

Semi-parametric Order-based Generalized Multivariate Regression

Milad Kharratzadeh, Mark Coates

McGill University

Abstract

In this paper, we consider a generalized multivariate regression problem where the responses are monotonic functions of linear transformations of predictors. We propose a semi-parametric algorithm based on the ordering of the responses which is invariant to the functional form of the transformation function. We prove that our algorithm, which maximizes the rank correlation of responses and linear transformations of predictors, is a consistent estimator of the true coefficient matrix. We also identify the rate of convergence and show that the squared estimation error decays with a rate of $o(1/\sqrt{n})$. We then propose a greedy algorithm to maximize the highly non-smooth objective function of our model and examine its performance through extensive simulations. Finally, we compare our algorithm with traditional multivariate regression algorithms over synthetic and real data.

Keywords: semi-parametric regression, generalized multivariate regression, rank correlation

1. Problem Setup

In linear multivariate regression, we have the following model:

$$\mathbf{y}_i^T = \mathbf{x}_i^T \mathbf{B} + \epsilon_i, \quad i = 1, \dots, n, \quad (1)$$

where $\mathbf{y}_i \in \mathbb{R}^{q \times 1}$ is the response vector ($q > 1$), $\mathbf{x}_i \in \mathbb{R}^{p \times 1}$ is the predictor vector, $\mathbf{B} \in \mathbb{R}^{p \times q}$ is the coefficient matrix, and $\epsilon_i \in \mathbb{R}^{q \times 1}$ represents the noise with i.i.d. elements that are independent of \mathbf{x}_i . In this paper, we consider the following extension of this problem:

$$\mathbf{y}_i^T = U_i(\mathbf{x}_i^T \mathbf{B} + \epsilon_i^T), \quad i = 1, \dots, n, \quad (2)$$

where $U_i: \mathbb{R} \rightarrow \mathbb{R}$ is a non-degenerate monotonic function called the *utility* or *link* function. When the input of U_i is a vector or a matrix, it is implied that U_i is applied separately on each individual element to give the output, which is a vector or matrix of the same size as the input. Without loss of generality, we assume that U_i is an increasing function. We propose a semi-parametric, rank-based approach to estimate \mathbf{B} which is invariant with respect to the functional form of U_i functions. Our approach only uses the ordering of the elements of \mathbf{y}_i , which makes it more robust to outliers and heavy-tailed noise compared to traditional regression algorithms. This also makes our approach applicable to cases where the numeric values of \mathbf{y}_i are not available, and only their ordering is known.

We show that it is possible to consistently estimate \mathbf{B} solely based on the ordering of the elements of \mathbf{y}_i . Our approach to estimating \mathbf{B} is based on maximizing Kendall's rank correlation of \mathbf{y}_i^T and $\mathbf{x}_i^T \mathbf{B}$. For notational simplicity, we assume that all the link functions are equal and denote them by

U ; however, all the results presented in this paper hold for the case where there is a separate link function, U_i , for each observation. Let us rewrite (2) in matrix form:

$$\mathbf{Y}_{n \times q} = U(\mathbf{X}_{n \times p} \mathbf{B}_{p \times q} + \mathbf{E}_{n \times q}), \quad (3)$$

where p is the number of predictors, q is the number of responses, and n denotes the number of instances. \mathbf{x}_i^T , \mathbf{y}_i^T , and $\boldsymbol{\epsilon}_i^T$ correspond, respectively, to the i -th rows of \mathbf{X} , \mathbf{Y} , and \mathbf{E} . To find \mathbf{B} , we propose to solve:

$$\hat{\mathbf{B}}_n = \arg \max_{\mathbf{B}} \underbrace{\frac{1}{n \binom{q}{2}} \sum_{i=1}^n \sum_{j=1}^q \sum_{k=1}^q \mathbf{1}(y_{ij} > y_{ik}) \mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j > \mathbf{x}_i^T \mathbf{b}_k)}_{S_n(\mathbf{B})}, \quad (4)$$

where \mathbf{b}_j denotes the j -th column of \mathbf{B} . The intuition behind this formulation is that since U is increasing and the error is i.i.d. and independent of \mathbf{x} , when we have $\mathbf{x}_i^T \mathbf{b}_j > \mathbf{x}_i^T \mathbf{b}_k$, it is more likely to have $y_{ij} > y_{ik}$ than the other way around. The term in the summation is zero for $j = k$. Maximizing $S_n(\mathbf{B})$ is equivalent to maximizing the average rank correlation of \mathbf{y}_i^T and $\mathbf{x}_i^T \mathbf{B}$ since $2S_n(\mathbf{B}) - 1$ corresponds to the average over the n observations of the Kendall rank correlation between \mathbf{y}_i^T and $\mathbf{x}_i^T \mathbf{B}$.

2. Motivating Examples and Related Work

2.1. Learning from non-linear measurements

In many practical settings, the measurements or observations are noisy non-linear functions of a linear transformation of an underlying signal. This could be due to the uncertainties and non-linearities of the measurement device or arise from the experimental design (e.g., censoring or quantization). In the statistics and economics literature, this model is known as the single-index model and it has been studied extensively [1, 2, 3, 4, 5, 6]. The response in the single-index model is univariate and the form of the link function is sometimes assumed known unknown.

In our model, the response is a vector (which leads to a multivariate regression inference problem) and we assume that the functional form of the link function is unknown. Also, as explained in more detail below, our inference approach only uses the ordering of the elements of the response vector. Recently, it has been shown that under certain assumptions (e.g., when the predictors are drawn from a Gaussian distribution), Lasso with non-linear measurements is equivalent to one with linear measurement with an equivalent input noise proportional to the non-linearity of the link function [7]. Thus, it has been suggested to use Lasso in the non-linear case as if the measurements were linear. Here, and under much more general conditions, we show that our algorithm performs better than a simple application of Lasso to the non-linear problem.

2.2. Learning from the ordering of responses

Our approach is particularly of interest in applications (e.g., surveys) where subjects order a set of items based on their preferences. Some examples include physicians ranking different life-sustaining therapies they are likely to remove from critically ill patients after they have already made the decision to gradually withdraw support [8, 9], or people ranking different types of sushi based on their preference [10]. In these scenarios, the underlying model cannot be learned by traditional regression techniques, which require a numerical response. However, our algorithm is directly applicable since it only uses the ordering of the elements of the response vector.

Even in the scenarios where the actual values of responses are available (e.g., numerical ratings), it is often more sensible to focus on the ordering rather than striving to learn based on the assigned numerical values. As discussed in [11], there is often no invariant and objective mapping between true preference and observed ratings among users or survey participants, since “each user uses his/her own mapping based on a subjective and variable criterion in his/her own mind”. Thus, the mappings might be inconsistent among different users. Moreover, the mappings might be inconsistent for a given user across different items; as noted in [12], only trained experts, e.g., wine tasters, are capable of providing consistent mappings for different items. By using the ordering in training and prediction, we minimize the effects of these inconsistencies.

2.3. Collaborative and content-based filtering

Our work is also related to the problem of personalized recommendation systems, but with important differences. Recommendation systems can be divided into three main categories: content-based filtering, collaborative filtering, and hybrid models; see [13, 14, 15] for recent surveys. Content-based filtering employs the domain knowledge of users (e.g. demographic information and user profile data) and items (e.g., genre or actors of a movie) to predict the ratings. Collaborative filtering does not use any user or item information except a partially observed rating matrix, with rows and vectors corresponding to users and items and matrix elements corresponding to ratings. In general, the rating matrix is extremely sparse, since each user, normally, does not experience and rate all items. Hybrid systems combine collaborative and content-based filtering, e.g., by making separate predictions with each filtering approach and averaging the results.

If the regression-based framework described in this paper were used in a recommendation system, it would predict each user’s ordering of a set of items based on a set of features for that user. These features could include demographic information, user profile data, or ratings of a fixed set of items. Contrary to content-based filtering, our approach does not need domain-specific knowledge about the features of items (e.g. relevant features are different for books and movies). This is potentially useful in applications where the items to be ranked are diverse in nature - for example, products on the online Amazon store. Also, as opposed to collaborative filtering, we can incorporate user profile data and provide predictions for new users even if they have provided no prior ratings (i.e., providing predictions only based on features such as demographic data). In Section 8, we provide a preference prediction task where neither collaborative nor content-based filtering is applicable since we want to make predictions for new users without using domain-specific knowledge about the items. Thus, our algorithm is different from the problems of collaborative and content-based filtering. It is important to stress that we introduce and study a general semi-parametric multivariate regression method which can be used in recommendation systems, but this is just one of multiple potential applications.

2.4. Maximum rank correlation estimation

In [16], Han considered a problem similar to (2) but with an important difference. His formulation, called Maximum Rank Correlation (MRC) estimation, was stated for the multiple regression setting, where y_i is real-valued rather than a q -vector, and his goal was to maximize the rank correlation across instances. Therefore, the goal in MRC estimation is to capture the ordering of $y_i, i = 1, \dots, n$ (across instances), whereas in this paper, our goal is to capture the ordering of $y_{ij}, j = 1, \dots, q$ for a fixed i (for a specific instance, across responses). This difference significantly changes the scope of the problem and its theoretical properties. Considering the ordering across responses enables us to model applications where an instance’s ordering (or rating) of a set of

items depends exclusively on its predictors. Also, as we see in the next section, the identifiability and consistency conditions for problem (2) differ significantly from those of the multiple regression problem.

There are extensions of the MRC approach (e.g., [17, 18]), but they all are in the multiple regression domain and only differ in how they define the objective function to solve the same problem. Our work differs from them for the same reasons mentioned above. Unlike exploded and ordered logit models [19, 20], our approach is semi-parametric and invariant to the functional form of U . Our work is also markedly different from the ‘learning to rank’ problem in the information retrieval literature, where the goal is to find relevant documents for a given query [21, 22].

