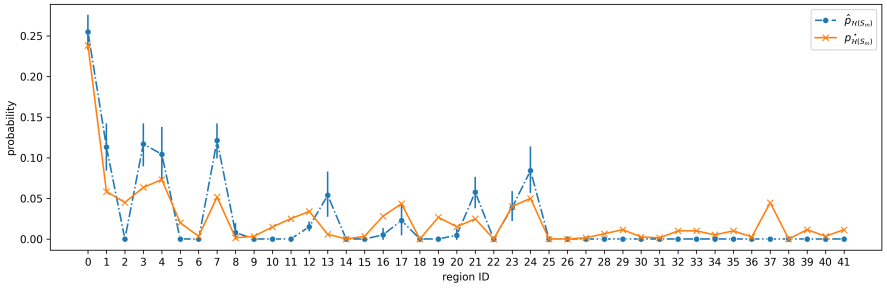

Figure 12: Preference distribution recovered by our method using the query results obtained from AMT vs. True Preference Distribution.

References

[1] Jens Abildtrup, E Audsley, M Fekete-Farkas, C Giupponi, Morten Gylling, P Rosato, and Mark Rounsevell. Socio-economic scenario development for the assessment of climate change impacts on agricultural land use: a pairwise comparison approach. *Environmental science & policy*, 9(2):101–115, 2006.

[2] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.

[3] AMT. Amazon mechanical turk. `https://www.mturk.com/`. Accessed: 2023-02-14.

[4] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL `https://proceedings.neurips.cc/paper/2001/file/f106b7f99d2cb30c3db1c3cc0fde9ccb-Paper.pdf`.

[5] Aniruddha Bhargava, Ravi Ganti, and Robert D. Nowak. Active algorithms for preference learning problems with multiple populations. *arXiv: Machine Learning*, 2016.

[6] R. C. Buck. Partition of space. *The American Mathematical Monthly*, 50(9):541–544, 1943. doi: 10.1080/00029890.1943.11991447. URL `https://doi.org/10.1080/00029890.1943.11991447`.

[7] Deng Cai, Xiaofei He, Jiawei Han, and Thomas S. Huang. Graph regularized nonnegative matrix factorization for data representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1548–1560, 2011. doi: 10.1109/TPAMI.2010.231.

[8] Greg Canal, Blake Mason, Ramya Koralkai Vinayak, and Rob Nowak. One for all: Simultaneous metric and preference learning over multiple users. *arXiv e-prints, arXiv:2207.03609*, 2022.

[9] Gregory Canal, Blake Mason, Ramya Korlakai Vinayak, and Robert Nowak. One for all: Simultaneous metric and preference learning over multiple users. *arXiv preprint arXiv:2207.03609*, 2022.

[10] Clyde H Coombs. Psychological scaling without a unit of measurement. *Psychological review*, 57(3):145, 1950.

[11] Luc Devroye. The equivalence of weak, strong and complete convergence in l1 for kernel density estimates. *The Annals of Statistics*, pages 896–904, 1983.

[12] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

[13] Cody Ding. Evaluating change in behavioral preferences: Multidimensional scaling single-ideal point model. *Measurement and Evaluation in Counseling and Development*, 49(1):77–88, 2016.

[14] John Fichtner. On deriving priority vectors from matrices of pairwise comparisons. *Socio-Economic Planning Sciences*, 20(6):341–345, 1986.

[15] Simon Foucart. Stability and robustness of $\ell 1$-minimizations with weibull matrices and redundant dictionaries. *Linear Algebra and its Applications*, 441:4–21, 2014. ISSN 0024-3795. doi: https://doi.org/10.1016/j.laa.2012.10.003. URL https://www.sciencedirect.com/science/article/pii/S0024379512007008. Special Issue on Sparse Approximate Solution of Linear Systems.

[16] Simon Foucart and Holger Rauhut. A mathematical introduction to compressive sensing. In *Applied and Numerical Harmonic Analysis*, 2013.

[17] Johannes Fürnkranz and Eyke Hüllermeier. Preference learning and ranking by pairwise comparison. In *Preference learning*, pages 65–82. Springer, 2010.

[18] Paul E. Green. Marketing applications of mds: Assessment and outlook. *Journal of Marketing*, 39(1):24–31, 1975. ISSN 00222429. URL http://www.jstor.org/stable/1250799.

[19] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742, 2006. doi: 10.1109/CVPR.2006.100.

[20] Reinhard Heckel, Max Simchowitz, Kannan Ramchandran, and Martin Wainwright. Approximate ranking from pairwise comparisons. In *International Conference on Artificial Intelligence and Statistics*, pages 1057–1066. PMLR, 2018.

[21] Daniel J Hopkins and Hans Noel. Trump and the shifting meaning of "conservative": Using activists' pairwise comparisons to measure politicians' perceived ideologies. *American Political Science Review*, 116(3):1133–1140, 2022.

[22] Kevin G Jamieson and Robert Nowak. Active ranking using pairwise comparisons. *Advances in neural information processing systems*, 24, 2011.

[23] Richard M. Johnson. Market segmentation: A strategic management tool. *Journal of Marketing Research*, 8(1):13–18, 1971. ISSN 00222437. URL http://www.jstor.org/stable/3149720.

[24] Mahsa Lotfi and Mathukumalli Vidyasagar. Compressed sensing using binary matrices of nearly optimal dimensions. *IEEE Transactions on Signal Processing*, 68:3008–3021, 2020. doi: 10.1109/TSP.2020.2990154.

[25] Andrew K Massimino and Mark A Davenport. As you like it: Localization via paired comparisons. *J. Mach. Learn. Res.*, 22:186–1, 2021.

