

Learning Preference Distributions From Distance Measurements

Gokcan Tatli, Rob Nowak and Ramya Korlakai Vinayak

Abstract

We introduce the problem of learning a distribution of user preferences over a set of items from noisy responses to distance queries. Rather than aiming to learn the preferences of each user, our goal is only to recover the overall distribution of user preferences. We show that distribution recovery can require just one response from each user. In contrast, learning the preferences of each user would require multiple responses from each user. Thus, learning preference distributions, rather than individual preferences, may be more practical in many applications. The preference distribution problem is formulated on a discrete domain in which items (e.g., products) and users' ideal preference points are located. We study both the noiseless and noisy settings in one dimension and provide sufficient conditions for identifiability of the underlying true distribution as a function of the set of items used for queries. We establish an upper bound on the total variation distance between the true distribution and the distribution learned via constrained least squares optimization problem for both noiseless and noisy settings. While the one-dimensional setting we consider is simple, our simulation results show that our proposed recovery technique extends to multidimensional settings and graph structures.

1 Introduction

We introduce the problem of *learning user preference distribution* from noisy distance queries from a large number of users where each user only answers one query. Consider the setup where each user has their own unknown ideal preference point in \mathbf{u} in domain \mathcal{D} and can answer distance queries of type $\text{Query}(\mathbf{x}, \mathbf{u}) := \text{dist}(\mathbf{x}, \mathbf{u}) + \varepsilon$, where $\mathbf{x} \in \mathcal{D}$ is the item that the user is being queried with, dist is an appropriate distance metric in domain \mathcal{D} and ε is noise. If each user could be queried multiple times, then one could reasonably localize the each user's preference point by triangulating the distance information from queries with items that are sufficiently spread out. However, in many applications, it is not possible to have such a rich number of observations per user. In such scenarios, where the number of observations per user is restricted, e.g., one query per user, it is impossible to localize each user's preference point. Instead, our aim is to address the following fundamental question – *can we learn the distribution of the user preference points over the population of users with only one observation per user?* (see Figure 1). Such a learned preference distribution can then be used in downstream tasks, e.g., testing if the user preference distribution is unimodal or bimodal, identifying top items that are user favourites over the population, learning the preference points for new users using fewer queries using the learned distribution as prior. It is a priori unclear whether learning the preference distribution should be possible with only one query per user, i.e., whether the underlying distribution is identifiable and, if so, under what conditions on the set of query points can we expect this to be identifiable? In particular, how many unique query items would we need and what choices of the set of query items enable learning in this setup?

Our contributions. We introduce the novel problem of learning user preference distribution from noisy distance measurements. We address the fundamental questions related to

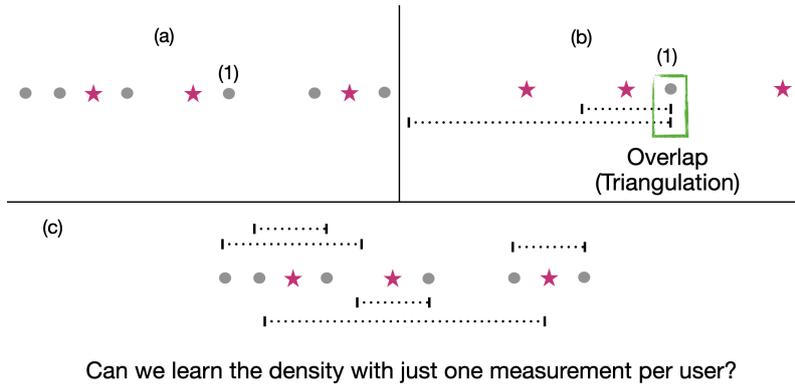


Figure 1: Depiction (a) a set of unknown user locations (gray circles) and items with known locations (pink stars); (b) Triangulation to localize a user, e.g., user denoted by 1, requires two distance measurements per user; (c) Distribution estimation instead of localization.

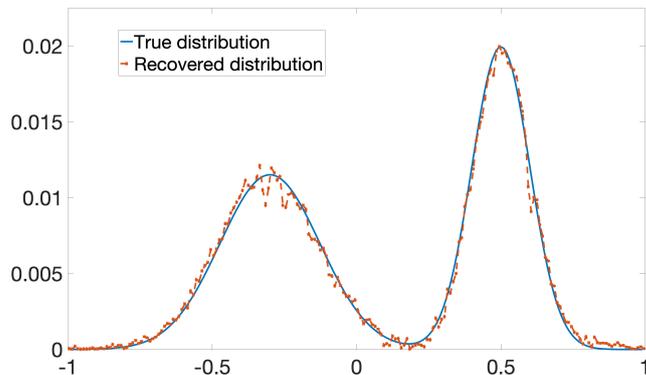


Figure 2: Recovery of underlying user preference distribution via distance measurements.

identifiability and recovery for 1D discrete setting (see Section 2.1 for details). We provide sufficient conditions for identifiability (Theorem 1) and provide recovery guarantees (Theorem 2) for both noiseless and noisy setups (Corollaries 1 and 2). We validate our results via simulations that demonstrate the efficacy of proposed recovery approaches (see Section 4). While the main focus of the paper is on the discrete 1D setting, we discuss extensions to 2D setup (Figure 9) and graph structured settings (Figure 10) in Section 5.

Related works. A closely related problem is that of localization in networks and source localization which has a rich line of literature [1, 10, 9, 15, 2] and more recently using triplet comparison in preference learning [11]. Consider a large number of tiny sensors in a geographical area with a few distance measuring stations spread across the area that can send beacons and receive noisy distance information from the sensors. If there are at least three stations located in generalized positions and each sensor is able to send distance information for each of the stations, one can reasonably triangulate each sensor, and depending on the noise levels, multiple measurements might be required per sensor for better localization. However, with a large number of tiny sensors, it might not be possible to take multiple measurements per sensor due to energy constraints. In scenarios where we are only able to get one measurement per sensor, it is impossible to localize each of them. Instead, in such cases, knowing the density of the sensors in the geographical area might suffice to test

whether it matches with a target distribution. User preference or ideal point learning has been well studied in the preference learning and ranking literature [4, 6, 8, 16, 14, 7, 17, 3] when each user is queried multiple times with comparison to learn their preference point or it is assumed that all users have the same preference point and the queries get distributed. While the model we study is simpler than those considered in preference learning via pairwise comparison queries, we believe the insights we develop by studying the restricted setup build the foundations for understanding more general sensing models.