3. Summary of Results

In Section 4, we outline the identifiability conditions on the true model. We show that the true matrix, \mathbf{B}^* , is identifiable up to a scaling and addition of a vector to all its columns. Moreover, we show that the maximizer of the optimization problem in (4) is a consistent estimator of the true matrix. In Section 5, we study the convergence rate and show that our estimator’s squared error decays faster than $o(1/\sqrt{n})$ if the Hessian of $E[S_n(\mathbf{B})]$ is negative definite at \mathbf{B}^* . In Section 6, we provide a greedy algorithm to estimate the maximizer of (4) in polynomial time and show how our algorithm can be extended to provide sparsity among the elements of \mathbf{B} . Finally, in Sections 7 and 8, we provide extensive experimental results using both synthetic and real datasets and show that our algorithm is successful in estimating the underlying true models and provides better prediction results compared to other applicable algorithms.

4. Strong Consistency

In this section, we show that the solution of (4) is strongly consistent under certain conditions. $S_n(\mathbf{B})$ is invariant to the multiplication of all elements of \mathbf{B} by a positive constant; i.e., for $c > 0$, $S_n(\mathbf{B}) = S_n(c\mathbf{B})$. The objective function also does not change if the same vector is added to all of the columns of \mathbf{B} ; i.e., for any $\boldsymbol{\beta} \in \mathbb{R}^{p \times 1}$, $S_n(\mathbf{B}) = S_n(\mathbf{B} + \boldsymbol{\beta}\mathbf{1}_{1 \times q})$, where $\mathbf{1}$ is a vector of all ones. These invariances are expected, since to have a semi-parametric estimate, we target maximizing the rank correlation and ranks are not affected when all the elements are multiplied by a positive constant (c), or are increased/decreased by the same amount ($\mathbf{x}^T\boldsymbol{\beta}$). In other words, since our estimation is semi-parametric in U (and thus, must be invariant to strictly monotonic transformations of observations), \mathbf{B} and $c\mathbf{B} + \boldsymbol{\beta}\mathbf{1}_{1 \times q}$ are equivalent. So, we assume that $\|\mathbf{B}\|_F = 1$ (normalization), and that the last column of \mathbf{B} is all zeros (subtracting the last column from all columns). We perform the optimization in (4) over the set

$$\mathcal{B} \triangleq \{\mathbf{B}_{p \times q} : \|\mathbf{B}\|_F = 1 \text{ and } \mathbf{B}_{i,q} = 0 \text{ for } i = 1, 2, \dots, p\},$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Let us denote the true coefficient matrix by \mathbf{B}^* and without loss of generality assume $\mathbf{B}^* \in \mathcal{B}$; otherwise, we can find $c > 0$ and $\boldsymbol{\beta} \in \mathbb{R}^{p \times 1}$ such that $c\mathbf{B}^* + \boldsymbol{\beta}\mathbf{1}_{1 \times q} \in \mathcal{B}$ which gives the same value for objective function as \mathbf{B}^* . Also, to have a non-degenerate problem, we assume that \mathbf{B}^* does not have rank 1, since in that case there exist two vectors $\mathbf{u} \in \mathbb{R}^{p \times 1}$, $\mathbf{v} \in \mathbb{R}^{q \times 1}$ such that $\mathbf{B}^* = \mathbf{u}\mathbf{v}^T$, and $\mathbf{y}^T = U(\mathbf{x}^T\mathbf{u}\mathbf{v}^T + \boldsymbol{\epsilon})$. Therefore, the ordering of the elements of \mathbf{y} will be either the same as the ordering of the elements of \mathbf{v} (if $\mathbf{x}^T\mathbf{u} > 0$) or the reverse of it (if $\mathbf{x}^T\mathbf{u} < 0$) with perturbations

due to the noise. So, different observed orderings of the elements of \mathbf{y} are caused merely by noise which is not of interest. Finally, we assume that no two columns of \mathbf{B}^* are equal, because in that case the expected values of the corresponding elements of \mathbf{y} will be the same.

Given the model in (3), to prove strong consistency, we need the following three conditions:

- (C1) U is a non-degenerate increasing function and changes value at least at one non-zero point (i.e., U is not a step function changing value only at 0).
- (C2) The elements of \mathbf{E} are i.i.d. random variables.
- (C3) The rows of \mathbf{X} are i.i.d. random vectors of size p , independent of the elements of \mathbf{E} , and have a distribution function F_X such that:
 - (C3.1) the support of F_X is not contained in any proper linear subspace of \mathbb{R}^p , and
 - (C3.2) for all $j \in \{1, 2, \dots, q\}$ the conditional distribution of x_j given the other components has everywhere positive Lebesgue density.
- (C4) \mathbf{B}^* is an interior point of \mathcal{B} . Moreover, \mathbf{B}^* has a rank higher than one and no two columns of it are equal.

As we show later, the second part of condition (C1) is needed in the proof of identifiability. However, for all practical purposes, the step function at 0 can be replaced by an approximate function changing value over $[-\epsilon, \epsilon]$ for some $\epsilon > 0$, for which the theoretical results we present in the following hold. Conditions (C3.1) and (C3.2) are also required for identifiability, and hold in many settings; e.g., when the rows of \mathbf{X} have a multivariate Gaussian distribution. In (C4), $\mathbf{B}^* \in \mathcal{B}$ implies that its last column is all zeros, and because no two columns are equal, we can conclude that every column except the last one has at least one non-zero element. For some known constant $\eta > 0$ which is less than all the absolute values of these non-zero elements, define

$$\mathcal{B}_\eta \triangleq \{\mathbf{B} : \mathbf{B} \in \mathcal{B}; \text{ and } \forall j \in \{1, \dots, p\} \quad \exists i \in \{1, \dots, p\} \text{ s.t. } |\mathbf{B}_{i,j}| \geq \eta\}.$$

Denoting the solution of (4) over the set \mathcal{B}_η by $\widehat{\mathbf{B}}_n$, we prove:

$$\lim_{n \rightarrow \infty} \widehat{\mathbf{B}}_n \rightarrow \mathbf{B}^*. \quad (5)$$

We conduct the proof in three steps. In Lemma 1, we prove the identifiability; in Lemma 2, we prove the convergence of $S_n(\mathbf{B})$ to the expected value of the rank correlation; and finally, we prove the consistency in Theorem 1.

Lemma 1 (Identifiability). *Given (C1)–(C4), \mathbf{B}^* attains the unique maximum of $E[S_n(\mathbf{B})]$ over the set \mathcal{B} .*

Proof. For a given \mathbf{x} such that $\mathbf{x}^T \mathbf{b}_j^* > \mathbf{x}^T \mathbf{b}_k^*$, we have:

$$P_{\epsilon|\mathbf{x}}(y_j > y_k) \geq P_{\epsilon|\mathbf{x}}(y_k > y_j), \quad (6)$$

where $\mathbf{y}^T = U(\mathbf{x}^T \mathbf{B}^* + \boldsymbol{\epsilon}^T)$, and \mathbf{b}_j^* is the j -th column of \mathbf{B}^* . For any matrix $\mathbf{B} \in \mathcal{B}$, we have:

$$\begin{aligned} & E_{\mathbf{x}, \boldsymbol{\epsilon}} \left[\mathbf{1}(y_j > y_k) \mathbf{1}(\mathbf{x}^T \mathbf{b}_j > \mathbf{x}^T \mathbf{b}_k) + \mathbf{1}(y_j < y_k) \mathbf{1}(\mathbf{x}^T \mathbf{b}_j < \mathbf{x}^T \mathbf{b}_k) \right] \\ &= E_{\mathbf{x}} \left[P_{\epsilon|\mathbf{x}}(y_j > y_k) \mathbf{1}(\mathbf{x}^T \mathbf{b}_j > \mathbf{x}^T \mathbf{b}_k) + P_{\epsilon|\mathbf{x}}(y_j < y_k) \mathbf{1}(\mathbf{x}^T \mathbf{b}_j < \mathbf{x}^T \mathbf{b}_k) \right]. \end{aligned} \quad (7)$$

For any j and k , \mathbf{b}_j^* and \mathbf{b}_k^* maximize this expected value, because for any given \mathbf{x} , the larger term between $P_{\epsilon|\mathbf{x}}(y_j > y_k)$ and $P_{\epsilon|\mathbf{x}}(y_j < y_k)$ is chosen in the expected value. Therefore, \mathbf{B}^* maximizes the expected value of each of the terms in (4) and consequently, maximizes $E_{\mathbf{x},\epsilon}[S_n(\mathbf{B})]$. Next, we show that \mathbf{B}^* is the unique maximizer of $E_{\mathbf{x},\epsilon}[S_n(\mathbf{B})]$ over the set \mathcal{B} .

For any $\tilde{\mathbf{B}} \in \mathcal{B}, \tilde{\mathbf{B}} \neq \mathbf{B}^*$, we show that

$$E_{\mathbf{x},\epsilon}[S_n(\mathbf{B}^*)] - E_{\mathbf{x},\epsilon}[S_n(\tilde{\mathbf{B}})] > 0. \quad (8)$$

For any pair of $j, k \in \{1, 2, \dots, q\}, k \neq j$, we can define the following two sets:

$$\begin{aligned} \mathcal{D}_1 &\triangleq \{\mathbf{x} \in \mathcal{R}^p : \mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*) > 0 \text{ and } \mathbf{x}^T(\tilde{\mathbf{b}}_j - \tilde{\mathbf{b}}_k) < 0\}, \\ \mathcal{D}_2 &\triangleq \{\mathbf{x} \in \mathcal{R}^p : \mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*) < 0 \text{ and } \mathbf{x}^T(\tilde{\mathbf{b}}_j - \tilde{\mathbf{b}}_k) > 0\}. \end{aligned}$$

If $(\mathbf{b}_j^* - \mathbf{b}_k^*) \neq c(\tilde{\mathbf{b}}_j - \tilde{\mathbf{b}}_k)$ for some $c > 0$, then both \mathcal{D}_1 and \mathcal{D}_2 have Lebesgue measure greater than 0. We claim that given $\tilde{\mathbf{B}} \in \mathcal{B}, \tilde{\mathbf{B}} \neq \mathbf{B}^*$ and (C4), a pair j and k such that $(\mathbf{b}_j^* - \mathbf{b}_k^*) \neq c(\tilde{\mathbf{b}}_j - \tilde{\mathbf{b}}_k)$ exists. We prove this by contradiction. Assume that such a pair does not exist. Then, since the last columns of both $\tilde{\mathbf{B}}$ and \mathbf{B}^* are zero, setting $k = q$ implies that there exists $c_i > 0$ such that $\mathbf{b}_i^* = c_i \tilde{\mathbf{b}}_i$ for $i = 1, \dots, q-1$. Also, for any $j, k \in \{1, \dots, q-1\}$, there must exist a $c_{j,k} > 0$ such that $(\mathbf{b}_j^* - \mathbf{b}_k^*) = c_{j,k}(\tilde{\mathbf{b}}_j - \tilde{\mathbf{b}}_k)$. Combining these two, we get $(1 - c_{j,k}/c_j)\mathbf{b}_j^* = (1 - c_{j,k}/c_k)\mathbf{b}_k^*$. Since this holds for all j and k , we can conclude that all columns of \mathbf{B}^* are multiples of each other and consequently, \mathbf{B}^* has rank 1, which is a contradiction. Therefore, there exists a pair j and k for which both \mathcal{D}_1 and \mathcal{D}_2 have Lebesgue measure greater than 0.