[26] Shigehiro Oishi, Jungwon Hahn, Ulrich Schimmack, Phanikiran Radhakrishan, Vivian Dzokoto, and Stephen Ahadi. The measurement of values across cultures: A pairwise comparison approach. *Journal of research in Personality*, 39(2):299–305, 2005.

[27] Camille Poignard, Tiago Pereira, and Jan Philipp Pade. Spectra of laplacian matrices of weighted graphs: Structural genericity properties. *SIAM Journal on Applied Mathematics*, 78(1):372–394, 2018. doi: 10.1137/17M1124474. URL https://doi.org/10.1137/17M1124474.

[28] Thomas L Saaty and Luis G Vargas. The possibility of group choice: pairwise comparisons and merging functions. *Social Choice and Welfare*, 38(3):481–496, 2012.

[29] Rolf Schneider. *Convex Bodies: The Brunn–Minkowski Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition, 2013. doi: 10.1017/CBO9781139003858.

[30] Nihar B Shah and Martin J Wainwright. Simple, robust and optimal ranking from pairwise comparisons. *The Journal of Machine Learning Research*, 18(1):7246–7283, 2017.

[31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[32] Adish Singla, Sebastian Tschiatschek, and Andreas Krause. Actively learning hemimetrics with applications to eliciting user preferences. In *International Conference on Machine Learning*, pages 412–420. PMLR, 2016.

[33] Gokcan Tatli, Rob Nowak, and Ramya Korlakai Vinayak. Learning preference distributions from distance measurements. In *2022 58th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1–8, 2022. doi: 10.1109/Allerton49937.2022.9929404.

[34] Yingfan Wang, Haiyang Huang, Cynthia Rudin, and Yaron Shaposhnik. Understanding how dimension reduction tools work: an empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization. *The Journal of Machine Learning Research*, 22(1):9129–9201, 2021.

[35] Fabian Wauthier, Michael Jordan, and Nebojsa Jojic. Efficient ranking from pairwise comparisons. In *International Conference on Machine Learning*, pages 109–117. PMLR, 2013.

[36] Austin Xu and Mark Davenport. Simultaneous preference and metric learning from paired comparisons. *Advances in Neural Information Processing Systems*, 33:454–465, 2020.

[37] Aron Yu and Kristen Grauman. Fine-grained visual comparisons with local learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 192–199, 2014.

[38] Aron Yu and Kristen Grauman. Semantic jitter: Dense supervision for visual comparisons via synthetic images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5570–5579, 2017.

Gokcan Tatli
University of Wisconsin-Madison
Madison, WI, USA
Email: `gtatli@wisc.edu`

Yi Chen
University of Wisconsin-Madison
Madison, WI, USA
Email: `reid.chen@wisc.edu`

Ramya Korlakai Vinayak
University of Wisconsin-Madison
Madison, WI, USA
Email: `ramya@ece.wisc.edu`

A Proofs

A.1 Proof of Proposition 1

We first note that we can write following

$$\mathbf{e}_j = \prod_{i \in \mathcal{K}_j}^{\odot} \mathbf{M}_{i,:}, \quad j = 1, \ldots, |\mathcal{H}(S_m)|. \tag{11}$$

Then, considering the structure of matrix $\mathbf{M}$, we note that

$$\sum_{i=1}^{2\binom{m}{2}} \lambda_i \mathbf{M}_{i,:} = \sum_{i=1}^{\binom{m}{2}} (\lambda_i - \lambda_{\binom{m}{2}+i}) \mathbf{M}_{i,:} + \left(\sum_{i=1}^{\binom{m}{2}} \lambda_{\binom{m}{2}+i} \right) \mathbf{1}$$

If $\sum_{i=1}^{\binom{m}{2}} (\lambda_i - \lambda_{\binom{m}{2}+i}) \mathbf{M}_{i,:} + \left(\sum_{i=1}^{\binom{m}{2}} \lambda_{\binom{m}{2}+i} \right) \mathbf{1} = 0$ holds only when $\lambda_i - \lambda_{\binom{m}{2}+i} = 0$ for all $i = 1, \ldots, \binom{m}{2}$ and $\sum_{i=1}^{\binom{m}{2}} \lambda_{\binom{m}{2}+i} = 0$, we can claim that $\mathbf{1}$ and $\mathbf{M}_{i,:}$'s for $i = 1, \ldots, \binom{m}{2}$ are linearly independent. Therefore, we suppose that

$$\sum_{i=1}^{2\binom{m}{2}} \lambda_i \mathbf{M}_{i,:} = \sum_{i=1}^{\binom{m}{2}} (\lambda_i - \lambda_{\binom{m}{2}+i}) \mathbf{M}_{i,:} + \left(\sum_{i=1}^{\binom{m}{2}} \lambda_{\binom{m}{2}+i} \right) \mathbf{1} = 0$$

which yields that

$$\left(\sum_{i=1}^{2\binom{m}{2}} \lambda_i \mathbf{M}_{i,:} \right)^{\odot \binom{m}{2}} = 0 \tag{12}$$

Now we consider all elements in given Hadamard product and we can write following Lemma.

Lemma 2. *Given a binary matrix $\mathbf{M} \in \{0,1\}^{2\binom{m}{2} \times |\mathcal{H}(S_m)|}$ and real coefficients λ_i's, we can write following*

$$\left(\sum_{i=1}^{2\binom{m}{2}} \lambda_i \mathbf{M}_{i,:} \right)^{\odot \binom{m}{2}} = \sum_{j=1}^{2\binom{m}{2}} \left(\sum_{i \in \mathcal{K}_j} \lambda_i \right)^{\binom{m}{2}} \mathbf{e}_j,$$

where $\mathcal{K}_j$ is the position of rows of $\mathbf{M}$ whose jth entry is 1.