The rest of the paper is organized as follows. In Section 2, we introduce the problem setup for the noiseless case and provide identifiability and recovery guarantees. In Section 3, we discuss the noisy setting and provide identifiability and recovery guarantees under noise. In Section 4, we present numerical simulations that corroborate our recovery guarantees. We then discuss extensions to 2D and continuous settings in Section 5 and we finally conclude in Section 6.

2 Noiseless case: Discrete 1D setting

2.1 Problem Setup

We begin by considering a discrete 1D setting with user preference points and item locations supported on m points, $[m] := \{1, 2, \dots, m\}$. Let \mathbf{p}^* denote the unknown true probability mass function (pmf) of the user preference points. Let $\mathcal{S} \subseteq [m]$ denote the locations of items that are used in queries, and these locations are known. Let N denote the number of users. Each user with an unknown preference point $u \in [m]$ can be queried for distance from one item $i \in \mathcal{S}$, $\text{Query}(i, u)$. In the noiseless setting, the answer provided is $\text{dist}(i, u) = |i - u|$. In the noisy setting the answers are $\text{dist}(i, u) + \varepsilon = |i - u| + \varepsilon$, where ε denotes the noise. We defer the discussion of noisy setting to Section 3 and focus on the noiseless setting here.

Let \mathbf{q} denote the conditional probability of distances observed given the items used in querying, $q_{ji} = q(\text{dist} = j | \text{item} = i) := \Pr(\text{dist} = j | \text{item} = i)$. The conditional probabilities \mathbf{q} and the true underlying distribution of user preferences \mathbf{p}^* are linearly related via measurement matrix \mathbf{M} whose rows correspond to possible locations of the user for the distance observed given the item queried in the noiseless setting. For example, suppose we query a user with an item located at position 4 and the observed distance is 3. In the noiseless setting, we can infer that the user preference point can be located in either positions 1 or 7. So the corresponding row of \mathbf{M} has 1's at these possible locations and 0's everywhere else. The conditional distribution of observing distance of 3 for a query with item 4 can be written as $q(\text{dist} = 3 | \text{item} = 4) = p_1 + p_7$. We have the following linear system of equations that connect the conditional distributions of distances given items, \mathbf{q} , to the underlying true distribution of user preference locations, \mathbf{p}^* ,

$$\mathbf{q} = \mathbf{M}\mathbf{p}^*. \quad (1)$$

Since the rows of \mathbf{M} correspond to the possible locations of user preferences based on distance observed given item queried, there can at most be two 1's in each row and the rest are 0's. For each position $i \in \mathcal{S}$, the number of corresponding rows in \mathbf{M} is given by the simple function

$$\max\{m - i + 1, i\}. \quad (2)$$

So, the number of rows of the matrix \mathbf{M} is given by $K := \sum_{i=1}^{|\mathcal{S}|} \max\{m - i + 1, i\}$.

Example 1. *As a concrete example, consider $m = 5$ and $\mathcal{S} = \{2, 3\}$. Then, the \mathbf{M}*

corresponding to this setup is as follows,

$$\mathbf{M} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{bmatrix},$$

where first 4 rows correspond to $\mathbf{q}(\cdot | \text{item} = 2)$ and remaining 3 rows correspond to $\mathbf{q}(\cdot | \text{item} = 3)$.

2.2 Identifiability

In this section we describe the conditions for identifiability in the noiseless setting. Recall from (1) that the relationship between the conditional probability of distances observed given the items queried, \mathbf{q} , and the true underlying pmf of user preferences \mathbf{p}^* is a linear system given by, $\mathbf{q} = \mathbf{M}\mathbf{p}^*$, where $\mathbf{M} \in \mathbb{R}^{K \times m}$, $\mathbf{p} \in \mathbb{R}^m$ and $\mathbf{q} \in \mathbb{R}^K$ are defined with respect to the setting in Section 2.1. Clearly, if \mathbf{M} has full column rank, then the underlying distribution \mathbf{p}^* is uniquely determined as follows,

$$\mathbf{p}^* = (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T \mathbf{q}. \quad (3)$$

In contrast, when \mathbf{M} does not have full column rank, the solution to the underlying distribution \mathbf{p} is not unique and therefore it is not identifiable. Therefore, we focus our attention on understanding when \mathbf{M} is guaranteed to be full-rank and hence the model is identifiable. We provide the following sufficient condition for identifiability,

Theorem 1 (Identifiability, Noiseless case). *Any query item set $\mathcal{S} \subseteq [m]$ of size $|\mathcal{S}| = 2$ is sufficient for the model in (1) to be identifiable.*

We note that for the item $i = 1$ (and similarly for $i = m$) the corresponding rows of the matrix \mathbf{M} has only one 1 in each row and thus it generates all the standard basis vectors. In contrast, an item at the mid-point only has $\lceil m/2 + 1 \rceil$ rows corresponding to it in \mathbf{M} . We further note that, the rows of \mathbf{M} corresponding to a single item i are independent. However, it is not a priori clear how many items together lead to a full column rank \mathbf{M} and the above theorem establishes that, surprisingly, any two items are sufficient. For details of the proof, see Appendix A.1.

2.3 Recovery Guarantees

If we had access to the true conditional distribution \mathbf{q} , we could recover \mathbf{p}^* via (3). Since we only have access to observed samples, we have to work with an estimated $\hat{\mathbf{q}}$. We propose using the following constrained least-squares optimization problem to estimate \mathbf{p} from observed samples,

$$\underset{\mathbf{p} \in \mathbb{R}^m}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{M}\mathbf{p} - \hat{\mathbf{q}}\|_2^2 \quad (4a)$$

$$\text{subject to} \quad \mathbf{1}^T \mathbf{p} = 1, \quad (4b)$$

$$p_i \geq 0, \quad i = 1, \dots, m. \quad (4c)$$

When the problem is identifiable, i.e., when \mathbf{M} is full rank, the above convex optimization problem is guaranteed to have a unique solution as the objective function is strongly convex.

Let \mathbf{p}_{sol} denote the solution obtained by solving the optimization problem (4). We show the following recovery guarantee in the noiseless setting under mild assumptions of independence across individuals and with random item from the subset \mathcal{S} being chosen for each query,

Theorem 2. *When the model in (1) is identifiable, with probability at least $1 - \delta$, the total variation distance between the recovered distribution \mathbf{p}_{sol} and the true distribution \mathbf{p}^* is bounded as follows,*

$$TV(\mathbf{p}^*, \mathbf{p}_{\text{sol}}) \leq \frac{\sqrt{mK} \text{cond}(\mathbf{M}, 1)}{2} \max \left\{ \sqrt{\frac{20r_{\max}}{n}}, \sqrt{\frac{25 \log(3/\delta)}{n}} \right\} \quad (5)$$

where $\text{cond}(\mathbf{M}, 1)$ is the condition number of \mathbf{M} with respect to l_1 -norm, n is the total number of users, m is the support size, $r_{\max} := \max_{i \in \mathcal{S}} \{\max\{m - i + 1, i\}\}$ and $K := \sum_{i \in \mathcal{S}} \max\{m - i + 1, i\}$ is the number of rows in \mathbf{M} .