Now, define the following two sets:

$$\begin{aligned} \mathcal{G}_1 &\triangleq \left\{ \mathbf{x} \in \mathcal{R}^p : \mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*) > 0 \text{ and } E[U(\mathbf{x}^T \mathbf{b}_j^* + \epsilon_j)] > E[U(\mathbf{x}^T \mathbf{b}_k^* + \epsilon_k)] \right\}, \\ \mathcal{G}_2 &\triangleq \left\{ \mathbf{x} \in \mathcal{R}^p : \mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*) < 0 \text{ and } E[U(\mathbf{x}^T \mathbf{b}_j^* + \epsilon_j)] < E[U(\mathbf{x}^T \mathbf{b}_k^* + \epsilon_k)] \right\}. \end{aligned} \quad (9)$$

Next, we show that $\mathcal{H}_1 = \mathcal{D}_1 \cap \mathcal{G}_1$ and/or $\mathcal{H}_2 = \mathcal{D}_2 \cap \mathcal{G}_2$ have Lebesgue measure greater than 0. This is trivial if U is strictly increasing, since $\mathcal{H}_1 = \mathcal{D}_1$ and $\mathcal{H}_2 = \mathcal{D}_2$. We show that \mathcal{H}_1 and/or \mathcal{H}_2 have positive measure in general. Take an $\mathbf{x} \in \mathcal{D}_1$; it is clear that for any $\alpha > 0$, $\alpha\mathbf{x}$ is also in \mathcal{D}_1 , and $-\alpha\mathbf{x}$ is in \mathcal{D}_2 . If we change α from 0^+ to $+\infty$, then $\alpha\mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*)$ changes from 0^+ to $+\infty$ and $-\alpha\mathbf{x}^T(\mathbf{b}_j^* - \mathbf{b}_k^*)$ changes from 0^- to $-\infty$. Since U is non-degenerate and changes value at a non-zero point, there exists a neighborhood \mathcal{A} such that for $\alpha \in \mathcal{A}$, we have $E[U(\alpha\mathbf{x}^T \mathbf{b}_j^* + \epsilon_j)] > E[U(\alpha\mathbf{x}^T \mathbf{b}_k^* + \epsilon_k)]$ and/or $E[U(-\alpha\mathbf{x}^T \mathbf{b}_j^* + \epsilon_j)] < E[U(-\alpha\mathbf{x}^T \mathbf{b}_k^* + \epsilon_k)]$. Thus \mathcal{H}_1 and/or \mathcal{H}_2 have Lebesgue measure greater than 0.

Without loss of generality, assume that \mathcal{H}_1 defined for $j = j'$ and $k = k'$ has Lebesgue measure greater than 0. In the following we show that $E[S_n(\mathbf{B}^*)] > E[S_n(\tilde{\mathbf{B}})]$, which proves the lemma. We

have:

$$\begin{aligned}
& n \binom{q}{2} \cdot E_{\mathbf{x}, \epsilon} [S_n(\mathbf{B}^*) - S_n(\tilde{\mathbf{B}})] \\
&= \sum_{i,j,k} E_{\mathbf{x}, \epsilon} \left[\mathbf{1}(y_{ij} > y_{ik}) \left(\mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j^* > \mathbf{x}_i^T \mathbf{b}_k^*) - \mathbf{1}(\mathbf{x}_i^T \tilde{\mathbf{b}}_j > \mathbf{x}_i^T \tilde{\mathbf{b}}_k) \right) \right. \\
&\quad \left. + \mathbf{1}(y_{ij} < y_{ik}) \left(\mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j^* < \mathbf{x}_i^T \mathbf{b}_k^*) - \mathbf{1}(\mathbf{x}_i^T \tilde{\mathbf{b}}_j < \mathbf{x}_i^T \tilde{\mathbf{b}}_k) \right) \right] \tag{10}
\end{aligned}$$

$$\geq \sum_i P(\mathbf{x}_i \in \mathcal{H}_1) E_{\mathbf{x}, \epsilon} \left[\mathbf{1}(y_{ij'} > y_{ik'}) - \mathbf{1}(y_{ij'} < y_{ik'}) \middle| \mathbf{x}_i \in \mathcal{H}_1 \right] \tag{11}$$

$$= \sum_i P(\mathbf{x}_i \in \mathcal{H}_1) E_{\mathbf{x}} \left[P_{\epsilon|\mathbf{x}_i}(y_{ij'} > y_{ik'}) - P_{\epsilon|\mathbf{x}_i}(y_{ij'} < y_{ik'}) \middle| \mathbf{x}_i \in \mathcal{H}_1 \right] > 0 \tag{12}$$

The summation in (10) is over all distinct pairs, j and k . The inequality in (11) is based on the argument made for (7); for any given i , j , and k , the expectation in (10) is greater than or equal to zero (because \mathbf{B}^* maximizes (7) for every \mathbf{x} and any pair j and k). Therefore, by removing all pairs of j and k such that $j \neq j'$ and $k \neq k'$, and conditioning only on $\mathbf{x}_i \in \mathcal{H}_1$, we get a smaller term. The first term in (12) is strictly greater than zero because \mathcal{H}_1 has positive measure, and hence $P(\mathbf{x}_i \in \mathcal{H}_1) > 0$. The second term is positive, because for any $\mathbf{x}_i \in \mathcal{G}_1$, the term inside $E_{\mathbf{x}}$ is strictly greater than zero. \square

Lemma 2 (Convergence). *Given (C1)–(C4), and denoting*

$$h_i(\mathbf{B}) = \frac{1}{\binom{q}{2}} \sum_{j=1}^q \sum_{k=1}^q \mathbf{1}(y_{ij} > y_{ik}) \mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j > \mathbf{x}_i^T \mathbf{b}_k),$$

we have:

$$S_n(\mathbf{B}) \xrightarrow{a.s.} E[h_i(\mathbf{B})]$$

Proof. Define

$$\begin{aligned}
h(\mathbf{B}) &\triangleq E[h_i(\mathbf{B})] \quad ; \quad D_\delta(\mathbf{B}) \triangleq \{\boldsymbol{\beta} \in \mathbb{R}^{p \times q} : \boldsymbol{\beta} \in \mathcal{B}, \|\boldsymbol{\beta} - \mathbf{B}\|_F < \delta\} \\
\bar{g}_i(\mathbf{B}, \delta) &\triangleq \sup_{\boldsymbol{\beta} \in D_\delta(\mathbf{B})} \{h_i(\boldsymbol{\beta}) - h(\boldsymbol{\beta})\} \quad ; \quad \underline{g}_i(\mathbf{B}, \delta) \triangleq \inf_{\boldsymbol{\beta} \in D_\delta(\mathbf{B})} \{h_i(\boldsymbol{\beta}) - h(\boldsymbol{\beta})\} \\
\bar{g}(\mathbf{B}, \delta) &\triangleq E[\bar{g}_i(\mathbf{B}, \delta)] \quad ; \quad \underline{g}(\mathbf{B}, \delta) \triangleq E[\underline{g}_i(\mathbf{B}, \delta)]
\end{aligned}$$

It is clear that $h_i(\mathbf{B})$ is bounded and moreover, given (C3), $h_i(\mathbf{B})$, and consequently $h(\mathbf{B})$, are continuous in $\mathbf{B} \in \mathcal{B}_\eta$ almost surely. Therefore, since \mathcal{B}_η is compact and separable, for any $\mathbf{B} \in \mathcal{B}_\eta$, there exists a sequence $\{\mathbf{B}_t\}$ in a countable dense subset of \mathcal{B}_η such that:

$$\lim_{t \rightarrow \infty} h_i(\mathbf{B}_t) = h_i(\mathbf{B}) \quad ; \quad \lim_{t \rightarrow \infty} h(\mathbf{B}_t) = h(\mathbf{B}).$$

Thus, almost surely:

$$\lim_{\delta \rightarrow 0} \bar{g}_i(\mathbf{B}, \delta) = h_i(\mathbf{B}) - h(\mathbf{B}) \quad ; \quad \lim_{\delta \rightarrow 0} \underline{g}_i(\mathbf{B}, \delta) = h_i(\mathbf{B}) - h(\mathbf{B}).$$

Taking the expected value of both sides of these limits, we get for all $\mathbf{B} \in \mathcal{B}_\eta$:

$$\lim_{\delta \rightarrow 0} \bar{g}(\mathbf{B}, \delta) = 0; \quad \lim_{\delta \rightarrow 0^-} \underline{g}(\mathbf{B}, \delta) = 0. \quad (13)$$

We use the Borel–Cantelli lemma [23] to prove the lemma.

$$\begin{aligned} & \sum_{n=1}^{\infty} P \left(\max_{\mathbf{B} \in \mathcal{B}_\eta} |S_n(\mathbf{B}) - h(\mathbf{B})| > \epsilon \right) \\ &= \sum_{n=1}^{\infty} P \left(\max_{\mathbf{B} \in \mathcal{B}_\eta} \left| \frac{1}{n} \sum_{i=1}^n (h_i(\mathbf{B}) - h(\mathbf{B})) \right| > \epsilon \right) \end{aligned} \quad (14)$$

$$\leq \sum_{n=1}^{\infty} P \left(\left| \frac{1}{n} \sum_{i=1}^n \sup_{\mathbf{B} \in \mathcal{B}_\eta} (h_i(\mathbf{B}) - h(\mathbf{B})) \right| > \epsilon \right) + \sum_{n=1}^{\infty} P \left(\left| \frac{1}{n} \sum_{i=1}^n \inf_{\mathbf{B} \in \mathcal{B}_\eta} (h_i(\mathbf{B}) - h(\mathbf{B})) \right| > \epsilon \right). \quad (15)$$