Lemma 3. *Given the binary matrix $\mathbf{M}$ in Section 2, for any $j \leq \binom{m}{2}$, we can find two columns $\mathbf{M}_{j_1}$ and $\mathbf{M}_{j_2}$ of matrix $\mathbf{M}$ such that only jth and $(\binom{m}{2}+j)$th entry of $\mathbf{M}_{j_1}$ and $\mathbf{M}_{j_2}$ differ.*

Proof: *Each hyperplane has to form neighboring regions by construction. Therefore, there has to be two columns $\mathbf{M}_{j_1}$ and $\mathbf{M}_{j_2}$ of $\mathbf{M}$ such that only jth and $(\binom{m}{2}+j)$th entry of $\mathbf{M}_{j_1}$ and $\mathbf{M}_{j_2}$ differ. To understand it better, we can consider a scenario where we delete any jth row of matrix $\mathbf{M}$ and call $\mathbf{M}^j$ to this new matrix. $\mathbf{M}^j$ has to have a pair of same columns. Otherwise, we would conclude that jth hyperplane does not form new regions, which is not possible. We can also refer to the fact that each hyperplane has to divide at least one previous region into two, when we consider adding one hyperplane at each time.*

Then, from Lemma 2, (12) yields that

$$\sum_{j=1}^{2\binom{m}{2}} \left(\sum_{i \in \mathcal{K}_j} \lambda_i \right)^{\binom{m}{2}} \mathbf{e}_j = 0,$$

which happens only if

$$\sum_{i \in \mathcal{K}_j} \lambda_i = 0, \quad j = 1, \dots, |\mathcal{H}(S_m)|.$$

From Lemma 3, it follows that we can find two numbers j_1 and j_2 for all $j = 1, \dots, |\mathcal{H}(S_m)|$ such that

$$\sum_{i \in \mathcal{K}_{j_1}} \lambda_i = \sum_{i \in \mathcal{K}_{j_2}} \lambda_i = 0,$$

where $j \in \mathcal{K}_{j_1}$, $\binom{m}{2} + j \in \mathcal{K}_{j_2}$ and $\mathcal{K}_{j_1} \backslash \{j\} = \mathcal{K}_{j_2} \backslash \{\binom{m}{2} + j\}$. Therefore, we conclude that $\lambda_j = \lambda_{\binom{m}{2}+j}$ for all $j = 1, \dots, l$. Now, (12) implies that $\sum_{i=1}^{\binom{m}{2}} \lambda_{\binom{m}{2}+i} = 0$, which proves that $\text{rank}(\mathbf{M}) = \binom{m}{2} + 1$. For the rest of the proof, we recall that half of the rows among $2\binom{m}{2}$ rows of $\mathbf{M}$ reflect the mass on the other side of each hyperplane. Basically, adding a row of all ones makes half of the rows redundant, since the rows representing the mass on the other side of each hyperplane are just flipped versions of rows representing the mass on the first side, i.e., $\mathbf{M}_{i+\binom{m}{2},:} = \mathbf{1}^T - \mathbf{M}_{i,:}$. We call $\mathbf{M}_{\text{half}}$ to the simplified version of $\mathbf{M}$. Then, we note that $\text{rank}(\mathbf{M}_{\text{half}}) = \text{rank}(\mathbf{M}) = \binom{m}{2} + 1$. Therefore, we cannot make further simplifications on $\mathbf{M}$ to get redundant rows.

Now, we consider the simplified version $\mathbf{M}_{\text{half}}$ and recall that any solution $\mathbf{p}^\star_{\mathcal{H}(S_m)}$ to the linear system of equations in (1) has to be in the probability simplex. Therefore, all possible $\hat{\mathbf{q}}_{\text{half}}$ vectors belong to the convex hull of columns of matrix $\mathbf{M}_{\text{half}}$, which we call $\text{conv}(\mathbf{M}_{\text{half}})$. Then, we apply Carathéodory's Theorem and write following expression. *Each element in conv($\mathbf{M}_{half}$) can be written as a convex combination of at most $\binom{m}{2} + 1$ columns of $\mathbf{M}_{half}$.* We can easily observe that the same property also applies to $\text{conv}(\mathbf{M})$ and $\hat{\mathbf{q}}$, as they share a one-to-one correspondence with $\mathbf{M}_{\text{half}}$ and $\hat{\mathbf{q}}_{\text{half}}$, respectively.

A.2 Graph Regularization

In this section, we discuss about the graph regularization that we proposed using in Section 3.2.1. We provide a standard graph regularizer without using volume weighting here to give a better intuition about graph regularizers and why we used volume weighting in Section 3.2.1. We start by defining following weight matrix $\mathbf{W}^{\text{unif}}$:

$$\mathbf{W}^{\text{unif}}_{i,j} = \|\mathbf{M}_{:,i} - \mathbf{M}_{:,j}\|_1^{-1}, \tag{13}$$

which is the inverse of the Hamming distance between nodes i and j. Accordingly, we can write following graph Laplacian regularizer:

$$\begin{aligned} R &= \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} |\mathbf{p}_i - \mathbf{p}_j|^2 \mathbf{W}^{\text{unif}}_{i,j} \\ &= \sum_{i=1}^{n} \mathbf{p}_i \mathbf{p}_i \mathbf{D}^{\text{unif}}_{i,i} - \sum_{i=1}^{n} \sum_{j=1}^{n} \mathbf{p}_i \mathbf{p}_j \mathbf{W}^{\text{unif}}_{i,j} = \mathbf{p}^T \mathbf{D}^{\text{unif}} \mathbf{p} - \mathbf{p}^T \mathbf{W}^{\text{unif}} \mathbf{p} = \mathbf{p}^T \mathbf{L}^{\text{unif}} \mathbf{p}, \end{aligned}$$

where $\mathbf{D}^{\text{unif}}_{i,i} = \sum_{j=1}^{n} \mathbf{W}^{\text{unif}}_{i,j}$, $\mathbf{D}^{\text{unif}}_{i,j} = 0$ when $i \neq j$ and $\mathbf{L}^{\text{unif}} = \mathbf{D}^{\text{unif}} - \mathbf{W}^{\text{unif}}$. Now, suppose that the spectral decomposition of $\mathbf{L}^{\text{unif}}$ can be written as $\mathbf{L}^{\text{unif}} = \sum_{i=1}^{l} \mu_i \mathbf{v}_i \mathbf{v}_i^T$, where $\mathbf{v}_i$'s are eigenvectors and μ_i's are the corresponding eigenvalues. We now further elaborate on spectral properties of Laplacian matrices and use following Lemma.