For a given support size m , the solution recovered by the optimization problem in (4) is consistent under total variation distance. The proof details are deferred to Appendix A.2.

3 Noisy case: Discrete 1D setting

In this section, we discuss the noise interference to the setting given in Section 2.1. In the noisy setting the answers to $\text{Query}(i, u)$ are $\text{dist}(i, u) + \varepsilon = |i - u| + \varepsilon$, where ε denotes independent noise. So, in the noisy setting, instead of true conditional distance distributions \mathbf{q} , we have noisy version $\mathbf{q}_{\text{noise}}$. Therefore, we can write the conditional distance distributions for item i in the noisy setting as follows,

$$\mathbf{q}_{\text{noise}}(\cdot | \text{item} = i) = f_i(\varepsilon) * \mathbf{q}(\cdot | \text{item} = i), \quad i \in \mathcal{S}, \quad (6)$$

where $*$ represents convolution operator, $f_i(\cdot)$ represents pmf of the distribution of noise on distances for item at position i and $\mathbf{q}(\cdot | \text{item} = i)$ is the true conditional distribution of distances for item i . We can represent convolution operation using Toeplitz matrices as follows,

$$\mathbf{q}_{\text{noise}}(\cdot | \text{item} = i) = \mathbf{T}_i^\varepsilon \mathbf{q}(\cdot | \text{item} = i), \quad i \in \mathcal{S}, \quad (7)$$

where \mathbf{T}_i^ε is the Toeplitz matrix capturing the noise ε for item at i . Using (7), we can re-write the linear system in (1) for the noisy setting as, $\mathbf{q}_{\text{noise}} = \mathbf{T}^\varepsilon \mathbf{M} \mathbf{p}^*$, where \mathbf{T}^ε is the appropriate concatenation of Toeplitz matrices of each item i . However, convolution results in a different support than support of f_i 's. To be more specific, since the noise interference causes blurring, the items close to the end points need accounting for spill overs. As an example, for a user with preference point at location 1 when queried with item $i = 1$, the noise in distance can only be positive since negative values are not possible. Therefore, $f_i(\cdot)$'s should be determined differently for i 's close to the boundary. Accordingly, we make modifications to the Toeplitz matrices to be able to express noise effect on \mathbf{q} with a matrix multiplication (see Example (2)). Let $\tilde{\mathbf{T}}_i^\varepsilon$ represent the modified versions of \mathbf{T}_i^ε . Then, we can rewrite linear system for the noisy setting as follows,

$$\mathbf{q}_{\text{noise}} = \tilde{\mathbf{T}}^\varepsilon \mathbf{M} \mathbf{p}^*. \quad (8)$$

3.1 Identifiability

The following proposition provides a sufficient condition for identifiability in the noisy setting.

Corollary 1 (Identifiability, Noisy case). *Under the noisy model in (8), \mathbf{p}^* is identifiable for any query item set $\mathcal{S} \subseteq [m]$ of size $|\mathcal{S}| \geq 2$ when the noise matrix \mathbf{T}^ϵ is full rank.*

The proof follows from that of Theorem 1 and full rankness of $\tilde{\mathbf{T}}^\epsilon$ which together makes $\tilde{\mathbf{T}}^\epsilon \mathbf{M}$ invertible.

Example 2. *Let us now take a look at some concrete examples of structures of matrix $\tilde{\mathbf{T}}_i^\epsilon$ by discussing an illustrative example. Consider a case where $m = 5$ and $S = \{1, 2\}$ with the following discrete noise model for items $i = 2, 3, 4$, $f(0) = 1 - 2\epsilon$, $f(-1) = f(1) = \epsilon$. For item $i = 1$, the noise is modified as $f(0) = 1 - 2\epsilon$, $f(1) = 2\epsilon$ and for item $i = 5$ as $f(0) = 1 - 2\epsilon$, $f(-1) = 2\epsilon$, to account for the boundary effect. The corresponding Toeplitz-like matrices are given below.*

$$[\tilde{\mathbf{T}}_1^\epsilon]_{5 \times 5} = \begin{bmatrix} 1 - 2\epsilon & \epsilon & 0 & 0 & 0 \\ 2\epsilon & 1 - 2\epsilon & \epsilon & 0 & 0 \\ 0 & \epsilon & 1 - 2\epsilon & \epsilon & 0 \\ 0 & 0 & \epsilon & 1 - 2\epsilon & 2\epsilon \\ 0 & 0 & 0 & \epsilon & 1 - 2\epsilon \end{bmatrix}.$$

$$[\tilde{\mathbf{T}}_2^\epsilon]_{4 \times 4} = \begin{bmatrix} 1 - 2\epsilon & \epsilon & 0 & 0 \\ 2\epsilon & 1 - 2\epsilon & \epsilon & 0 \\ 0 & \epsilon & 1 - 2\epsilon & 2\epsilon \\ 0 & 0 & \epsilon & 1 - 2\epsilon \end{bmatrix}.$$

Then, modified Toeplitz matrix $\tilde{\mathbf{T}}^\epsilon$ can be written as

$$\tilde{\mathbf{T}}^\epsilon = \begin{bmatrix} \tilde{\mathbf{T}}_1^\epsilon & \mathbf{0}_{5 \times 4} \\ \mathbf{0}_{4 \times 5} & \tilde{\mathbf{T}}_2^\epsilon \end{bmatrix}.$$

Note that rank of $\tilde{\mathbf{T}}^\epsilon$ for this illustrative example is a function of ϵ . Whereas $\tilde{\mathbf{T}}^{\frac{1}{6}}$ satisfies full rank condition, $\tilde{\mathbf{T}}^{\frac{1}{3}}$ does not satisfy.