\mathcal{B}_η is a compact set and also $\mathcal{B}_\eta \subset \bigcup_{\beta \in \mathcal{B}_\eta} D_\delta(\beta)$ for any small δ . However, for a compact set, every open cover has a finite subcover. Thus, for a finite L , there exist $\beta_1, \beta_2, \dots, \beta_L$, where $\beta_l \in \mathcal{B}_\eta$, and:

$$\mathcal{B}_\eta \subset \bigcup_{l=1}^L D_{\delta_l}(\beta_l) \quad (16)$$

Also, assume that for $l = 1, \dots, L$, we have $|\bar{g}(\beta_l, \delta_l)| < \epsilon/2$ and $|\underline{g}(\beta_l, \delta_l)| < \epsilon/2$. This is possible to achieve because of (13). Now, from (14)–(15), we have:

$$\sum_{n=1}^{\infty} P \left(\max_{\mathbf{B} \in \mathcal{B}_\eta} |S_n(\mathbf{B}) - h(\mathbf{B})| > \epsilon \right) \quad (17)$$

$$\begin{aligned} & \leq \sum_{n=1}^{\infty} \sum_{l=1}^L P \left(\left| \frac{1}{n} \sum_{i=1}^n \bar{g}_i(\beta_l, \delta_l) \right| > \epsilon \right) + \sum_{n=1}^{\infty} \sum_{l=1}^L P \left(\left| \frac{1}{n} \sum_{i=1}^n \underline{g}_i(\beta_l, \delta_l) \right| > \epsilon \right) \\ & \leq \sum_{n=1}^{\infty} \sum_{l=1}^L P \left(\left| \frac{1}{n} \sum_{i=1}^n \bar{g}_i(\beta_l, \delta_l) - \bar{g}(\beta_l, \delta_l) \right| > \epsilon/2 \right) \\ & \quad + \sum_{n=1}^{\infty} \sum_{l=1}^L P \left(\left| \frac{1}{n} \sum_{i=1}^n \underline{g}_i(\beta_l, \delta_l) - \underline{g}(\beta_l, \delta_l) \right| > \epsilon/2 \right). \end{aligned} \quad (18)$$

For each l , we can invoke the Small Law of Large Numbers for U -statistics [24, Chapter 5] to get:

$$\frac{1}{n} \sum_{i=1}^n \bar{g}_i(\beta_l, \delta_l) \xrightarrow{a.s.} \bar{g}(\beta_l, \delta_l) \quad ; \quad \frac{1}{n} \sum_{i=1}^n \underline{g}_i(\beta_l, \delta_l) \xrightarrow{a.s.} \underline{g}(\beta_l, \delta_l) \quad (19)$$

In the following, we use the Borel–Cantelli lemma [23] twice to get the desired result. Reordering

the summands, we can re-write the first term in (18) as follows:

$$\sum_{l=1}^L \underbrace{\sum_{n=1}^{\infty} P \left(\left| \frac{1}{n} \sum_{i=1}^n \bar{g}_i(\beta_l, \delta_l) - \bar{g}(\beta_l, \delta_l) \right| > \epsilon/2 \right)}_{(\star)}.$$

According to the Borel-Cantelli lemma, the almost sure convergence in (19) implies that each of terms marked by (\star) is finite, and thus their sum is finite (since L is finite). The same argument can be made for the second term in (18). Therefore, $\sum_{n=1}^{\infty} P(\max_{\mathbf{B} \in \mathcal{B}_\eta} |S_n(\mathbf{B}) - h(\mathbf{B})| > \epsilon)$ is less than something finite, and hence, is finite. Applying the Borel-Cantelli result proves the lemma. \square

Theorem 1. *Given (C1)–(C4), the solution of (4) over the set \mathcal{B}_η , $\hat{\mathbf{B}}_n$, is strongly consistent; i.e.,*

$$\hat{\mathbf{B}}_n \rightarrow \mathbf{B}^* \quad \text{almost surely.}$$

Proof. Given the results of Lemmas 1 and 2, we are now ready to prove the consistency of the solution of (4) over \mathcal{B}_η . We do so by showing that any set, $\mathcal{B}_0 \subset \mathbb{R}^{p \times q}$, that contains \mathbf{B}^* , also contains $\hat{\mathbf{B}}_n$ as $n \rightarrow \infty$.

Define $\mathcal{B}_1 \triangleq \mathcal{B}_\eta - (\mathcal{B}_0 \cap \mathcal{B}_\eta)$. \mathcal{B}_1 is compact and there exists $\zeta = h(\mathbf{B}^*) - \max_{\mathbf{B} \in \mathcal{B}_1} h(\mathbf{B})$ which is always greater than 0 (because \mathbf{B}^* attains the unique maximum and $\mathbf{B}^* \notin \mathcal{B}_1$). From the result of Lemma 2, we know that for any ζ , there is an N , such that for $n > N$, $|S_n(\mathbf{B}) - h(\mathbf{B})| < \zeta/2$ for all $\mathbf{B} \in \mathcal{B}_\eta$ with probability 1. This implies that $\hat{\mathbf{B}}_n$ cannot be in \mathcal{B}_1 ; because otherwise, we get $h(\mathbf{B}^*) - S_n(\hat{\mathbf{B}}_n) > \zeta/2$ which is in contradiction with almost sure convergence. Thus, $\hat{\mathbf{B}}_n \in \mathcal{B}_0$ with probability 1, and since this is true for any \mathcal{B}_0 , we have $\hat{\mathbf{B}}_n \rightarrow \mathbf{B}^*$ almost surely. \square

5. Rate of Convergence

For ease of notation, let $\boldsymbol{\theta} \in \mathbb{R}^{p(q-1)}$ be the vectorization of the matrix $\mathbf{B} \in \mathcal{B}$, except the last column which is assumed to be all zero. Thus,

$$\boldsymbol{\theta} \triangleq (B_{1,1}, B_{2,1}, \dots, B_{p,1}, \dots, B_{1,q-1}, \dots, B_{p,q-1}). \quad (20)$$

For $\mathbf{B} \in \mathcal{B}$, the corresponding $\boldsymbol{\theta}$ is in Θ , the set of d -dimensional vectors with norm 1. So, we can denote \mathbf{B} and its columns as functions of $\boldsymbol{\theta}$, and write:

$$h(\mathbf{z}, \boldsymbol{\theta}) = \sum_{j=1}^q \sum_{k=1}^q \mathbf{1}(y_j > y_k) \mathbf{1}(\mathbf{x}^T \mathbf{b}_j(\boldsymbol{\theta}) > \mathbf{x}^T \mathbf{b}_k(\boldsymbol{\theta})), \quad (21)$$

where $\mathbf{z} = (\mathbf{y}, \mathbf{x}) \in \mathbb{R}^{p+q}$ is the joint vector of predictors and responses for an instance, and \mathbf{b}_j denotes the j 'th column of \mathbf{B} . Let $\hat{\boldsymbol{\theta}}_n$ correspond to $\hat{\mathbf{B}}_n$ and $\boldsymbol{\theta}_0$ correspond to \mathbf{B}^* . Based on (C4) we know that $\boldsymbol{\theta}_0$ is an interior point of Θ . In the previous section, we showed that $\hat{\boldsymbol{\theta}}_n \xrightarrow{a.s.} \boldsymbol{\theta}_0$. In this section, we study the rate of convergence and show that $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\|^2 \leq o_p(1/\sqrt{n})$, where $\|\cdot\|$ is the Euclidean norm.

The following Lemma plays a critical role in establishing the rate of convergence. Its proof, using results from [25, 26, 27], is included in the Appendix.

Lemma 3. For $\boldsymbol{\theta}$ in an $o_p(1)$ neighborhood of $\boldsymbol{\theta}_0$, and $S(\boldsymbol{\theta}) \triangleq E_{\mathbf{z}}[h(\mathbf{z}, \boldsymbol{\theta})]$, we have:

$$S_n(\boldsymbol{\theta}) = S(\boldsymbol{\theta}) + S_n(\boldsymbol{\theta}_0) - S(\boldsymbol{\theta}_0) + o_p(1/\sqrt{n}). \quad (22)$$

Proof. See the Appendix. \square

For the next Theorem we require that there exists an $o_p(1)$ neighborhood \mathcal{A} of $\boldsymbol{\theta}_0$ and a constant $\kappa > 0$ for which $S(\boldsymbol{\theta}) - S(\boldsymbol{\theta}_0) \leq -\kappa\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2$ for all $\boldsymbol{\theta} \in \mathcal{A}$. Assume ∇ and ∇_2 denote the first and second order derivatives with respect to $\boldsymbol{\theta}$. The existence of \mathcal{A} is guaranteed since $\boldsymbol{\theta}_0$ is an interior point of Θ . Also, $\kappa > 0$ exists if $\nabla_2 S(\boldsymbol{\theta}_0)$ is negative definite, because $\nabla S(\boldsymbol{\theta}_0) = \mathbf{0}$ (since S is maximized at $\boldsymbol{\theta}_0$) and in $o_p(1)$ neighborhood of $\boldsymbol{\theta}_0$, Taylor expansion gives us:

$$S(\boldsymbol{\theta}) - S(\boldsymbol{\theta}_0) = \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}_0)^T \nabla_2 S(\boldsymbol{\theta}_0)(\boldsymbol{\theta} - \boldsymbol{\theta}_0) + o_p(\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2). \quad (23)$$

Since $\nabla_2 S(\boldsymbol{\theta}_0)$ is negative definite, there exists a positive α such that $\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}_0)^T \nabla_2 S(\boldsymbol{\theta}_0)(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \leq -\alpha\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2$. In an $o_p(1)$ neighborhood of $\boldsymbol{\theta}_0$, setting $\kappa = (1 - \epsilon)\alpha$ for any small positive ϵ gives $S(\boldsymbol{\theta}) - S(\boldsymbol{\theta}_0) \leq -\alpha\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2 + o_p(\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2) \leq -\kappa\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2$. In sum, a sufficient condition for the existence of $\kappa > 0$ is the following:

(C5) The matrix $\nabla_2 S(\boldsymbol{\theta}_0)$ is negative definite.