Lemma 4. *Graph Laplacian matrices are positive semi-definite by the Gershgorin circle theorem. Furthermore, the eigenvectors of the Laplacian matrix $\mathbf{L}^{unif}$ corresponding to zero eigenvalues are spanned by $\mathbf{1}$, which is referred to as constant vectors in [27].*

Then, we can rewrite Laplacian regularizer in (3) as

$$\mathbf{p}^T \mathbf{A}^T \mathbf{L}^{\text{unif}} \mathbf{A} \mathbf{p} = \mathbf{p}^T \sum_{i=1}^{l} \mu_i \mathbf{A}^T \mathbf{v}_i \mathbf{v}_i^T \mathbf{A} \mathbf{p} = \sum_{i=1}^{l} \mu_i (\mathbf{p}^T (\mathbf{A}^T \mathbf{v}_i))^2,$$

where $\mathbf{A}$ is a diagonal matrix with the entries $\mathbf{A}_{i,i} = \frac{1}{\alpha_i}$ and $\sum_i \mathbf{A}_{i,i} = 1$. Laplacian regularizer $\mathbf{L} = \mathbf{A}^T \mathbf{L}^{\text{unif}} \mathbf{A}$ penalizes $\mathbf{p}$ so that potential $\mathbf{p}$ values correlated to vectors $\mathbf{A}^T \mathbf{v}_i$'s are diminished. We can rephrase it as follows: regularizer penalizes $\mathbf{p}$ so that potential $\mathbf{A}^{-1}\mathbf{p}$ values correlated to eigenvectors $\mathbf{v}_i$'s are diminished. Therefore, $\mathbf{v}_i$'s corresponding to larger eigenvalues cause more penalty. From Lemma 4, it follows that Laplacian matrix $\mathbf{L}$ corresponding to zero eigenvalues are spanned by $\mathbf{A}^{-1}\mathbf{1}$. In [27], authors also point out that the multiplicity of the eigenvalue is equal to the number of connected components in the graph, which is clearly 1 in our graph structure induced by $\mathbf{M}$, since the regions in $\mathcal{H}(S_m)$ are connected. We note that the eigenvectors of $\mathbf{L}^{\text{unif}}$ are mutually orthogonal by spectral theory. We observe that orthogonal eigenvectors of nonzero eigenvalues would force the candidate of the solution $\mathbf{p}$ to be similar to uniform distribution by punishing possible directions other than $\mathbf{1}$. However, we note that regions in $\mathcal{H}(\mathcal{S}_m)$ are not similar to an equally spaced grid. Therefore, we use a weighted version of the regularizer in (3) with respect to the volumes of the regions in $\mathcal{H}(S_m)$ instead of $\mathbf{L}^{\text{unif}}$ and punish possible directions other than $\mathbf{A}^{-1}\mathbf{1}$, i.e. $\overline{\alpha}$.

A.3 Proof of Proposition 2

Recall that $f(\mathbf{p})$ is the objective function $\frac{1}{2}||\mathbf{M}\mathbf{p} - \hat{\mathbf{q}}||_2^2 + \frac{\lambda}{2}\mathbf{p}^T\mathbf{L}\mathbf{p}$. If we guarantee that

$$\frac{\partial^2 f}{\partial \mathbf{p}^2} = 2\mathbf{M}^T\mathbf{M} + 2\lambda\mathbf{L} \succ 0, \tag{14}$$

we can show that solution to the convex optimization problem in (5) is unique. Therefore, we first focus on matrix $\mathbf{L}$. From Lemma 4, null space of $\mathbf{L}^{\text{unif}}$ is spanned by $\mathbf{1}$. Since $\mathbf{A}$ is a full rank matrix, null space of $\mathbf{L} = \mathbf{A}^T\mathbf{L}^{\text{unif}}\mathbf{A}$ is spanned by $\mathbf{A}^{-1}\mathbf{1}$. All entries of $\mathbf{A}^{-1}\mathbf{1}$ are nonnegative since $\mathbf{A}^{-1}$ is a diagonal matrix with nonnegative entries. Now, we have following

$$\begin{aligned} \mathbf{M}^T\mathbf{M} &\succeq 0, \\ \mathbf{L} &\succeq 0, \\ \mathbf{M}^T\mathbf{M} + \lambda\mathbf{L} &\succeq 0. \end{aligned}$$

If $ker(\mathbf{M}^T\mathbf{M}) \neq ker(\mathbf{L})$, we can guarantee that $\mathbf{M}^T\mathbf{M} + \lambda\mathbf{L} \succ 0$. $\mathbf{M}^T\mathbf{M}$ is already positive semidefinite and $\mathbf{A}^{-1}\mathbf{1}$ cannot be an eigenvector for $\mathbf{M}^T\mathbf{M}$, since all nonzero entries of $\mathbf{M}^T\mathbf{M}$ have same sign. Therefore, $\mathbf{M}^T\mathbf{M} + \lambda\mathbf{L}$ is always positive definite.

