
Sparse Linear Isotonic Models

Sheng Chen

Department of Computer Science and Engineering, University of Minnesota, Twin Cities
{shengc, banerjee}@cs.umn.edu

Arindam Banerjee

Abstract

In machine learning and data mining, linear models have been widely used to model the response as parametric linear functions of the predictors. To relax such stringent assumptions made by parametric linear models, additive models consider the response to be a summation of unknown transformations applied on the predictors; in particular, additive isotonic models (AIMs) assume the unknown transformations to be monotone. In this paper, we introduce *sparse linear isotonic models* (SLIMs) for high-dimensional problems by hybridizing ideas in parametric sparse linear models and AIMs, which enjoy a few appealing advantages over both. In the high-dimensional setting, a two-step algorithm is proposed for estimating the sparse parameters as well as the monotone functions over predictors. Under mild statistical assumptions, we show that the algorithm can accurately estimate the parameters. Promising preliminary experiments are presented to support the theoretical results.

1 Introduction

Linear models of the parametric form $y = \langle \tilde{\boldsymbol{\theta}}, \mathbf{x} \rangle + \epsilon$ have enjoyed the enormous popularity in both machine learning and data mining communities for decades, due to its simplicity and interpretability. Here $y \in \mathbb{R}$ and $\mathbf{x} = [x_1, x_2, \dots, x_p]^T \in \mathbb{R}^p$ are observed response and predictors respectively, and $\epsilon \in \mathbb{R}$ is a zero-mean additive noise. It is well known that the unknown parameter $\tilde{\boldsymbol{\theta}} = [\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_p]^T$ can be estimated from multiple instances of (\mathbf{x}, y) , denoted by $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$,

using least squares regression (typically when $n > p$). In recent years, linear models also prove to be able to survive the high-dimensional setting (i.e., $n \ll p$), by exploiting the sparsity of $\tilde{\boldsymbol{\theta}}$ (i.e., $\tilde{\boldsymbol{\theta}}$ has few non-zero entries). Under very mild statistical assumptions, various estimators have been proposed to find $\tilde{\boldsymbol{\theta}}$ with provable guarantees on estimation error, such as Lasso [30, 32] and Dantzig selector [6, 5]. Despite the prevalent success of linear models, modern data often arise from complex environments in which the linear correlation could break down, leading to poor performance of linear models. Progress has been made to relax the stringent assumption of linear models by allowing nonlinearity. In particular, [1] consider the following *additive isotonic models* (AIMs),

$$y = \sum_{j=1}^p f_j(x_j) + \epsilon, \quad (1)$$

where $\{f_j\}_{j=1}^p \triangleq \mathcal{F}$ is a set of *monotone* univariate functions. To estimate \mathcal{F} , a commonly-used procedure is *cyclic pooled adjacent violators* (CPAV). At each iteration of CPAV, *isotonic regression* is called to estimate one f_j and its solution can be efficiently found by the *pooled adjacent violators algorithm* (PAVA) [3]. Though the non-linearity can be captured by \mathcal{F} , one need to specify the monotonicity for each f_j (either increasing or decreasing) in advance, which could be unknown in real-world applications, and enumerating all possible combinations can be computationally prohibitive. In high dimension, the estimation of \mathcal{F} becomes even more challenging, because the number of monotone functions is very large.

To address the challenges in AIMs, we propose the *sparse linear isotonic models* (SLIMs), which assume

$$\mathbb{E}[y|\mathbf{x}] = \sum_{j=1}^p \tilde{\theta}_j f_j(x_j) = \langle \tilde{\boldsymbol{\theta}}, \mathbf{f}(\mathbf{x}) \rangle, \quad (2)$$

where $\mathbf{f}(\mathbf{x}) \triangleq \tilde{\mathbf{x}} = [f_1(x_1), \dots, f_p(x_p)]^T$. SLIMs combine the parametric form from the sparse linear models with the monotone transformations from AIMs, and generalize the assumption of additive noise ϵ to the