Theorem 2. Assume that there exists an $o_p(1)$ neighborhood \mathcal{A} of $\boldsymbol{\theta}_0$ and a constant $\kappa > 0$ for which $S(\boldsymbol{\theta}) - S(\boldsymbol{\theta}_0) \leq -\kappa\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|^2$ for all $\boldsymbol{\theta} \in \mathcal{A}$. Then, the squared estimation error decays with a rate faster than $1/\sqrt{n}$: $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\|^2 \leq o_p(1/\sqrt{n})$.

Proof. By definition of $\hat{\boldsymbol{\theta}}_n$ we have:

$$0 \leq S_n(\hat{\boldsymbol{\theta}}_n) - S_n(\boldsymbol{\theta}_0). \quad (24)$$

Rewriting this inequality using the result of Lemma 3 gives:

$$0 \leq S(\hat{\boldsymbol{\theta}}_n) - S(\boldsymbol{\theta}_0) + o_p(1/\sqrt{n}) \leq -\kappa\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\|^2 + o_p(1/\sqrt{n}). \quad (25)$$

which gives us $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\|^2 \leq o_p(1/\sqrt{n})$. \square

Corollary 1. Given **(C1)**—**(C5)**, we have $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\|^2 \leq o_p(1/\sqrt{n})$.

6. Optimization Algorithm

In the previous section, we showed that solving (4) provides a consistent estimate of \mathbf{B}^* . However, the objective function is very non-smooth (the sum of many step functions) and finding $\hat{\mathbf{B}}_n$ can be challenging. In this section, we propose a fast, greedy algorithm to solve (4).

First, consider the following maximization problem:

$$\hat{x} = \arg \max_{x \in \mathbb{R}} \sum_{t=1}^T \mathbf{1}(u_t + v_t x > 0), \quad (26)$$

where $\{u_t\}_{t=1}^T$ and $\{v_t\}_{t=1}^T$ are given real numbers. This objective function is a piece-wise constant function changing values at $\{-u_t/v_t\}_{t=1}^T$ (one step function changes value at each of these points).

Therefore, this function takes $O(T)$ different values. Computing each of these values requires $O(T)$ operations, and thus, finding \hat{x} by computing all the possible function values requires $O(T^2)$ operations.

We propose an $O(T \log T)$ algorithm to find \hat{x} . The algorithm works as follows. First, we sort the sequence $\{-u_t/v_t\}_{t=1}^T$, in $O(T \log T)$. Then, we start x from a value less than the first sorted point, i.e., $\min\{-u_t/v_t; t = 1, \dots, T\}$, and at each step, move forward to the next smallest point. At each step, we cumulatively add or subtract 1 depending on the sign of v_t (i.e., depending on whether one of the step functions went from 0 to 1 or vice-versa) and keep track of the largest cumulative sum seen so far, and the value of x for which that maximum happened. After going through all T points, the largest observed cumulative sum is equal to the maximum value of the objective function, and the corresponding value of x is \hat{x} . Since the objective function is piece-wise constant, its maximum is attained over an interval; we set \hat{x} to the center of this interval. Therefore, we can solve (4) in $O(T \log T) + O(T)$ which is equivalent to $O(T \log T)$. The details of the algorithm are summarized below under Algorithm 1. We use this algorithm repeatedly to solve (4).

Algorithm 1 — for solving (26)

Inputs: $u_t, v_t, t = 1, 2, \dots, T$; **Output:** \hat{x}

Compute: $r_t = -u_t/v_t, t = 1, 2, \dots, T$

Sort $r_t : r_{(1)} \leq r_{(2)} \leq \dots \leq r_{(T)}$

▷ **Notation:** $v_{(i)}$: corresponding v of $r_{(i)}$

for $i = 1, 2, \dots, T + 1$ **do**

if $i = 1$ **then**

$\hat{x} = r_{(1)} - 1$

$m \leftarrow \sum_{t=1}^T \mathbf{1}(u_t + v_t \hat{x} > 0)$

$s \leftarrow \sum_{t=1}^T \mathbf{1}(u_t + v_t \hat{x} > 0)$

else if $i = T + 1$ **then**

$s \leftarrow s + \text{sign}(v_{(i)})$

if $s > m$ **then**

$\hat{x} = r_{(T)} + 1$

$m \leftarrow s$

end if

else

$s \leftarrow s + \text{sign}(v_{(i)})$

if $s > m$ **then**

$\hat{x} = (r_{(i-1)} + r_{(i)})/2$

$m \leftarrow s$

end if

end if

end for

We use an alternating maximization scheme to solve (4). We go through the $p \times q$ elements of \mathbf{B} one-by-one and update them to maximize $S_n(\mathbf{B})$ while the other elements are kept fixed. We show that each of these optimization problems are of the form (26) and can be solved easily. Assume

that we want to maximize $S_n(\mathbf{B})$ while all elements except B_{rs} are fixed. Remember that:

$$S_n(\mathbf{B}) = \frac{1}{n \binom{q}{2}} \sum_{i=1}^n \sum_{j=1}^q \sum_{k=1}^q \mathbf{1}(y_{ij} > y_{ik}) \mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j > \mathbf{x}_i^T \mathbf{b}_k)$$

where \mathbf{b}_j and \mathbf{b}_k , respectively, denote the j -th and k -th columns of \mathbf{B} . B_{rs} appears in the sum when either $j = s$ or $k = s$. Simple calculations show that:

$$S_n(\mathbf{B}) = \frac{1}{n \binom{q}{2}} \sum_{i=1}^n \sum_{j=1}^q \left(\mathbf{1}(y_{ij} > y_{is}) \mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j > \mathbf{x}_i^T \mathbf{b}_s) \right. \\ \left. + \mathbf{1}(y_{ij} < y_{is}) \mathbf{1}(\mathbf{x}_i^T \mathbf{b}_j < \mathbf{x}_i^T \mathbf{b}_s) \right) + c,$$

where c denotes the sum of all terms that do not depend on B_{rs} . It is now easy to see that maximizing $S_n(\mathbf{B})$ with respect to B_{rs} is an instance of (26); c disappears in the optimization, and depending on whether $y_{ij} > y_{is}$ or $y_{ij} < y_{is}$, only one of the terms in each set of parentheses remains. In maximizing $S_n(\mathbf{B})$ with respect to B_{rs} , we have $T = nq$; if $y_{ij} > y_{is}$, $u_t = \mathbf{x}_i^T (\mathbf{b}_j - \mathbf{b}_s) + x_{ir} B_{rs}$, and $v_t = -x_{rs}$; and if $y_{ij} < y_{is}$, $u_t = \mathbf{x}_i^T (\mathbf{b}_s - \mathbf{b}_i) - x_{ir} B_{rs}$, and $v_t = x_{rs}$. Computing each u_t and v_t requires $O(p)$ operations, and thus computing all u_t and v_t for $t = 1, \dots, nq$ requires $O(nqp)$ operations. In total, maximizing $S_n(\mathbf{B})$ with respect to one of the elements of \mathbf{B} requires $O(nq(p + \log(nq)))$ operations; $O(nqp)$ to compute the coefficients needed for formulating (26) and $O(nq \log(nq))$ to solve it.

In each full round of the alternating maximization procedure, we go over all pq elements of \mathbf{B} . So, a full round requires $O(npq^2(p + \log(nq)))$ operations. After each round we adjust \mathbf{B} (by subtracting its last column from all its columns and then normalizing it) so that $\mathbf{B} \in \mathcal{B}$. As argued before, this adjustment does not change the objective function. We continue this alternating maximization, until a point when a full round over all elements does not increase $S_n(\mathbf{B})$ —an indication of reaching a fixed point. Since $S_n(\mathbf{B})$ is the sum of $O(nq^2)$ step functions, it takes $O(nq^2)$ different values. Therefore, our algorithm stops at most in $O(nq^2)$ steps (since S_n is increased in each round). Putting all these together, we conclude that our whole algorithm requires $O(n^2pq^4(p + \log(nq)))$ operations, which is polynomial in all three parameters. The pseudo-code of the algorithm is presented in Algorithm 2.

Our algorithm is greedy and it may converge to a local maximum. We can alleviate this problem to some degree by starting the algorithm from different random initial points and choosing the best result. In the next section, we show that our proposed alternating maximization scheme is successful in providing very good estimates of the true coefficient matrix.

6.1. Imposing Sparsity

So far, we formulated the problem of finding \mathbf{B} with no extra structural constraint. In many settings (e.g., high-dimensional, high noise, limited data), solving these problems with no extra constraints results in poor predictive performance, often due to overfitting. Here, we explain how we can impose sparsity on the coefficient matrix in our alternating maximization scheme. The ℓ_0 norm is the precise metric to measure sparsity, but is often relaxed to the ℓ_1 norm to make optimization problems convex. With our proposed formulation, it is possible to use the ℓ_0 norm directly. Let us revisit the optimization problem in (26), but this time with an additional ℓ_0 penalty

Algorithm 2 — for solving (4)

Inputs: $\mathbf{x}_i, \mathbf{y}_i, i = 1, \dots, n$; **Output:** $\widehat{\mathbf{B}}_n$
 Random Initialization: $\mathbf{B} \leftarrow \text{randn}(p, q)$
while $S_n(\mathbf{B})$ changes **do**
 for $r = 1, \dots, p; s = 1, \dots, q$ **do**
 Maximize $S_n(\mathbf{B})$ w.r.t. B_{rs} with
 other elements fixed using Alg. 1
 end for
 $\mathbf{b} \leftarrow \mathbf{B}(:, q)$ \triangleright \mathbf{b} is the last column of \mathbf{B}
 $\mathbf{B} \leftarrow (\mathbf{B} - \mathbf{b}\mathbf{1}_{1 \times q}) / \|\mathbf{B} - \mathbf{b}\mathbf{1}_{1 \times q}\|_F$
end while
 $\widehat{\mathbf{B}}_n \leftarrow \mathbf{B}$

on x . We have

$$\widehat{x} = \arg \max_{x \in \mathbb{R}} \underbrace{\sum_{t=1}^T \mathbf{1}(u_t + v_t x > 0)}_{f(x)} - \lambda \|x\|_0, \quad (27)$$

where λ is the regularization parameter and $\|x\|_0 = 0$ if $x = 0$ and $\|x\|_0 = 1$ if $x \neq 0$. Solving this problem is very similar to solving the unconstrained problem in (26). We only need to solve the problem without the constraint (say we get the solution \widehat{x}) and then compare the value of $f(\widehat{x})$ with $f(0)$. If $f(\widehat{x}) > f(0)$, then we choose \widehat{x} as the solution of the constrained problem and if $f(\widehat{x}) \leq f(0)$, we choose 0 as the solution. Thus, the solution of the constrained problem can be achieved with $O(1)$ extra computations compared to the unconstrained case. Therefore, we can impose element-wise sparsity on \mathbf{B} (i.e., having the penalty term $\lambda \sum_{i,j} \|B_{ij}\|_0$) by solving (27) instead of (26) in each step of the alternating maximization.