Remark 1. (Worst case error) *The noise matrices $\tilde{\mathbf{T}}^\epsilon$ are not always full rank. As a simple example for worst case scenario, we consider a case where $m = 5$, $S = 1$ and all distance values are equally possible. Here, $\tilde{\mathbf{T}}^\epsilon$ becomes $\frac{1}{5} \mathbf{1}_{[5 \times 1]} \mathbf{1}_{[1 \times 5]}$, which is clearly a rank one matrix.*

3.2 Recovery Guarantees

We use the following modified optimization problem to estimate the underlying distribution of user preference locations,

$$\underset{\mathbf{p} \in \mathbb{R}^m}{\text{minimize}} \quad \frac{1}{2} \|\tilde{\mathbf{T}}^\epsilon \mathbf{M} \mathbf{p} - \hat{\mathbf{q}}_{\text{noise}}\|_2^2 \quad (9a)$$

$$\text{subject to} \quad \mathbf{1}^T \mathbf{p} = 1, \quad (9b)$$

$$p_i \geq 0, \quad i = 1, \dots, m. \quad (9c)$$

Similar to the noiseless setting, the above optimization problem has unique solution when the model is identifiable, i.e., when $\tilde{\mathbf{T}}_i^\epsilon \mathbf{M}$ is full rank. The following corollary provides an

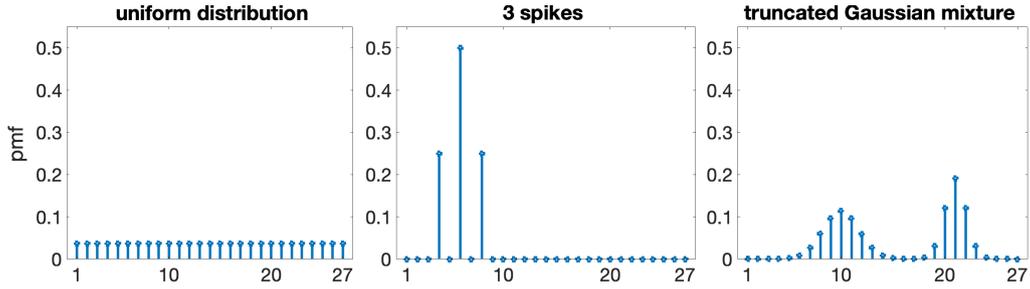


Figure 3: True distributions used for simulations in the discrete setting for both noiseless and noisy cases.

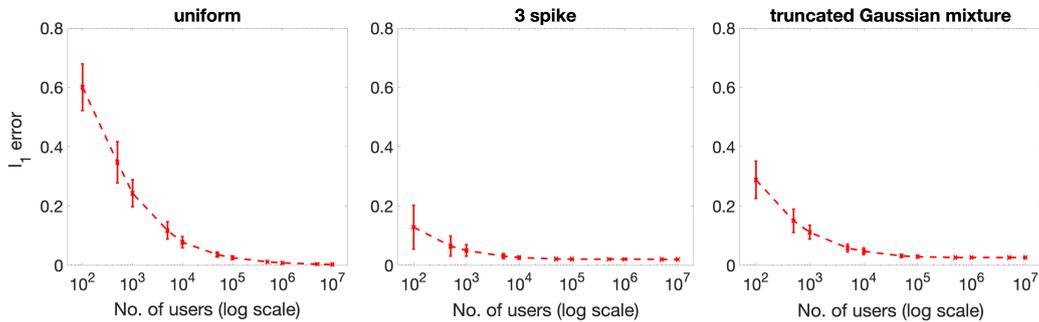


Figure 4: Error in recovered distributions compared to the true distributions for various underlying distributions for noiseless case in 1D discrete setting.

upper bound on the total variation error between the optimal solution \mathbf{p}_{sol} to the above optimization problem and the true distribution \mathbf{p}^* under mild assumptions of independence across individuals and with random item from the subset \mathcal{S} being chosen for each query.

Corollary 2 (Recovery guarantee, Noisy case). *When the model in (8) is identifiable, with probability at least $1 - \delta$, the total variation distance between the recovered distribution \mathbf{p}_{sol} and the true distribution \mathbf{p}^* is bounded as follows,*

$$TV(\mathbf{p}^*, \mathbf{p}_{sol}) \leq \sqrt{mK} \frac{\text{cond}(\mathbf{M}, 1) \text{cond}(\tilde{\mathbf{T}}^\epsilon, 1)}{2} \times \max \left\{ \sqrt{\frac{20r_{\max}}{n}}, \sqrt{\frac{25 \log(3/\delta)}{n}} \right\},$$

where $\text{cond}(\mathbf{M}, 1)$ and $\text{cond}(\tilde{\mathbf{T}}^\epsilon, 1)$ denote the condition numbers of the matrices \mathbf{M} and $\tilde{\mathbf{T}}^\epsilon$ with respect to l_1 -norm, n is the total number of users, m is the support size, $r_{\max} := \max_{i \in \mathcal{S}} \{\max\{m - i + 1, i\}\}$ and $K := \sum_{i \in \mathcal{S}} \max\{m - i + 1, i\}$ is the number of rows in \mathbf{M} .

4 Simulations

In this section we provide numerical simulations that show the recovery of user preference distributions in 1D for both discrete and continuous setups and for both noiseless and noisy cases. We solve the constrained least square optimization problems in (4) and (9) using exponential gradient descent. We initialize the algorithm with uniform distribution except

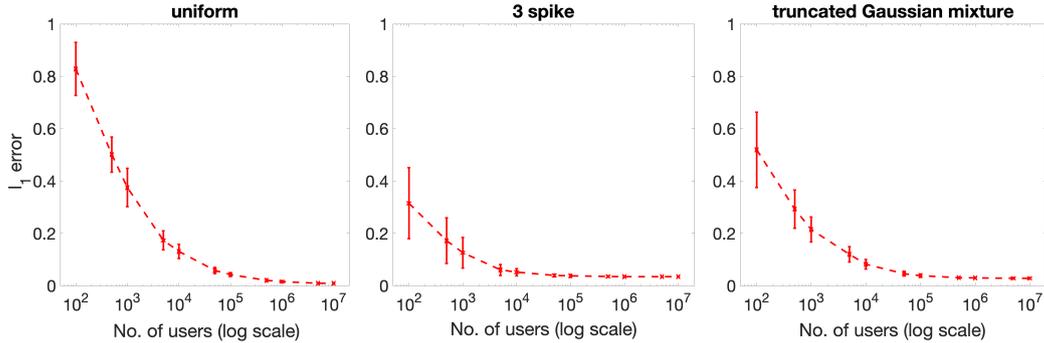


Figure 5: Error in recovered distributions compared to the true distributions for various underlying distributions for noisy case in 1D discrete setting.

when the true underlying distribution is uniform, in which case we use a triangular-like distribution to initialize the algorithm. For the learning rate, we use 0.5 for all scenarios and terminate the iterations when either the maximum iterations of 10^4 or the tolerance value, 10^{-10} , is reached. All the simulations are run on Matlab version R2018b [12].