conditional expectation form $\mathbb{E}[y|\mathbf{x}]$. Throughout the paper, the parameter $\boldsymbol{\theta}$ is assumed to be *s-sparse*. For identifiability, we also assume w.l.o.g. that each f_j is monotonically increasing (as the monotonicity can be flipped by changing signs of θ_j), and properly normalized such that every $\tilde{x}_j = f_j(x_j)$ is zero-mean and unit-variance. Note that without losing any representational power of AIM, the assumption of increasing f_j avoids the pre-specification of monotonicity for each f_j as required in (1). For such hybrid model, given n i.i.d. samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, our goal is to estimate both $\boldsymbol{\theta}$ and \mathcal{F} . Since the hidden predictor $\tilde{\mathbf{x}}$ is inaccessible, brutally fitting data into a linear model could result in a poor estimate of $\boldsymbol{\theta}$. In this work, we design a two-step algorithm to accomplish this goal, which estimates $\boldsymbol{\theta}$ followed by \mathcal{F} . The estimation of $\boldsymbol{\theta}$ is inspired by the *rank-based* approaches for structure learning of graphical models. At the high level, those approaches do not rely on the exact values of samples generated from the graphical model, in order to learn its structure. Instead they resort to rank correlations (e.g., *Kendall's tau correlation* [16]) that are invariant under monotonically increasing transformation, so that observing \mathbf{x} and $\tilde{\mathbf{x}}$ makes no difference to the method. By leveraging a similar idea, we propose the Kendall's tau Dantzig selector (KDS) to estimate $\boldsymbol{\theta}$, with certain Kendall's tau correlation coefficients appropriately plugged in. Under some distributional assumptions, we show that this estimator is guaranteed to recover a normalized version of $\boldsymbol{\theta}$ with small error. After $\boldsymbol{\theta}$ is estimated, we have a CPAV-type algorithm tailored for estimating transformations \mathcal{F} , which efficiently extends CPAV at little cost.

To sum up, we highlight a few merits of SLIM as follows. Firstly, as aforementioned, SLIM need not specify the monotonicity of f_j whereas AIM requires. Secondly the two-step estimation for SLIM is particularly useful in high-dimensional settings. The estimation of $\boldsymbol{\theta}$ may identify many “don't-care” f_j 's as their corresponding θ_j 's are zero, thus reducing the problem size of estimating \mathcal{F} . Besides, estimating $\tilde{\boldsymbol{\theta}}$ will suffice if one only focuses on variable selection. For the ease of exposition, we introduce a few notations which will be used in the rest of the paper. We let $\mathbf{y} = [y_1, y_2, \dots, y_n]^T \in \mathbb{R}^n$ be the response vector, $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^T \in \mathbb{R}^{n \times p}$ be the observed design matrix, and denote its columns by $\mathbf{x}^j \in \mathbb{R}^n$. Similarly $\tilde{\mathbf{X}}$, $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}^j$ will denote the hidden counterpart of \mathbf{X} , \mathbf{x}_i and \mathbf{x}^j . Matrix is bold capital, and the corresponding bold lowercase is reserved for its rows (columns) with suitable subscripts (superscripts), and its entries are plain lowercase with subscripts indexing both row and column. In general, vectors are bold lowercase while scalars are plain lowercase. For a matrix, $\|\cdot\|_2$ denotes its spectral norm (i.e. the largest eigenvalue)

and $\|\cdot\|_{\max}$ denotes the value of the largest entry in magnitude. The rest of the paper is organized as follows: we first review the related work in Section 2, and provide an overview of estimation for SLIM in Section 3. Next we analyze the recovery of $\tilde{\boldsymbol{\theta}}$ and present the algorithmic details for estimating \mathcal{F} in Section 4. In Section 5, we demonstrate the effectiveness of SLIM through experiments. Section 6 is dedicated to the conclusion. Due to page limit, most of the proofs are deferred to the supplementary material.