7. Simulation Study

7.1. Consistency of Optimizer

In Section 4, we proved that the global maximizer of $S_n(\mathbf{B})$ is a consistent estimator of the coefficient matrix. However, $S_n(\mathbf{B})$ is a highly non-smooth objective function, and it is difficult to find its global maximum in general. Our proposed greedy algorithm in Section 6 is only guaranteed to reach a local maximum. In this section, by extensive simulations, we show that our alternating maximization is able to provide estimates very close to the true matrix for large values of n which is in conformity with consistency.

We use the model in (3) with various utility functions and noise distributions to generate the synthetic data. For fixed p and q (for the simulations in this section, we set $p = q = 5$), we change n from 2^3 to 2^{17} in powers of two, and test the consistency of our alternating maximization scheme, i.e., how close the estimated coefficient matrix, $\widehat{\mathbf{B}}_n$, gets to the true one in the model, \mathbf{B}^* . To

measure the similarity of $\tilde{\mathbf{B}}_n$ and \mathbf{B}^* , we use the following two measures:

$$M_1(\tilde{\mathbf{B}}_n, \mathbf{B}^*) = \|\tilde{\mathbf{B}}_n - \mathbf{B}^*\|_F^2 = \sum_{i,j} (\tilde{B}_{ij} - B_{ij}^*)^2,$$

$$M_2(\tilde{\mathbf{B}}_n, \mathbf{B}^*) = \frac{\sum_{i=1}^p \sum_{j=1}^q (\tilde{B}_{ij} - \tilde{b})(B_{ij}^* - b^*)}{\sqrt{\sum_{i,j} (\tilde{B}_{ij} - \tilde{b})^2} \sqrt{\sum_{i,j} (B_{ij}^* - b^*)^2}},$$

where $\tilde{b} = \sum_{i,j} \tilde{B}_{ij}/pq$ and $b^* = \sum_{i,j} B_{ij}^*/pq$. Thus, M_1 measures how close the elements of the estimated coefficient matrix are to the true elements, and M_2 measures how correlated the elements of the two matrices are.

We generate the data as follows. First, we generate an $n \times p$ predictor matrix, \mathbf{X} , with rows independently drawn from $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_X)$, where the (i, j) -th element of $\boldsymbol{\Sigma}_X$ is defined as $\sigma_{i,j}^X = 0.7^{|j-i|}$. This is a common model for predictors in the literature [29, 30, 31]. The true coefficient matrix, \mathbf{B}^* , has 75% non-zero elements which are drawn independently from a normal distribution $\mathcal{N}(0, 1)$. To satisfy condition **(C4)**, we then subtract its last column from all columns and normalize it so that $\mathbf{B}^* \in \mathcal{B}$.

We consider three types of noise:

- (E1)** The elements of \mathbf{E} are drawn independently from a Gaussian distribution $\mathcal{N}(0, 1)$.
- (E2)** The elements of \mathbf{E} are drawn independently from a t -student distribution with $\nu = 1$.
- (E3)** The elements of \mathbf{E} are drawn independently according to a Gaussian mixture model; each element is drawn from a $\mathcal{N}(0, .2)$ with probability .8 and from a $\mathcal{N}(1, .2)$ with probability .2.

After sampling, to have a consistent signal-to-noise ratio in various scenarios, the elements of \mathbf{E} are scaled accordingly to satisfy $\|\mathbf{E}\|_F / \|\mathbf{X}\mathbf{B}^*\|_F = 0.2$; i.e., the norm of the noise matrix is 20% of the norm of the signal matrix. **E1** corresponds to a general setting where the noise is Gaussian; **E2** simulates a heavy-tailed noise which is present in many practical settings; and **E3** simulates a case where 20% of the data points are corrupted with larger noise (i.e., simulating outliers).

We also consider three different utility functions:

- (U1)** Identity: $U(x) = x$.
- (U2)** Sigmoid: $U(x) = 1/(1 + e^{-x/5})$.
- (U3)** Piecewise constant: $U(x) = \lfloor x \rfloor$.

The identity utility function transforms the problem in (3) to a regular multivariate regression problem. The sigmoid utility function has the nice property of diminishing marginal returns and is of special interest in many applications including economics [32, 33] and network utility maximization [34]. Piecewise constant functions are useful in choice models [35, Chapter 17] and also can be used in modeling the settings where the responses are ordinal, e.g., surveys asking subjects' levels of preference (say 1–10 corresponding to “very poor” to “excellent”) for a set of responses.

Finally, we set $\mathbf{Y} = U(\mathbf{X}\mathbf{B}^* + \mathbf{E})$, where U is applied individually to the elements of its input. For any given \mathbf{X} and \mathbf{Y} , we run Algorithm 2 described in Section 6 with 10 different starting points and choose the result with the highest objective value. The starting points are chosen by drawing elements independently from a standard normal distribution. For each n , we generate ten sets of

\mathbf{X} and \mathbf{Y} as described above and report the median of each of the two performance measures over these ten sets (mean results are similar). The results are shown in Figure 1.

As n becomes larger, the sum-of-squares error, M_1 , goes to 0 and the correlation, M_2 , goes to 1. Therefore, in these cases, our greedy algorithm provides estimates with increasing similarity to the true matrix which is in agreement with consistency. We are increasing n to very large values to show this consistency; however, the algorithm provides very good solutions for all practical purposes at smaller values of n . For instance, for **(U2)** and $n = 128$, the averages of M_1 and M_2 medians for all three noise settings are respectively 0.052 and 0.95. Since $\|\mathbf{B}^*\|_F^2 = 1$, this means that the sum of all squared errors is only 5.2% of the sum of the squares of the true elements, and that the estimated and true coefficients are highly correlated.

7.2. Variable Selection

The results in Figure 1 show that Algorithm 2 provides consistent estimates of \mathbf{B}^* in different settings; if an element of \mathbf{B}^* is zero, then the corresponding estimated element becomes very close to 0 for large values of n . However, in the absence of extra constraints, these estimates may not be exactly zero. Thus, to identify the relevant variables in a model, one has to use ad-hoc thresholding techniques, which is not desirable. The penalized formulation described in Section 6.1, imposes sparsity on the elements of the estimated matrix, resulting in joint variable selection and estimation. In Table 1, we compare the variable selection and estimation of solutions of the penalized version of (4) with those of the unconstrained problem, for $n = 128$ and the sigmoid utility function. Data is generated as before except \mathbf{B}^* has 50% non-zero elements. For the penalized version, we set $\lambda = 5$. We measure the variable selection performance by the following two metrics:

$$\text{Signed Sensitivity} = \frac{\sum_{i,j} \mathbf{1}[B_{i,j}^* \cdot \hat{B}_{i,j} > 0]}{\sum_{i,j} \mathbf{1}[B_{i,j}^* \neq 0]},$$

$$\text{Specificity} = \frac{\sum_{i,j} \mathbf{1}[B_{i,j}^* = 0] \cdot \mathbf{1}[\hat{B}_{i,j} = 0]}{\sum_{i,j} \mathbf{1}[B_{i,j}^* = 0]}.$$

We observe that although both versions of the algorithm produce solutions almost equally close to the true matrix (in terms of M_1), the penalized version provides much better variable selection.




Figure 1: Examining the consistency of the alternating maximization scheme for different noise distributions and utility functions. For both similarity metrics, medians over 10 runs are reported.

In the original model, most of the elements are non-zero (some of them very close to zero), resulting in a poor specificity. The solutions of the penalized version have both high sensitivity and specificity.

Problem	Metric	E1	E2	E3
Original	M_1	0.059	0.013	0.059
	Sensitivity	1	1	1
	Specificity	0.36	0.37	0.34
Penalized	M_1	0.048	0.016	0.050
	Sensitivity	0.89	0.92	0.87
	Specificity	0.71	0.87	0.78

Table 1: Variable selection comparison. For all metrics, medians over 20 runs are reported.

7.3. Comparison with Other Multivariate Regressions Algorithms

Compared with traditional linear multivariate regression algorithms (such as LASSO), our method has two important differences. First, the objective function is defined as the Kendall’s rank correlation between responses and their estimates rather than the sum-of-squares error (with some regularization). The second difference is in the assumed underlying model. We consider a model where responses are related to the predictors through an unknown, potentially non-linear function, U , whereas in the traditional techniques, the relationship is assumed to be linear (or in the generalized regression setting, a known function, e.g., log, of inputs).

When the underlying problem structure is similar to (3), and the objective is the rank correlation, it is expected that our algorithm performs better than traditional models, since it matches the structure and objective better. To verify this expected result, we perform an experiment on simulated data and compare our algorithm with some the state-of-the-art multivariate regression techniques.