4.1 Discrete 1D Setting

Let $m = 27$, that is, the support for the user preference points and item locations are $\{1, 2, \dots, 27\}$. The query items are in the set $\mathcal{S} = \{12, 16\}$ which gives an \mathbf{M} with condition number ≈ 5.03 . We ran simulations for three different underlying discrete distributions: (a) uniform, (b) three spikes, and (c) truncated Gaussian for user preference locations as shown in Figure 3. We vary the number of users from 100 to 10 million and repeat each setting 50 times.

4.1.1 Noiseless case

Figure 4 shows the l_1 -error between recovered distributions and the true underlying distributions for each of the cases in the noiseless setup.

4.1.2 Noisy case

We assume following noise model with support size 3, $f(0) = \frac{2}{3}$, $f(-1) = f(1) = \frac{1}{6}$, and construct corresponding Toeplitz-like matrix as described in Section 3.1. Figure 5 shows the l_1 -error (twice the total variation distance) between recovered distributions and the true underlying distributions for each of the cases.

4.2 Continuous 1D Setting

We also provide simulations for continuous underlying distributions with user preference points and item locations supported on interval of $[-1, 1]$ for three different distributions (a) uniform, (b) three spikes, and (c) truncated Gaussian for user preference locations as shown in Figure 6. We choose 10 points uniformly at random from $[-1, 1]$ for the query item set \mathcal{S} . We recover the underlying distribution by discretizing the interval using 0.01 as bin width. We vary the number of users from 100 to 10^7 on the log-scale and repeat each setting 50 times.

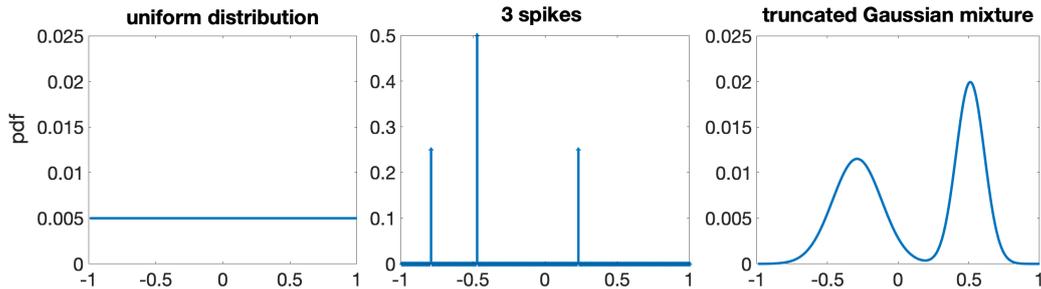


Figure 6: True distributions used for simulations in the continuous setting for both noiseless and noisy cases.

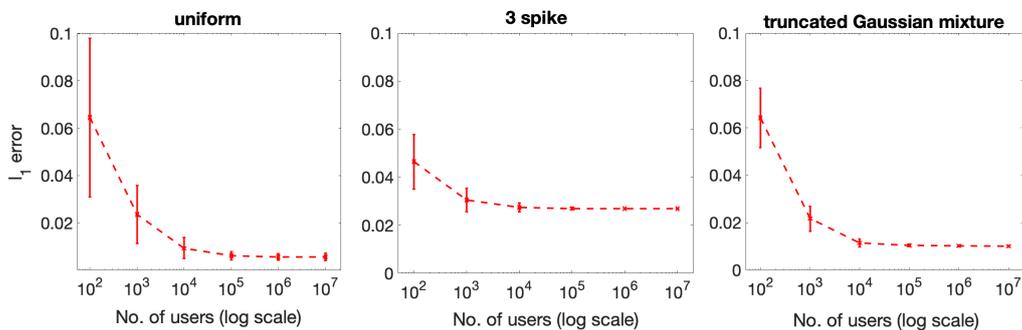


Figure 7: Wasserstein-1 error (earth mover’s distance) between recovered distributions and the true distributions for various underlying distributions for noiseless case in 1D continuous setting.

4.2.1 Noiseless case

Figure 7 shows the Wasserstein-1 error (earth mover’s distance) between the recovered distribution and the true underlying distribution for varying number of users in the noiseless case for different underlying true distributions.

4.2.2 Noisy case

Zero mean Gaussian noise with variance 0.01^2 is used as additive noise to the true distances to obtain noisy distances as answers to distance queries. These noisy answers are used to estimate $\hat{\mathbf{q}}_{\text{noise}}$ by rounding them to nearest bin. We then construct corresponding Toeplitz-like noise matrix and apply the optimization problem in (9) to recover the discretized version of underlying distribution. Figure 8 shows the Wasserstein-1 error (Earth Mover’s Distance) between the recovered distribution and the true underlying distribution for varying number of users for the noisy setting for different underlying true distributions.

5 Extensions to 2D setting and Graphs

While we have focused on the 1D discrete settings so far, our approach of the learning distribution of user preferences over the population can be extended to higher dimensions and to other underlying structures. In this section we discuss how the ideas presented in the paper extends to 2D discrete setting and to graph structured data. We also provide simulations to

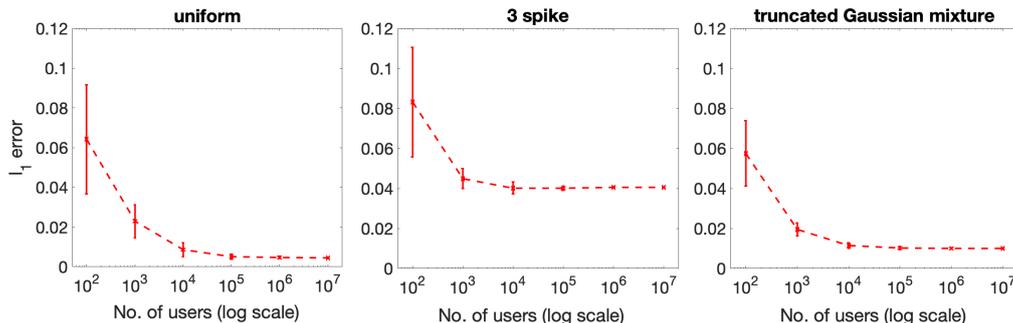


Figure 8: Wasserstein-1 error (earth mover’s distance) between recovered distributions and the true distributions for various underlying distributions for noisy case in 1D continuous setting.

demonstrate that our learning approach indeed extends beyond the 1D discrete setting to these settings. The detailed analysis to obtain theoretical guarantees of identifiability and recovery guarantees for these extended settings are deferred to future work.