A.4 Proof of Theorem 4.1

Let $\mathbf{p}_{sol}$ represents the solution to the optimization setting in (6), which is a constrained least square optimization problem with unit simplex constraint. We note that $\mathbf{M}\mathbf{p}_{sol}$ is the projection of $\hat{\mathbf{q}}$ onto the closed convex set $C_{\mathbf{M}}$ under ℓ_2 distance, which we call $\text{Proj}_{C_M}(\hat{\mathbf{q}})$, where

$$C_{\mathbf{M}} := \text{conv}(\mathbf{M}e_1, \ldots, \mathbf{M}e_{|\mathcal{H}(S_m)|}).$$

Now, we can write

$$\begin{aligned} ||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_2 &= ||\mathbf{M}^\dagger(\mathbf{q}^\star - \text{Proj}_{C_M}(\hat{\mathbf{q}}))||_2 \\ &\leq ||\mathbf{M}^\dagger||_2||\mathbf{q}^\star - \text{Proj}_{C_M}(\hat{\mathbf{q}})||_2 \\ &\overset{(a)}{\leq} ||\mathbf{M}^\dagger||_2||\mathbf{q}^\star - \hat{\mathbf{q}}||_2 \\ &\leq ||\mathbf{M}^\dagger||_2||\mathbf{q}^\star - \hat{\mathbf{q}}||_1, \end{aligned} \tag{15}$$

where the inequality (a) is due the fact that projection onto closed convex sets is contracting (Thm. 1.2.2.[29]). Then, we note that 2 $\text{TV}(\mathbf{p}^\star_{\mathcal{H}(S_m)}, \mathbf{p}_{sol}) = ||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_1$, and use $l_1 - l_2$ norm inequality to obtain the following from (15),

$$\begin{aligned} ||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_1 &\leq \sqrt{|\mathcal{H}(S_m)|}||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_2 \\ &\leq \sqrt{|\mathcal{H}(S_m)|}||\mathbf{M}^\dagger||_2||\hat{\mathbf{q}} - \mathbf{q}^\star||_2 \\ &\leq \sqrt{2|\mathcal{H}(S_m)|\binom{m}{2}}||\mathbf{M}^\dagger||_1||\hat{\mathbf{q}} - \mathbf{q}^\star||_1 \\ &= \sqrt{2|\mathcal{H}(S_m)|\binom{m}{2}}\frac{\text{cond}(\mathbf{M}, 1)}{||\mathbf{M}||_1}||\hat{\mathbf{q}} - \mathbf{q}^\star||_1. \end{aligned}$$

Note that the term $||\hat{\mathbf{q}} - \mathbf{q}^\star||_1$ is the sum of l_1- distances between the empirical and the true conditional distributions of pairwise comparisons for items in the query set S. We then note that $||\mathbf{M}||_1 = \binom{m}{2}$ and use the bound on the l_1-norm between empirical distribution and the true distribution for discrete distributions on finite support from Lemma 1 to complete the proof.

A.5 Proof of Theorem 4.2

Multiplication of each element in unit simplex with matrix $\mathbf{R}$ defines following closed convex set,

$$C_{\mathbf{R}} := \text{conv}(\mathbf{R}e_1, \mathbf{R}e_2 \ldots, \mathbf{R}e_{|\mathcal{H}(S_m)|}).$$

Then, the unique solution $\mathbf{p}_{sol}$ to the optimization setting in (8) can be expressed as

$$\mathbf{p}_{sol} = \mathbf{R}^{-1}\text{Proj}_{C_{\mathbf{R}}}(\mathbf{b}), \tag{16}$$

where $\mathbf{b} = \mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*$. Therefore,

$$\mathbf{R}\mathbf{p}_{sol} = \text{Proj}_{C_{\mathbf{R}}}(\mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)}). \tag{17}$$

We can write

$$\begin{aligned}
\|\mathbf{p}_{sol} - \mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 &\leq \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}_{sol} - \mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \\
&= \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathrm{Proj}_{C_{\mathbf{R}}}(\mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)})\|_2 \quad (18) \\
&\leq \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \quad (19) \\
&= \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}(\mathbf{R}^T\mathbf{R} - \lambda\mathbf{L})\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \\
&\leq \|\mathbf{R}^{-1}\|_2 \|\lambda\mathbf{R}^{-T}\mathbf{L}\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \\
&\leq \lambda\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2, \quad (20)
\end{aligned}$$

where (18) is from (17), (19) is due to contracting property of projection onto closed convex sets and (20) holds since $\|\mathbf{A}\|_2 \leq 1$ by construction. Then, we use $l_1 - l_2$ norm inequality and (20) to obtain following

$$\begin{aligned}
\mathrm{TV}(\mathbf{p}^\star_{\mathcal{H}(S_m)}, \mathbf{p}_{sol}) = \frac{1}{2}||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_1 &\leq \frac{\sqrt{|\mathcal{H}(S_m)|}}{2}||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_2 \\
&\leq \frac{\lambda}{2}\sqrt{|\mathcal{H}(S_m)|}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2.
\end{aligned}$$

A.6 Proof of Theorem 4.3

Following similar lines with the proof of Theorem 4.2, we can write

$$\mathbf{R}\mathbf{p}_{sol} = \mathrm{Proj}_{C_{\mathbf{R}}}(\mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}}).$$