The setup is as follows. We assume a model as in (3) with 10 predictors and 10 responses. The elements of the predictor matrix are sampled from a uniform distribution between 0 and 100. Matrix \mathbf{B} is sparse with density 60% and its elements are drawn from a standard uniform distribution. The utility function is a sigmoid: $U(x) = 1/(1 + e^{-x})$. We consider three noise settings as in **(E1)**–**(E3)**. The learning is done over 30 instances and the test is done on a separate set of 20 instances. For all algorithms, the regularization parameters are achieved via 5-fold cross-validation. The algorithms against which we compare our method are Least Squares (LS), a robust version of LS called Least Trimmed Squares (LTS) [36], LASSO [37], Sparse Reduced Rank Regression (SRRR) [38], elastic net [39], and regressions with ridge regularization.

We run each experiment 100 times and report the median and 95% confidence intervals of the improvements in the test rank correlation in Table 2. In essence, we run a paired hypothesis test comparing our algorithm against each of the algorithms in Table 2, and report the median, 2.5’tth percentile, and 97.5’tth percentile of the test statistic, $c_1 - c_2$, where c_1 and c_2 are respectively the test rank correlations of our algorithm and the test rank correlation of the other algorithm. We observe that in all cases, our algorithm performs statistically significantly better than other algorithms.

8. Application to Real Data

In this section, we study two problems with real data where the training set is in the form of an ordering and/or the objective is to order a set of items.

Improvement over		LS	LTS	LASSO	SRRR	ElasticNet	Ridge
E1	median	0.11	0.18	0.09	0.57	0.52	0.50
	95% CI	[0.02, 0.30]	[0.09, 0.51]	[0.01, 0.26]	[0.16, 1.2]	[0.27, 0.86]	[0.24, 0.93]
E2	median	0.11	0.16	0.09	0.62	0.54	0.51
	95% CI	[0.05, 0.41]	[0.07, 0.37]	[0.04, 0.34]	[0.24, 1.2]	[0.27, 0.90]	[0.24, 0.96]
E3	median	0.11	0.17	0.08	0.59	0.48	0.50
	95% CI	[0.01, 0.32]	[0.07, 0.44]	[0.00, 0.29]	[0.21, 1.22]	[0.23, 0.93]	[0.19, 0.88]

Table 2: Improvements achieved over state-of-the-art multivariate regression algorithms by using our proposed rank-based semi-parametric multivariate regression. The objective is the rank correlation between estimates and true values of responses. We report the statistics for $c_1 - c_2$, where c_1 and c_2 are respectively the test rank correlations of our algorithm and the test rank correlation of the other algorithm.

8.1. Sushi Dataset

First, we consider the sushi preference dataset¹. This dataset includes the preference ordering of 10 sushi types² by 5000 users and the demographic information about these users. For each user we keep the following features: gender (male/female), age group (15–19, 20–29, 30–39, 40–49, 50–59, 60+), region in which the user had lived for the longest period until 15 years old (11 in total), and the region in which the user currently lives (11 in total). We represent each feature with a binary indicator vector. For example, for the gender, we use (0, 1) for males and (1, 0) for females. We show different age groups by (1, 0, 0, 0, 0), (0, 1, 0, 0, 0, 0), . . . , (0, 0, 0, 0, 1). Similarly, the regions are represented by binary vectors of size 11. Thus, in total, each user has a feature vector of size 30 representing his/her demographic information.

The goal of our prediction task is to estimate the ordering of the 10 sushi types for a new user only based on his/her demographic information. Note that neither collaborative filtering nor content-based filtering is applicable to this problem, since first, the new user may not have rated or ranked any sushi types beforehand and second, we do not include extra domain knowledge about each sushi type. On the other hand, our regression-based framework is suitable for this prediction task. From the 5000 users, we choose 2500 of them at random as the training set and keep the rest as the test set. We take the average Kendall correlation between the rows of predicted and true orderings for the users in the test set as the performance metric. We repeat this random division of users into training and testing groups 100 times to achieve bootstrapped confidence intervals for the performance metric.

In order to compare our algorithm to other regression-based algorithms, we need to transform the ordering into ratings. We use the technique described in [11] and assign the ratings 1/11, 2/11, 3/11, . . . , 10/11 to the least preferred to most preferred items. Then, for the algorithms that performed well in the simulation study (see Table 2), we follow the same training and testing procedure as explained above. We also compare the results to a K nearest neighbor technique (KNN) where the feature vector of a new user is compared with the available users to identify the K most similar users (in terms of the Euclidean distance), and then its ratings are calculated by

¹<http://www.kamishima.net/sushi/>

²Ebi, Anago, Maguro, Ika, Uni, Sake, Tamago, Toro, Tekka-maki, and Kappa-maki

averaging the ratings of those K neighbors. The parameters of all these models are found via 5-fold cross-validation. The results are shown in Table 3. As we observe, our algorithm outperforms other with high statistical significance.

	Order-based	LASSO	LS	SRRR	KNN
Median	0.34	0.31	0.31	0.18	0.31
95% CI	[0.33, 0.35]	[0.30, 0.32]	[0.30, 0.32]	[0.02, 0.27]	[0.31, 0.32]

Table 3: Comparison of median and 95% confidence intervals of the performance metric (average Kendall correlations between rows of estimated and true responses).

8.2. BIXI Dataset

We use a dataset providing information about Montreal’s bicycle sharing system called BIXI. The data contains the number of available bikes in each of the 400 installed stations for every minute. We use the data collected for the first three weeks of June 2012. From this dataset we first form the features as follows. We allocate two features to each station corresponding to the number of arrivals and departures of bikes to or from that station for every hour. We define two learning tasks: using the number of arrivals and departures in the last hour, estimate the ordering of the number of arrivals and the ordering of the number of departures in the next hour. The result of this estimation is useful in identifying the stations with the highest incoming or outgoing traffic; the BIXI management team can then provide these stations with more bikes or remove the extra bikes accordingly.

Mathematically, we want to estimate \mathbf{B} such that the ordering of elements in the rows of \mathbf{Y} and \mathbf{XB} are as similar as possible, where $\mathbf{X}_{t,j}$ and $\mathbf{X}_{t,400+j}$ respectively show the number of arrivals and departures in hour t at station j and $\mathbf{Y}_{t,j}$ either show the number of arrivals or departures in hour $t + 1$ at station j .

From the roughly 500 hourly data points during the first three weeks of June, we use the first 300 data points for training and use the other 200 points for testing. We compare our algorithm to the multivariate regression algorithms that performed well in Table 2 (LS and LTS cannot be used since we have more features than instances and are replaced by ridge regression as the baseline) and a KNN algorithm. All the parameters of the algorithms are found via 5-fold cross-validation. In Table 4, we compare the average Kendall correlations between corresponding rows of estimated and true response matrices. We observe that our algorithm provides a better estimate (by around 10% to 20%) of the ordering of the stations in terms of the number of arrivals and departures in the next hour.

	Order-based	LASSO	SRRR	Ridge	KNN
Arrivals	0.39	0.33	0.32	0.31	0.36
Departures	0.40	0.33	0.33	0.29	0.35

Table 4: Comparison of Kendall correlations between rows of estimated and true responses.

9. Conclusions

In this paper, we considered a generalized regression problem where the responses are monotonic functions of a linear transformation of the predictors. We proposed a semi-parametric method

based on rank correlation, which is invariant with respect to the functional form of the underlying monotonic function, to estimate the linear transformation. We showed that the solution to our formulated problem is a consistent estimator of the true matrix and identified the convergence rate. To find the solution, we need to maximize a highly non-smooth function. We proposed a greedy algorithm to solve that problem, and showed its success in estimating the true coefficient matrix through simulations over a variety of noise distributions and utility functions. Finally, we presented a penalized version of our problem which has the same computational complexity as the original problem, but results in better variable selection and more interpretable models.

Appendix A. Proof of Lemma 3

Definition 1. Assume that \mathbf{z} is a random vector with a distribution over \mathbb{R}^d . For a given Θ , we say $\mathcal{F} = \{f(\mathbf{z}, \boldsymbol{\theta}), \boldsymbol{\theta} \in \Theta\}$ is a P -degenerate class of functions over \mathbb{R}^d if $E_{\mathbf{z}}[f(\mathbf{z}, \boldsymbol{\theta})] = 0$ for all $\boldsymbol{\theta} \in \Theta$.

Definition 2 (Definition (2.7) from [25]). A class of functions, \mathcal{F} , is called Euclidean for envelope F if there exist constants A and V such that we have: if $0 < \epsilon \leq 1$ and if μ is a measure such that $\int F d\mu < \infty$, then there are functions f_1, f_2, \dots, f_k in \mathcal{F} such that (i) $k \leq A\epsilon^{-V}$ and (ii) \mathcal{F} is covered by the union of closed balls with radius $\epsilon \int F d\mu$ and centers f_1, \dots, f_k . In other words, for each $f \in \mathcal{F}$, there is an f_i with $\int |f - f_i| d\mu \leq \epsilon \int F d\mu$. A and V must not depend on μ .

Lemma 4 (Example (2.11) from [25]). *The class of piece-wise constant functions that are bounded by a fixed function F is Euclidean with envelope F .*

Lemma 5 (Corollaries 17 and 21 from [26]).