5.1 2D Setting

Consider a 2D grid of $[m_1, m_2] := \{1, 2, \dots, m_1\} \times \{1, 2, \dots, m_2\}$. The problem of learning user preference distribution in 2D can be modelled in a similar way, where the vectorized conditional probabilities \mathbf{q} and the true underlying distribution of user preferences \mathbf{p}^* are linearly related via measurement matrix \mathbf{M} , which is constructed in a similar fashion to that of 1D with appropriate adjustments for 2D geometry. The rows of the measurement matrix \mathbf{M} can be obtained by vectorizing the distance matrices. As an example, consider a 4×4 grid. The row of \mathbf{M} corresponding to $q(\text{dist} = 2 | \text{item} = (2, 2))$ can be obtained by vectorizing the following matrix:

$$m^{2,2} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

Figure 9 shows the l_1 -error between the recovered distribution and true distribution for an example 2D setting supported on a 20×20 grid (with 5 repetitions for each case).

We further note that, for continuous settings, one can use the discretization approach as done in the simulation section with appropriate binning as we demonstrated for the 1D setting in the simulations (Section 4.2). However, for higher dimensions, such discretization based approach can become computationally challenging as the grid size scales exponentially with dimension. Another approach in the continuous case is to consider learning approaches based on kernel density estimators. Detailed analysis of identifiability and recovery guarantees for continuous set-up is out of the scope of this paper and is deferred to future work.

5.2 Graph structures

Another setting where our approach extends to is learning the distribution of user preference locations over a graph, where an underlying known graph structure determines the connectivity across the users (sensors) and items (beacons/anchors). Let $A \in \{0, 1\}^{m \times m}$

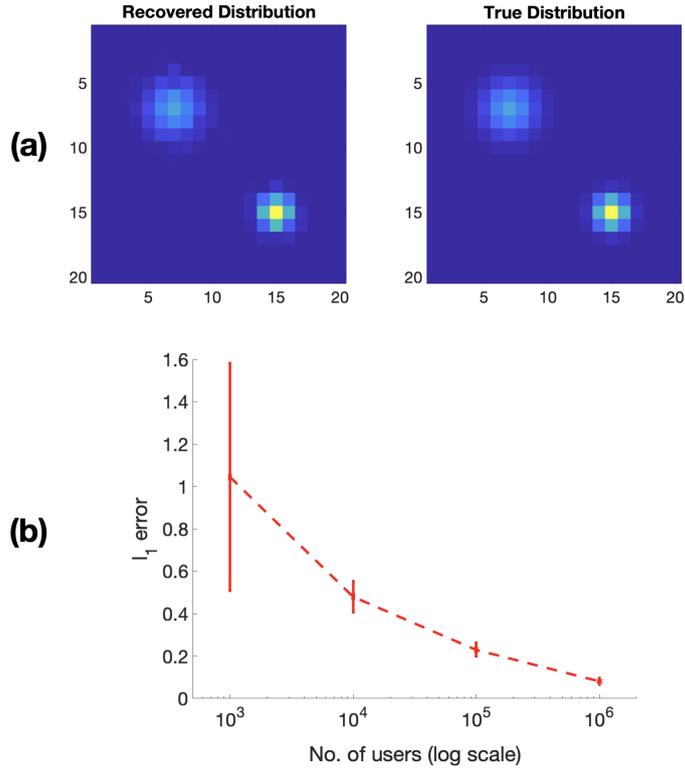


Figure 9: Example of recovery of distribution of user locations for discrete 2D setting. (a) Recovered and true distributions (with 10^6 users). (b) l_1 -error in recovered distribution as function of number of users (noiseless setting).

denote the unweighted and undirected adjacency matrix of a graph on m nodes. The distances between two nodes in this case are measured as the shortest path between the two nodes. We can extend the idea of setting up the linear system of equations relating the conditional distributions of distances given the query nodes and the true distribution of user locations on the nodes by constructing the measurement matrix \mathbf{M} that depends on the underlying graph A . As an example, consider a graph on 6 nodes shown in Figure 10. The rows of \mathbf{M} corresponding to $q(\cdot|\text{item} = \text{node } 3)$ and $q(\cdot|\text{item} = \text{node } 6)$ are given by:

$$\mathbf{M} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Figure 10 shows the l_1 -error between the recovered distribution and true distribution for this example for a varying number of users in the noiseless setting (with 50 repetitions for each case).

The sufficient conditions for identifiability for distribution over a graph depend on the structure of the underlying graph and the locations of query nodes. The 1D discrete setting considered in this paper can be seen as a line graph for which any two nodes in the query set are sufficient for identifiability. However, if the underlying graph is a complete graph,