We, again, start with bounding ℓ_2 norm error and write

$$\begin{aligned}
\|\mathbf{p}_{sol} - \mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 &\leq \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}_{sol} - \mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \\
&= \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathrm{Proj}_{C_{\mathbf{R}}}(\mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}})\|_2 \\
&\leq \|\mathbf{R}^{-1}\|_2 \|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}}\|_2 \\
&\leq \|\mathbf{R}^{-1}\|_2 (\|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)} \quad (21) \\
&+ \mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}}\|_2) \\
&\leq \|\mathbf{R}^{-1}\|_2 (\|\mathbf{R}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)}\|_2 \quad (22) \\
&+ \|\mathbf{R}^{-T}\mathbf{M}^T\mathbf{M}\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{R}^{-T}\mathbf{M}^T\hat{\mathbf{q}}\|_2) \\
&\leq \|\mathbf{R}^{-1}\|_2^2(\lambda\|\mathbf{L}\|_2 + \|\mathbf{M}^T\|_2\|\mathbf{q}^\star - \hat{\mathbf{q}}\|_2), \quad (23)
\end{aligned}$$

where (23) is from the fact that $\mathbf{M}\mathbf{p}_{\mathcal{H}(S_m)} = \mathbf{q}^\star$ and (20). Then, we can simply write following inequalities

$$\begin{aligned}
\mathrm{TV}(\mathbf{p}^\star_{\mathcal{H}(S_m)}, \mathbf{p}_{sol}) &= \frac{1}{2}||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_1 \\
&\leq \frac{\sqrt{|\mathcal{H}(S_m)|}}{2}||\mathbf{p}^*_{\mathcal{H}(S_m)} - \mathbf{p}_{sol}||_2 \\
&\leq \frac{\sqrt{|\mathcal{H}(S_m)|}}{2}\|\mathbf{R}^{-1}\|_2^2(\lambda\|\mathbf{L}\|_2 + \|\mathbf{M}^T\|_2\|\mathbf{q}^\star - \hat{\mathbf{q}}\|_2)\|\mathbf{L}\|_2 \\
&\leq \frac{\lambda}{2}\sqrt{|\mathcal{H}(S_m)|}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2 + \frac{\sqrt{|\mathcal{H}(S_m)|}}{2}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{M}^T\|_2\|\mathbf{q}^\star - \hat{\mathbf{q}}\|_2 \\
&\leq \frac{\lambda}{2}\sqrt{|\mathcal{H}(S_m)|}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{L}\|_2 + \frac{\sqrt{|\mathcal{H}(S_m)|r}}{2}\|\mathbf{R}^{-1}\|_2^2\|\mathbf{M}^T\|_2\|\mathbf{q}^\star - \hat{\mathbf{q}}\|_1.
\end{aligned}$$

Lastly, we apply Lemma 1 and complete the proof.

A.7 Proof of Proposition 3

We refer to the fact that each column of matrix $\mathbf{M}$ gives us position of restricted regions in terms of binary coordinates with respect to hyperplanes. By construction, corresponding entries of $\mathbf{q}^\star$, which we denote as $\mathbf{q}^\star_{\mathcal{K}_j}$, have $\mathbf{p}^\star_{\mathcal{H}(S_m)_j}$ as a nonnegative summand. Therefore, we can deduce following

$$\mathbf{p}^\star_{\mathcal{H}(S_m)_j} \leq \min_{i \in \mathcal{K}_j} \mathbf{q}^\star_{a_i b_i}, \quad j = 1, \ldots, |\mathcal{H}(S_m)|. \tag{24}$$

Using this fact and nonnegativity of entries of matrix $\mathbf{M}$, we can write following set of inequalities

$$\mathbf{M} \underbrace{\begin{bmatrix} \min_{i \in \mathcal{K}_1} \mathbf{q}^\star_{a_i b_i} \\ \vdots \\ \min_{i \in \mathcal{K}_{j-1}} \mathbf{q}^\star_{a_i b_i} \\ \mathbf{p}^\star_{\mathcal{H}(S_m)_j} \\ \min_{i \in \mathcal{K}_{j+1}} \mathbf{q}^\star_{a_i b_i} \\ \vdots \\ \min_{i \in \mathcal{K}_l} \mathbf{q}^\star_{a_i b_i} \end{bmatrix}}_{\mathbf{Q}^j} \geq \mathbf{q}^\star, \quad j = 1, \ldots, |\mathcal{H}(S_m)| \tag{25}$$

which enables us to lower bound each entry $\mathbf{p}^\star_{\mathcal{H}(S_m)_j}$ for $j = 1, \ldots, |\mathcal{H}(S_m)|$. Here, $\mathbf{Q}^j$ represents the vector constructed with minimum $\mathbf{q}_{a_i b_i}$'s over different sets and $\mathbf{p}^\star_{\mathcal{H}(S_m)_j}$. We let $\mathbf{Q}^j_0$ represent the vector that jth entry of $\mathbf{Q}^j$ is replaced with 0 and note that each inequality in (25) can be rewritten as follows

$$\mathbf{M}^T_{k,:} \mathbf{Q}^j \geq \mathbf{q}^\star_{a_k b_k}, \quad k = 1, \ldots, 2\binom{m}{2},$$

where the vector $\mathbf{M}_{k,:}$ represents kth row of $\mathbf{M}$. Equivalently, using standard basis vectors, we can write

$$\mathbf{p}^\star_{\mathcal{H}(S_m)_j} \mathbf{M}^T_{k,:} \mathbf{e}_j \geq \mathbf{q}^\star_{a_k b_k} - \mathbf{M}^T_{k,:} \mathbf{Q}^j_0, \quad k = 1, \ldots, 2\binom{m}{2},$$

which provides us following bound

$$\mathbf{p}^\star_{\mathcal{H}(S_m)_j} \geq \max\{\max_{i \in \mathcal{K}_j} \mathbf{q}^\star_{a_i b_i} - \mathbf{M}^T_{i,:} \mathbf{Q}^j_0, 0\}, \quad j = 1, \ldots, |\mathcal{H}(S_m)|. \tag{26}$$

Combining (24) and (26), we obtain following expression

$$\max_{i \in \mathcal{K}_j} \mathbf{q}^\star_{a_i b_i} - \mathbf{M}^T_{i,:} \mathbf{Q}^j_0 \leq \mathbf{p}^\star_{\mathcal{H}(S_m)_j} \leq \min_{i \in \mathcal{K}_j} \mathbf{q}^\star_{a_i b_i}. \tag{27}$$