2 Related Work

AIM was initially proposed in [1]. [23] established the asymptotic properties of the CPAV procedure. The high-dimensional counterpart of AIMs (i.e., assuming most of f_j 's are zero), Lasso ISO (LISO), was studied by [10], where a modified CPAV is used to achieved the sparsity in \mathcal{F} . [8] considered a semiparametric additive isotonic model by introducing an additional parametric model into (1). On the other hand, [14] considered an additive model of the same form as (1) for general \mathcal{F} . With suitable smoothing operator on f_j 's, a coordinate descent procedure called *backfitting* can be applied to estimating \mathcal{F} . In high-dimensional regime, [28] correspondingly investigated the sparse additive models (SpAMs), which is solved a backfitting algorithm with extra soft-thresholding steps. Many other efforts have been spent by relying on the smoothness of f_j 's, including [18], [24], [15], and etc.

The method we use to estimate $\tilde{\boldsymbol{\theta}}$ is closely related to the high-dimensional structure learning of graphical models. For sparse Gaussian graphical model, [25] proposed a neighborhood selection procedure for estimating the graph structure, which iteratively regresses each variable against the rest via Lasso. The neighborhood Dantzig selector [34] shares the similar spirit with this approach, which switches Lasso to Dantzig selector. Recent progress has shown that these approaches continue to work for some non-Gaussian distributions, such as *nonparanormal distribution* [21], by using rank correlations to approximate the latent correlation matrix [33, 19]. Similar results have been further generalized to *transelliptical distribution* [20, 12, 13].

3 Overview of Two-Step Algorithm

In this section, we present an overview of the two-step algorithm for the estimation of SLIM, which first estimates $\tilde{\boldsymbol{\theta}}$ and then \mathcal{F} .

For the estimation of $\tilde{\boldsymbol{\theta}}$, if the hidden design matrix $\tilde{\mathbf{X}}$ could be observed, Dantzig selector [6] can be used to

estimate $\tilde{\boldsymbol{\theta}}$ as normal linear models,

$$\hat{\boldsymbol{\theta}}_{\text{orc}} = \underset{\boldsymbol{\theta} \in \mathbb{R}^p}{\operatorname{argmin}} \|\boldsymbol{\theta}\|_1 \text{ s.t. } \left\| \frac{1}{n} \tilde{\mathbf{X}}^T (\tilde{\mathbf{X}}\boldsymbol{\theta} - \mathbf{y}) \right\|_{\infty} \leq \gamma_n, \quad (3)$$

where γ_n is a tuning parameter. A key observation from (3) is that instead of exactly knowing $\tilde{\mathbf{X}}$ and \mathbf{y} , it is sufficient to be given the (approximate) value of $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n}$ and $\frac{\tilde{\mathbf{X}}^T \mathbf{y}}{n}$ in order for (3) to work. Note that the quantity $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n}$ and its expectation $\tilde{\boldsymbol{\Sigma}} = \mathbb{E}[\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T]$ also arise in the structure learning of nonparanormal graphical models. Specifically if $\tilde{\mathbf{x}}$ follows a multivariate Gaussian $\mathcal{N}(\mathbf{0}, \tilde{\boldsymbol{\Sigma}})$, then the observed predictor \mathbf{x} , represented as $\mathbf{f}^{-1}(\tilde{\mathbf{x}}) \triangleq [f_1^{-1}(\tilde{x}_1), \dots, f_p^{-1}(\tilde{x}_p)]^T$, is by definition a nonparanormal distribution $NPN(\tilde{\boldsymbol{\Sigma}}, \mathbf{f}^{-1})$, in which $\tilde{\boldsymbol{\Sigma}}$ is often called *latent correlation matrix*. Simply speaking, the nonparanormal distribution models the random vector whose coordinates are element-wise monotone transformations of a Gaussian random vector. To estimate $\tilde{\boldsymbol{\Sigma}}$ without knowing \mathbf{f} or \mathbf{f}^{-1} , Kendall's tau correlation coefficient [16] plays a key role in rank-based methods. Given data $\mathbf{X} = [x_{ij}] \in \mathbb{R}^{n \times p}$, we define the *sample Kendall's tau correlation matrix* $\hat{\mathbf{T}} = [\hat{t}_{ij}] \in \mathbb{R}^{p \times p}$ as

$$\hat{t}_{ij} = \sum_{1 \leq k, k' \leq n} \frac{\operatorname{sign}((x_{ki} - x_{k'i})(x_{kj} - x_{k'j}))}{n(n-1)}, \quad (4)$$