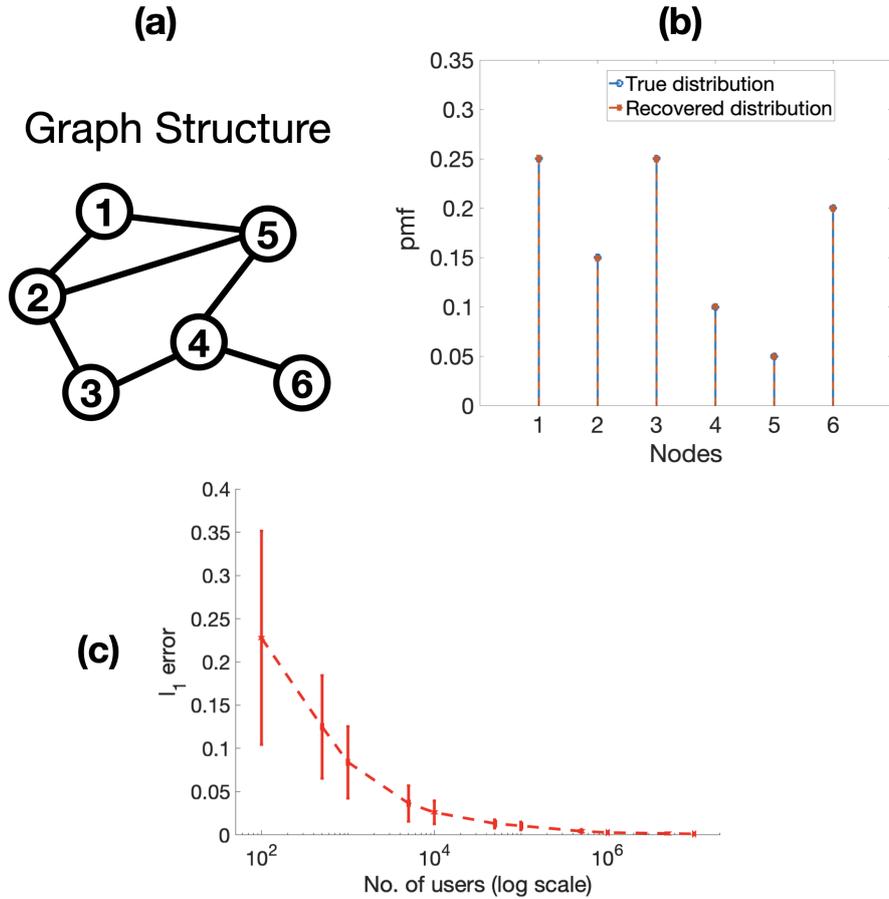


Figure 10: Example of recovery of distribution of user locations on a graph via distance measurements. (a) Known underlying graph structure on 6 nodes. (b) Example of true and recovered distribution (with 10^6 users). (c) l_1 -error in recovered distribution as function of number of users (noiseless setting).

then resolving the mass on individual nodes would not be possible without using almost all the nodes for querying since all nodes are distance 1 from each other.

6 Conclusion and Future Work

In this paper, we introduced the novel problem of learning user preference distribution from noisy answers to distance queries from a large number of users, where each user is limited to answering one query. We study both the noiseless and noisy setups in 1D discrete setting and provide sufficient conditions for identifiability. We also provide recovery guarantees for learning the underlying distribution under total variation distance via constrained least squares optimization problem. While the setting considered is simple, the proposed technique for recovering the distribution of preferences can be easily extended to continuous settings, higher dimensions, and for graph structures. A detailed study of identifiability guarantees and analysis of recovery guarantees for these broader settings is left to future work.

References

- [1] James Aspnes, Tolga Eren, David Kiyoshi Goldenberg, A Stephen Morse, Walter Whiteley, Yang Richard Yang, Brian DO Anderson, and Peter N Belhumeur. A theory of network localization. *IEEE Transactions on Mobile Computing*, 5(12):1663–1678, 2006.
- [2] R Michael Buehrer, Henk Wymeersch, and Reza Monir Vaghefi. Collaborative sensor network localization: Algorithms and practical issues. *Proceedings of the IEEE*, 106(6):1089–1114, 2018.
- [3] 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.
- [4] Gregory S Carpenter and Kent Nakamoto. Consumer preference formation and pioneering advantage. *Journal of Marketing research*, 26(3):285–298, 1989.
- [5] Luc Devroye. The equivalence of weak, strong and complete convergence in l_1 for kernel density estimates. *The Annals of Statistics*, pages 896–904, 1983.
- [6] Johannes Fürnkranz and Eyke Hüllermeier. Preference learning and ranking by pairwise comparison. In *Preference learning*, pages 65–82. Springer, 2010.
- [7] 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.
- [8] Kevin G Jamieson and Robert Nowak. Active ranking using pairwise comparisons. *Advances in neural information processing systems*, 24, 2011.
- [9] Usman A Khan, Soumya Kar, and José MF Moura. Distributed sensor localization in random environments using minimal number of anchor nodes. *IEEE Transactions on Signal Processing*, 57(5):2000–2016, 2009.
- [10] Guoqiang Mao, Barış Fidan, and Brian DO Anderson. Wireless sensor network localization techniques. *Computer networks*, 51(10):2529–2553, 2007.
- [11] Andrew K Massimino and Mark A Davenport. As you like it: Localization via paired comparisons. *J. Mach. Learn. Res.*, 22:186–1, 2021.
- [12] MATLAB. *version R2018b*. The MathWorks Inc., Natick, Massachusetts, 2018.
- [13] Rolf Schneider. *Convex Bodies: The Brunn–Minkowski Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition, 2013. doi: 10.1017/CBO9781139003858.
- [14] 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.
- [15] Jing Wang, Ratan K Ghosh, and Sajal K Das. A survey on sensor localization. *Journal of Control Theory and Applications*, 8(1):2–11, 2010.
- [16] Fabian Wauthier, Michael Jordan, and Nebojsa Jojic. Efficient ranking from pairwise comparisons. In *International Conference on Machine Learning*, pages 109–117. PMLR, 2013.

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

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

Rob Nowak
University of Wisconsin-Madison
Madison, WI, USA
Email: rdnowak@wisc.edu

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

A Proofs

Here we provide proofs for the results presented in Sections 2 and 3.

A.1 Proof of Theorem 1

We suppose that \mathbf{M}^i represents rows of \mathbf{M} regarding to item position i . Then, for any two item positions i and j , without loss of generality, we assume $i < j$ and $i \leq \frac{m}{2}$, where

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}^i \\ \mathbf{M}^j \end{bmatrix}$$

Step I: First we subtract column $i + k$ from its symmetric column $i - k$ with respect to i for all $k = 1, \dots, i - 1$. These operations generate the following subset of standard basis vectors, $\{e_i, e_{i+1} \dots e_m\}$, in the rows of \mathbf{M}^i .

Step II: The target is to generate the remaining standard basis vectors from rows of \mathbf{M}^j to be able to conclude that \mathbf{M} is full rank. We subtract generated standard basis vectors, $\{e_i, e_{i+1} \dots e_m\}$'s, from corresponding rows of \mathbf{M}^j to cancel them out. We call $\overline{\mathbf{M}}$ to this new matrix containing a subset of standard basis vectors as its rows.

Now, we consider rows of $\overline{\mathbf{M}}$ corresponding to same indexed rows with \mathbf{M}^j and call this submatrix $\overline{\mathbf{M}}^j$. If we can prove that the right of anti-diagonal is all zeros for matrix $\overline{\mathbf{M}}^j$, then we can generate $\{e_1, e_2 \dots e_{i-1}\}$ via straightforward row operations and conclude that \mathbf{M} is full column rank. The rest of the proof focuses on verifying that the right of anti-diagonal is all zeros for matrix $\overline{\mathbf{M}}^j$.

Step III: We start by noting that $\overline{\mathbf{M}}^j$ has all zeros in columns i to m , due to row operations in Step II. However, column operations in Step I generate some nonzero values on the latter version $\overline{\mathbf{M}}^j$ of \mathbf{M}^j . If $j \geq 2i - 1$, it is easy to see that those emerging nonzero values are belong to the left of $\overline{\mathbf{M}}^j$'s anti-diagonal, since they result from subtracting column $i + k$ from its symmetric smaller indexed column $i - k$ for all $k = 1, \dots, i - 1$, where column $i + k$ has its only nonzero value 1 on the anti-diagonal of \mathbf{M}^j .