(i) *If \mathcal{F} is Euclidean for envelope F and \mathcal{G} is Euclidean for envelope G , then $\mathcal{F} + \mathcal{G}$ is Euclidean with envelope $F + G$.*

(ii) *If \mathcal{F} is a uniformly bounded class of functions, then for each finite measure ν , the class $\nu\mathcal{F}$ is Euclidean.*

Lemma 6 (Corollary 8 from [27]). *Let \mathbf{z} be a d -dimensional random vector, $\mathcal{F} = \{f(\mathbf{z}, \boldsymbol{\theta}), \boldsymbol{\theta} \in \Theta\}$ be a class of P -degenerate functions over \mathbb{R}^d , and $\boldsymbol{\theta}_0$ be a point in Θ for which $f(\mathbf{z}, \boldsymbol{\theta}_0) = 0$ for all \mathbf{z} . If (i) \mathcal{F} is Euclidean for an envelope F satisfying $E[F^2] < \infty$ and (ii) $E_{\mathbf{z}}[f(\mathbf{z}, \boldsymbol{\theta})] \rightarrow 0$ as $\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}_0$, then uniformly over $o_p(1)$ neighborhoods of $\boldsymbol{\theta}_0$, for i.i.d. samples \mathbf{z}_i ,*

$$\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i, \boldsymbol{\theta}) = o_p(1/\sqrt{n}). \quad (\text{A.1})$$

Lemma 3. *For $\boldsymbol{\theta}$ in an $o_p(1)$ neighborhood of $\boldsymbol{\theta}_0$, and $S(\boldsymbol{\theta}) \triangleq E_{\mathbf{z}}[h(\mathbf{z}, \boldsymbol{\theta})]$, we have:*

$$S_n(\boldsymbol{\theta}) = S(\boldsymbol{\theta}) + S_n(\boldsymbol{\theta}_0) - S(\boldsymbol{\theta}_0) + o_p(1/\sqrt{n}). \quad (\text{A.2})$$

Proof. For this proof, we use results from [26, 25, 27]. Define:

$$f(\mathbf{z}, \boldsymbol{\theta}) \triangleq h(\mathbf{z}, \boldsymbol{\theta}) - h(\mathbf{z}, \boldsymbol{\theta}_0) - S(\boldsymbol{\theta}) + S(\boldsymbol{\theta}_0) \quad (\text{A.3})$$

Since $E_{\mathbf{z}}[f(\mathbf{z}, \boldsymbol{\theta})] = S(\boldsymbol{\theta}) - S(\boldsymbol{\theta}_0) - S(\boldsymbol{\theta}) + S(\boldsymbol{\theta}_0) = 0$, $\mathcal{F} = \{f(\mathbf{z}, \boldsymbol{\theta}), \boldsymbol{\theta} \in \Theta\}$ is P -degenerate. All the functions in $\mathcal{H} = \{h(\mathbf{z}, \boldsymbol{\theta}) - h(\mathbf{z}, \boldsymbol{\theta}_0), \boldsymbol{\theta} \in \Theta\}$ are piecewise constant and thus \mathcal{H} is Euclidean

(Lemma 4). Also, $h(\mathbf{z}, \boldsymbol{\theta})$ is uniformly bounded since $|h(\mathbf{z}, \boldsymbol{\theta})| < q^2$. Thus, $\{-S(\boldsymbol{\theta}) + S(\boldsymbol{\theta}_0), \boldsymbol{\theta} \in \Theta\}$ is Euclidean, because it is the expected value of a Euclidean class of uniformly bounded functions (Lemma 5). Since the sum of two Euclidean classes is also Euclidean (Lemma 5), we conclude that \mathcal{F} is Euclidean. An envelope for \mathcal{F} is $4q^2$, since each of its four comprising summand functions is less than q^2 .

We have shown that all the conditions of Lemma 6 hold: \mathcal{F} is a Euclidean class of P -degenerate functions with a constant envelope; $f(\mathbf{z}, \boldsymbol{\theta}_0) = 0$; and $E_{\mathbf{z}} [f(\mathbf{z}, \boldsymbol{\theta})] = 0$. Application of the Lemma gives:

$$\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i, \boldsymbol{\theta}) = o_p(1/\sqrt{n}), \quad (\text{A.4})$$

for $\boldsymbol{\theta}$ in an $o_p(1)$ neighborhood of $\boldsymbol{\theta}_0$. Replacing f with its summands gives the result. \square

Acknowledgement

This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] K.-C. Li and N. Duan, “Regression analysis under link violation,” *The Annals of Statistics*, vol. 17, pp. 1009–1052, 1989.
- [2] H. Ichimura, “Semiparametric least squares (sls) and weighted sls estimation of single-index models,” *Journal of Econometrics*, vol. 58, no. 1, pp. 71–120, 1993.
- [3] Y. Xia and W. Härdle, “Semi-parametric estimation of partially linear single-index models,” *Journal of Multivariate Analysis*, vol. 97, no. 5, pp. 1162–1184, 2006.
- [4] M. Delecroix, M. Hristache, and V. Patilea, “On semiparametric m-estimation in single-index regression,” *Journal of Statistical Planning and Inference*, vol. 136, no. 3, pp. 730–769, 2006.
- [5] X. Yi, Z. Wang, C. Caramanis, and H. Liu, “Optimal linear estimation under unknown non-linear transform,” *arXiv preprint arXiv:1505.03257*, 2015.
- [6] P. Radchenko, “High dimensional single index models,” *Journal of Multivariate Analysis*, vol. 139, pp. 266–282, 2015.
- [7] C. Thrampoulidis, E. Abbasi, and B. Hassibi, “Lasso with non-linear measurements is equivalent to one with linear measurements,” in *Advances in Neural Information Processing Systems*, 2015, pp. 3402–3410.
- [8] N. A. Christakis and D. A. Asch, “Biases in how physicians choose to withdraw life support,” *The Lancet*, vol. 342, no. 8872, pp. 642–646, 1993.
- [9] E. Festic, M. E. Wilson, O. Gajic, G. D. Divertie, and J. T. Rabatin, “Perspectives of physicians and nurses regarding end-of-life care in the intensive care unit,” *Journal of Intensive Care Medicine*, vol. 27, no. 1, pp. 45–54, 2012.

- [10] T. Kamishima, H. Kazawa, and S. Akaho, “A survey and empirical comparison of object ranking methods,” in *Preference learning*. Springer, 2011, pp. 181–201.
- [11] T. Kamishima and S. Akaho, “Nantonac collaborative filtering: A model-based approach,” in *Proc. ACM Conf. on Recommender Systems*, 2010, pp. 273–276.
- [12] O. Luaces, G. F. Bayón, J. R. Quevedo, J. Díez, J. J. Del Coz, and A. Bahamonde, “Analyzing sensory data using non-linear preference learning with feature subset selection,” in *Proc. European Conf. Machine Learning*, 2004, pp. 286–297.
- [13] P. Lops, M. De Gemmis, and G. Semeraro, “Content-based recommender systems: State of the art and trends,” in *Recommender Systems Handbook*. Springer, 2011, pp. 73–105.
- [14] J. Lee, M. Sun, and G. Lebanon, “A comparative study of collaborative filtering algorithms,” *arXiv preprint arXiv:1205.3193*, 2012.
- [15] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez, “Recommender systems survey,” *Knowledge-Based Systems*, vol. 46, pp. 109–132, 2013.
- [16] A. K. Han, “Non-parametric analysis of a generalized regression model: the maximum rank correlation estimator,” *Journal of Econometrics*, vol. 35, no. 2, pp. 303–316, 1987.
- [17] C. Cavanagh and R. P. Sherman, “Rank estimators for monotonic index models,” *Journal of Econometrics*, vol. 84, no. 2, pp. 351–381, 1998.
- [18] J. Abrevaya, “Pairwise-difference rank estimation of the transformation model,” *Journal of Business & Economic Statistics*, vol. 21, no. 3, pp. 437–447, 2003.
- [19] P. D. Allison and N. A. Christakis, “Logit models for sets of ranked items,” *Sociological Methodology*, vol. 24, pp. 199–228, 1994.
- [20] M. Hernandez Alava, J. Brazier, D. Rowen, and A. Tsuchiya, “Common scale valuations across different preference-based measures: Estimation using rank data,” *Medical Decision Making*, vol. 33, no. 6, pp. 839–852, 2013.
- [21] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li, “Learning to rank: from pairwise approach to listwise approach,” in *Proc. Int. Conf. on Machine Learning*, 2007, pp. 129–136.
- [22] T.-Y. Liu, *Learning to rank for information retrieval*. Springer Science & Business Media, 2011.
- [23] A. Klenke, *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.
- [24] R. J. Serfling, *Approximation theorems of mathematical statistics*. John Wiley & Sons, 2009.
- [25] A. Pakes and D. Pollard, “Simulation and the asymptotics of optimization estimators,” *Econometrica*, vol. 57, pp. 1027–1057, 1989.
- [26] D. Nolan and D. Pollard, “U-processes: rates of convergence,” *The Annals of Statistics*, vol. 15, pp. 780–799, 1987.

- [27] R. P. Sherman, “Maximal inequalities for degenerate u-processes with applications to optimization estimators,” *The Annals of Statistics*, vol. 22, pp. 439–459, 1994.
- [28] A. W. Van der Vaart, *Asymptotic statistics*. Cambridge University Press, 2000.
- [29] M. Yuan, A. Ekici, Z. Lu, and R. Monteiro, “Dimension reduction and coefficient estimation in multivariate linear regression,” *Journal of the Royal Statistical Society: Series B*, vol. 69, no. 3, pp. 329–346, 2007.
- [30] J. Peng, J. Zhu, A. Bergamaschi, W. Han, D.-Y. Noh, J. R. Pollack, and P. Wang, “Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer,” *The Annals of Applied Statistics*, vol. 4, no. 1, pp. 53–77, 2010.
- [31] A. J. Rothman, E. Levina, and J. Zhu, “Sparse multivariate regression with covariance estimation,” *Journal of Computational and Graphical Statistics*, vol. 19, no. 4, pp. 947–962, 2010.
- [32] M. Friedman and L. J. Savage, “The utility analysis of choices involving risk,” *Journal of Political Economy*, vol. 56, no. 4, pp. 279–304, 1948.
- [33] A. Tversky and D. Kahneman, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.
- [34] M. Fazel and M. Chiang, “Network utility maximization with nonconcave utilities using sum-of-squares method,” in *Proc. IEEE Conf. on Decision and Control*, 2005, pp. 1867–1874.
- [35] W. H. Greene, *Econometric Analysis (7th Edition)*. Prentice Hall, 2011.
- [36] P. Čížek and J. Á. Víšek, “Least trimmed squares,” in *XploRe – Application Guide*. Springer Berlin Heidelberg, 2000, pp. 49–63.
- [37] R. J. Tibshirani, “Regression shrinkage and selection via the lasso: a retrospective,” *Journal of the Royal Statistical Society: Series B*, vol. 73, no. 3, pp. 273–282, 2011.
- [38] L. Chen and J. Z. Huang, “Sparse reduced-rank regression for simultaneous dimension reduction and variable selection,” *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1533–1545, 2012.
- [39] H. Zou and T. Hastie, “Regularization and variable selection via the elastic net,” *Journal of the Royal Statistical Society: Series B*, vol. 67, no. 2, pp. 301–320, 2005.