and its transformed version $\hat{\boldsymbol{\Sigma}} = [\hat{\sigma}_{ij}] \in \mathbb{R}^{p \times p}$,

$$\hat{\sigma}_{ij} = \sin\left(\frac{\pi}{2} \hat{t}_{ij}\right), \quad (5)$$

One straightforward yet critical property of $\hat{\mathbf{T}}$ and $\hat{\boldsymbol{\Sigma}}$ is the invariance to monotone increasing transformations on columns of \mathbf{X} , indicating that the two quantities remain unchanged if \mathbf{X} is replaced by $\tilde{\mathbf{X}}$ in the definitions. More importantly, later analysis will reveal for the class of transelliptical distributions (a generalization of nonparanormal distribution) the closeness between the transformed sample Kendall's tau correlation matrix $\hat{\boldsymbol{\Sigma}}$ and the latent correlation matrix $\tilde{\boldsymbol{\Sigma}}$, thus $\hat{\boldsymbol{\Sigma}}$ can serve as an approximation to $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n}$ as $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n} \approx \tilde{\boldsymbol{\Sigma}}$ in expectation. For $\frac{\tilde{\mathbf{X}}^T \mathbf{y}}{n}$ and its expectation $\tilde{\boldsymbol{\beta}} = \mathbb{E}[y\tilde{\mathbf{x}}] = \tilde{\boldsymbol{\Sigma}}\tilde{\boldsymbol{\theta}}$, we similarly define the *sample Kendall's tau correlation vector* $\hat{\mathbf{b}} \in \mathbb{R}^p$ and its transformation $\hat{\boldsymbol{\beta}}$

$$\hat{b}_j = \sum_{1 \leq k, k' \leq n} \frac{\operatorname{sign}((x_{kj} - x_{k'j})(y_k - y_{k'}))}{n(n-1)}, \quad (6)$$

$$\hat{\beta}_j = \sin\left(\frac{\pi}{2} \hat{b}_j\right), \quad (7)$$

and use $\hat{\boldsymbol{\beta}}$ as a replacement for $\frac{\tilde{\mathbf{X}}^T \mathbf{y}}{n}$. Therefore the estimation of $\tilde{\boldsymbol{\theta}}$ can proceed with (3) by replacing $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n}$

Algorithm 1 Estimating $\tilde{\boldsymbol{\theta}}$ and \mathcal{F} for SLIM

Input: $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{y} \in \mathbb{R}^n$, tuning parameter γ_n

Output: Estimated $\tilde{\boldsymbol{\theta}}$ for $\boldsymbol{\theta}$ and $\hat{\mathcal{F}}$ for \mathcal{F}

- 1: Compute the transformed sample Kendall's tau correlation matrix $\hat{\boldsymbol{\Sigma}}$ and vector $\hat{\boldsymbol{\beta}}$ using (4) - (7)
 - 2: Estimate $\tilde{\boldsymbol{\theta}}$ via Kendall's tau Dantzig selector (8)
 - 3: $\tilde{\boldsymbol{\theta}} := \hat{\sigma}_y \tilde{\boldsymbol{\theta}}$, where $\hat{\sigma}_y$ is the sample variance of \mathbf{y}
 - 4: Estimate the hidden design $\tilde{\mathbf{X}}$ via (9)
 - 5: $\hat{\mathcal{F}} := \{\hat{f}_j\}_{j=1}^p$, where \hat{f}_j is given by (10)
 - 6: **Return** $\tilde{\boldsymbol{\theta}}$ and $\hat{\mathcal{F}}$
-

and $\frac{\tilde{\mathbf{X}}^T \mathbf{y}}{n}$ with $\hat{\boldsymbol{\Sigma}}$ and $\hat{\boldsymbol{\beta}}$ respectively, which leads to the following estimator which we call *Kendall's tau Dantzig selector* (KDS),

$$\tilde{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \mathbb{R}^p}{\operatorname{argmin}} \|\boldsymbol{\theta}\|_1 \text{ s.t. } \left\| \hat{\boldsymbol{\Sigma}}\boldsymbol{\theta} - \hat{\boldsymbol{\beta}} \right\|_{\infty} \leq \gamma_n. \quad (8)$$

Later it will be shown in the analysis that the $\tilde{\boldsymbol{\theta}}$ only approximates the direction of $\boldsymbol{\theta}$, and the scale should be attached on the final estimate $\tilde{\boldsymbol{\theta}}$ by calculating the sample variance of \mathbf{y} .