Otherwise, if $j < 2i - 1$, we note that the same rule as above applies for emerging nonzero values originating from subtracting column $i + k$ from its symmetric column $i - k$ for $k = 1, \dots, j - i$. For remaining pairs corresponding to $k = j - i + 1, \dots, i - 1$, we need to dig more to show that target locations, $i - k$'s for $k = j - i + 1, \dots, i - 1$, are on the left of anti-diagonal. We take any pair $(i - k, i + k)$ for $k = j - i + 1, \dots, i - 1$, where column $i + k$ is the source of nonzero entry on the target column $i - k$. Considering their row index, we realize that anti-diagonal element in the same row belongs to column $j - (i + k - j) = 2j - i - k$. Then, we note that the same row's anti-diagonal element has a bigger column index than emerging nonzero value on column $i - k$, i.e.,

$$2j - i - k > i - k, \tag{10}$$

since $j > i$. Therefore, we can say that emerging nonzero value on the column $i - k$ is on the left of anti-diagonal.

A.2 Proof of Theorem 2

Recall that the optimization setting (4) is constrained least square optimization where the constraint is the unit simplex. Multiplication of \mathbf{M} with each element in unit simplex defines a closed convex set, which can be expressed as follows,

$$C_{\mathbf{M}} := \text{conv}(\mathbf{M}e_1, \dots, \mathbf{M}e_m).$$

Clearly, if $\hat{\mathbf{q}} \in C_{\mathbf{M}}$, then $\mathbf{M}\mathbf{p}_{sol} = \hat{\mathbf{q}}$ and $\mathbf{p}_{sol} = \mathbf{M}^\dagger \hat{\mathbf{q}}$, where $\mathbf{M}^\dagger := (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T$. Otherwise, if $\hat{\mathbf{q}} \notin C_{\mathbf{M}}$, we note that $\mathbf{M}\mathbf{p}_{sol} = \text{Proj}_{C_{\mathbf{M}}}(\hat{\mathbf{q}})$, where the projection is under l_2 -distance. Let $\hat{\mathbf{q}}_M := \text{Proj}_{C_{\mathbf{M}}}(\hat{\mathbf{q}})$ for ease of notation.

Noting that $\mathbf{p}^* = \mathbf{M}^\dagger \mathbf{q}$ and $\mathbf{p}_{sol} = \mathbf{M}^\dagger \hat{\mathbf{q}}_M$, we first show that the solution \mathbf{p}_{sol} to (4) is close to the true \mathbf{p}^* under l_2 -norm with the following bound,

$$\begin{aligned} \|\mathbf{p}^* - \mathbf{p}_{sol}\|_2 &= \|\mathbf{M}^\dagger(\mathbf{q} - \hat{\mathbf{q}}_M)\|_2 \leq \|\mathbf{M}^\dagger\|_2 \|\mathbf{q} - \hat{\mathbf{q}}_M\|_2 \\ &\stackrel{(a)}{\leq} \|\mathbf{M}^\dagger\|_2 \|\mathbf{q} - \hat{\mathbf{q}}\|_2 \leq \|\mathbf{M}^\dagger\|_2 \|\mathbf{q} - \hat{\mathbf{q}}\|_1, \end{aligned} \quad (11)$$

where the inequality (a) follows from the fact that projection onto closed convex sets is contracting (Thm. 1.2.2.[13]). Noting the fact that $2 \text{TV}(\mathbf{p}^*, \mathbf{p}_{sol}) = \|\mathbf{p}^* - \mathbf{p}_{sol}\|_1$, and using the inequality between the l_1 - and l_2 -norms, we obtain the following bound from (11),

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

Note that the term $\|\hat{\mathbf{q}} - \mathbf{q}\|_1$ is the sum of l_1 -distances between the empirical and the true conditional distributions of distances for $|\mathcal{S}|$ items in the query set \mathcal{S} . We then note that $\|\mathbf{M}\|_1 = |\mathcal{S}|$ and use the following bound on the l_1 -norm between empirical distribution and the true distribution for discrete distributions on finite support from Lemma 3 in [5],

Lemma 1. (Lemma 3 in [5]) Let $\mathbf{p} \in \Delta^{k-1}$ be discrete distribution with support size k . Let $\hat{\mathbf{p}}$ denote the empirical estimate of this distribution from N i.i.d samples drawn from \mathbf{p} . Then, for all $\varepsilon \geq \sqrt{20k/N}$,

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

This completes the proof of Theorem 2.

A.3 Proof outline of Corollary 2

The matrix \mathbf{M} in the optimization problem (4) is replaced in (9) by $\tilde{\mathbf{T}}^\varepsilon \mathbf{M}$. The proof for recovery guarantee in the noisy setting follows by similar steps to that of the noiseless setting in Appendix A.2. Using the fact that, $(\tilde{\mathbf{T}}^\varepsilon \mathbf{M})^\dagger = \mathbf{M}^\dagger (\tilde{\mathbf{T}}^\varepsilon)^{-1}$, we can obtain the following bound,

$$\|\mathbf{p}^* - \mathbf{p}_{sol}\|_1 \leq \sqrt{mK} \|\mathbf{M}^\dagger (\tilde{\mathbf{T}}^\varepsilon)^{-1}\|_1 \|\mathbf{q} - \hat{\mathbf{q}}\|_1 \quad (12)$$

Then, using sub-multiplicative property of induced matrix norms, we obtain the following,

$$\begin{aligned} \|\mathbf{p}^* - \mathbf{p}_{sol}\|_1 &\leq \sqrt{mK} \|\mathbf{M}^\dagger\|_1 \|(\tilde{\mathbf{T}}^\varepsilon)^{-1}\|_1 \|\mathbf{q} - \hat{\mathbf{q}}\|_1 \\ &= \sqrt{mK} \frac{\text{cond}(\mathbf{M}, 1)}{\|\mathbf{M}\|_1} \frac{\text{cond}(\tilde{\mathbf{T}}^\varepsilon, 1)}{\|\tilde{\mathbf{T}}^\varepsilon\|_1} \|\mathbf{q} - \hat{\mathbf{q}}\|_1 \end{aligned}$$

We note that $\|\tilde{\mathbf{T}}^\varepsilon\|_1 = 1$ and $\|\mathbf{M}\|_1 = |\mathcal{S}|$ by construction. We then apply Lemma 1 for bounding $\|\mathbf{q} - \hat{\mathbf{q}}\|_1$ which leads to the result in Corollary 2.