A.8 Proof of Proposition 4

For any $\mathbf{q}^\star_{a_i b_i}$, we can say that

$$|\hat{\mathbf{q}}_{a_i b_i} - \mathbf{q}^\star_{a_i b_i}| \leq \sqrt{\frac{\log(2/\delta')}{2n_i}} \tag{28}$$

holds with probability at least $1-\delta'$ by Hoeffding's Inequality. Therefore, we want to bound the probability that

$$|\hat{\mathbf{q}}_{a_i b_i} - \mathbf{q}^\star_{a_i b_i}| \geq \sqrt{\frac{\log(2/\delta')}{2n_i}}$$

holds at least for one i. Therefore, we want to bound

$$\Pr\left(\bigcup_i \left\{ |\hat{\mathbf{q}}_{a_i b_i} - \mathbf{q}^\star_{a_i b_i}| \geq \sqrt{\frac{\log(2/\delta')}{2n_i}} \right\}\right) \leq \sum_i \Pr\left(\left\{ |\hat{\mathbf{q}}_{a_i b_i} - \mathbf{q}^\star_{a_i b_i}| \geq \sqrt{\frac{\log(2/\delta')}{2n_i}} \right\}\right) \quad (29)$$

$$\leq 2\binom{m}{2}\delta', \quad (30)$$

where (29) is from union bound and (30) is due to (28). Picking $\delta = 2\binom{m}{2}\delta'$, we conclude that

$$|\hat{\mathbf{q}}_{a_i b_i} - \mathbf{q}^\star_{a_i b_i}| \leq \sqrt{\frac{\log(4\binom{m}{2}/\delta)}{2n_i}}, \quad \forall i,$$

holds with probability at least $1-\delta$. Inserting it to the result in Proposition 3, we complete the proof.

B Additional Experimental Details

B.1 Zappos dataset

In order to obtain the embedding of the shoes, we modify VGG11 [31] so that the output is a 21-dimensional vector, which is the number of classes in the UT Zappos 50K dataset. After reaching 97% training accuracy, we extract the representation in the penultimate layer of the network, which is 512 dimensional, and use PaCMAP [34] with default parameters to obtain the 2D embedding of the 5 shoes, shown in Figure 13.

Then, we create tasks with the interface shown in Figure 14 on Amazon Mechanical Turk [3] where each worker is asked to answer 15 preference questions. 10 of which are the 10 pairs formed by 5 shoes in Figure 11 and the rest are repeated queries used to measure the consistency of workers. We first asked 100 different workers to complete this task so that we can have answers to all 10 pairwise queries and find regions where workers are located. Therefore, we can form an estimate of $\mathbf{p}^\star_{\mathcal{H}(S_m)}$. In case of an inconsistency, where a worker may fall into not a unique region, but multiple regions, these regions share the probability equally, i.e. if a worker falls into p regions, then each region has a probability of $\frac{1}{p}$.

Then, we gather answers from 1000 unique workers so that each pair of shoes is answered by 10 folds of 100 different people. We use our methods to recover the $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ 10 times, each time with one of the ten folds of 100 different people. Figure 12 shows our recovered $\hat{\mathbf{p}}_{\mathcal{H}(S_m)}$ compares to $\mathbf{p}^\star_{\mathcal{H}(S_m)}$. Figure 15 shows regions and 2D embedding of the dataset. Each color represents one of the 21 types of shoes in the dataset. Each region is labeled with 4 numbers, which correspond to regions ID, median, minimum, and maximum of the probabilities in that region respectively.

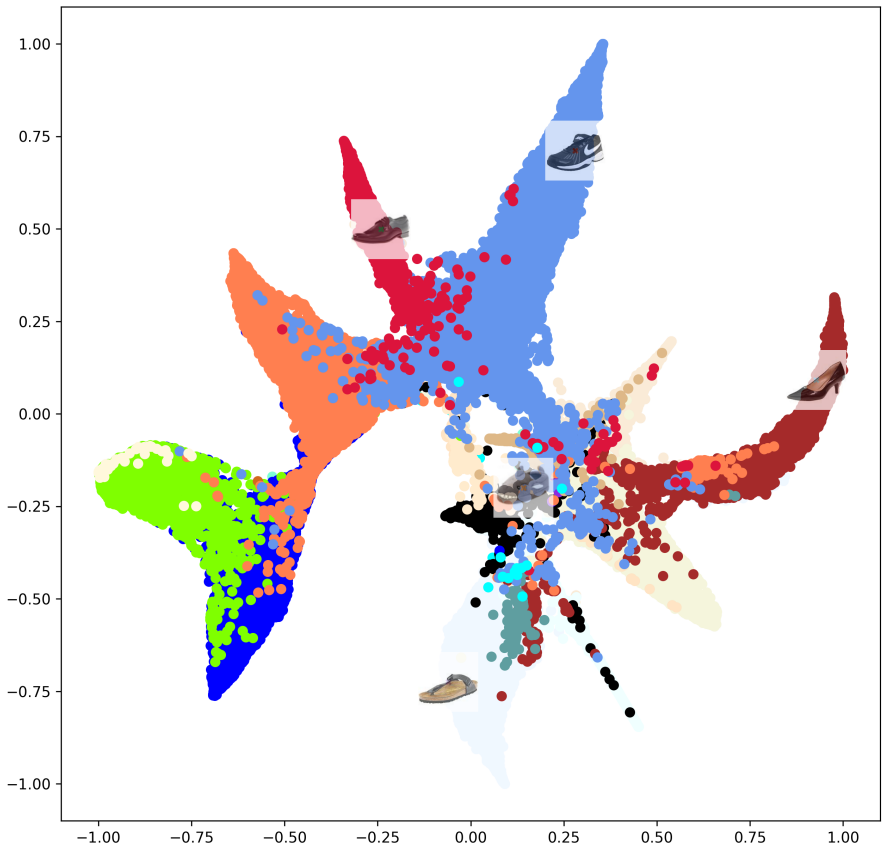

Figure 13: 2D embedding of the 5 shoes obtained using penultimate layer of modified VGG11 and PaCMAP. Each color represents a type of shoes. 5 shoes we used for experiments are also located.

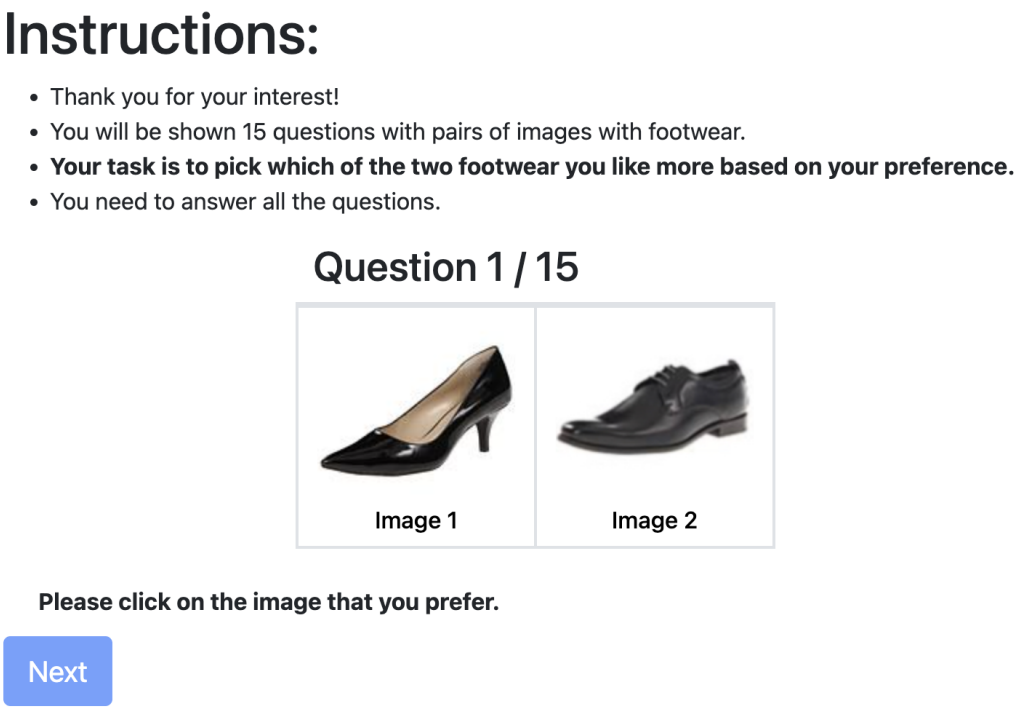

Figure 14: AMT Task interface

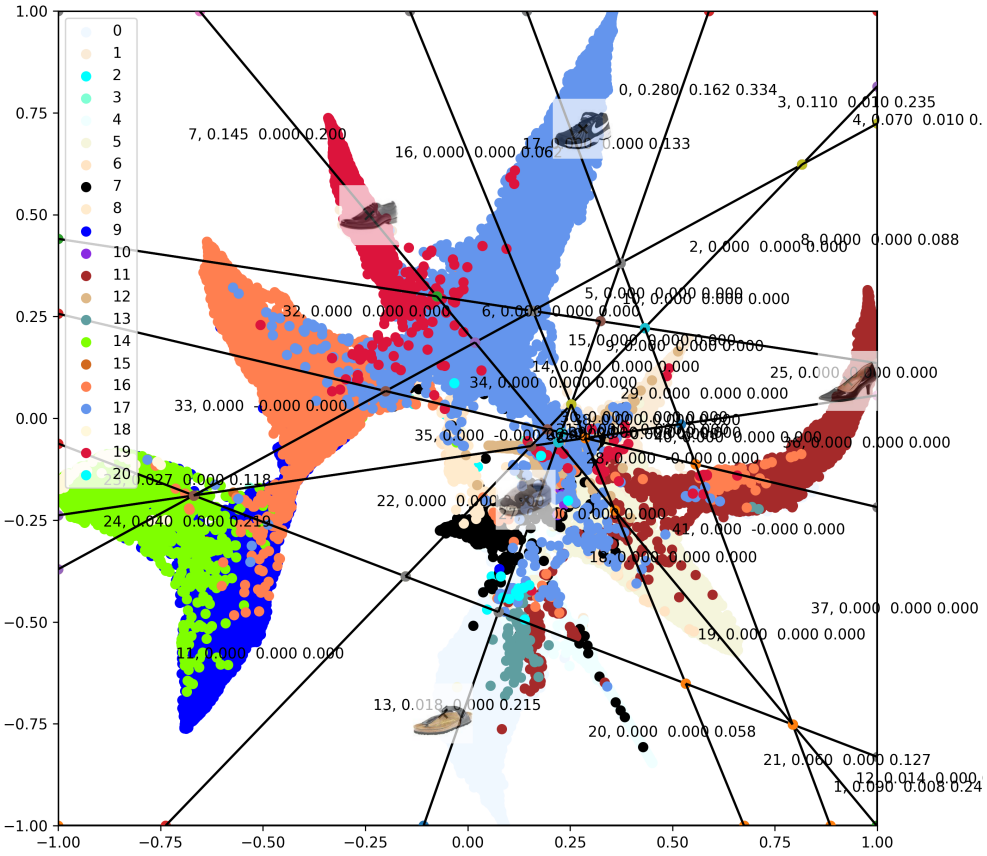

Figure 15: 2D embedding of shoes with restricted regions by hyperplanes