To estimate the transformations \mathcal{F} , one needs to first find out an $\tilde{\mathbf{X}} = [\hat{x}_{ij}]$ that approximates the hidden design $\tilde{\mathbf{X}} = [f_j(x_{ij})]$ for the observed $\mathbf{X} = [x_{ij}]$, which essentially gives us the estimated values of each f_j at n points x_{1j}, \dots, x_{nj} . To be specific, we fit $\tilde{\mathbf{X}}$ into \mathbf{y} and the estimated $\tilde{\boldsymbol{\theta}}$ through the following convex program,

$$\tilde{\mathbf{X}} = \underset{\mathbf{Z} \in \mathbb{R}^{n \times p}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{Z}\tilde{\boldsymbol{\theta}} - \mathbf{y}\|_2^2 \quad (9)$$

s.t. $\mathbf{z}^j \in \mathcal{M}(\mathbf{x}^j)$, $\mathbf{1}^T \mathbf{z}^j = 0$, $\|\mathbf{z}^j\|_2 \leq \sqrt{n}$, $\forall 1 \leq j \leq p$,

where $\mathcal{M}(\mathbf{x}) = \{\mathbf{v} \mid v_i \geq v_j \text{ iff } x_i \geq x_j, \forall 1 \leq i, j \leq p\}$. In order to get the f_j defined everywhere, we need to interpolate the n estimated points $\hat{x}_{1j}, \dots, \hat{x}_{nj}$. In the algorithm, we simply use nearest-neighbor interpolation as follows,

$$\hat{f}_j(x) = \sum_{i=1}^n \hat{x}_{ij} \cdot \mathbb{I}\left\{i = \underset{1 \leq k \leq n}{\operatorname{argmin}} |x_{kj} - x|\right\}, \quad (10)$$

where $\mathbb{I}\{\cdot\}$ is the indicator function that outputs one if the predicate is true and zero otherwise. Other interpolation technique, e.g., linear/spline interpolation, can be applied in the need of certain desired properties of f_j . The full estimation algorithm is given in Algorithm 1.

4 Statistical and Algorithmic Analysis

In this section, we detail the Algorithm 1 in several aspects. We analyze the recovery guarantee of the KDS

for estimating $\tilde{\boldsymbol{\theta}}$. Under the assumption of transelliptically distributed (\mathbf{x}, y) and the so-called *sign sub-Gaussian condition*, we show that the sample complexity of KDS can be sharpened compared with [33, 19]. To estimate \mathcal{F} , we present a CPAV-type algorithm for solving (9), where each step can be solved almost at no more cost than isotonic regression.

4.1 Recovery guarantee of $\tilde{\boldsymbol{\theta}}$

In this subsection, we consider the estimation of $\tilde{\boldsymbol{\theta}}$. The KDS (8) can be casted as a linear program, which can be solved efficiently by many optimization algorithms [7, 17]. Hence we focus on the statistical aspect of KDS. From Section 3, we know that the success of KDS relies on $\tilde{\boldsymbol{\Sigma}}$ and $\tilde{\boldsymbol{\beta}}$, which replace $\frac{\tilde{\mathbf{X}}^T \tilde{\mathbf{X}}}{n}$ and $\frac{\tilde{\mathbf{X}}^T \mathbf{y}}{n}$ in the Dantzig selector (3). Hence we first investigate the property of $\tilde{\boldsymbol{\Sigma}}$ and $\tilde{\boldsymbol{\beta}}$. The definition (4) - (7) are sample versions of (transformed) Kendall's tau correlation matrix and vector. Here we define their population counterparts.

Definition 1 Given (\mathbf{x}, y) and its independent copy (\mathbf{x}', y') , the population Kendall's tau correlation matrix $\mathbf{T} = [t_{ij}] \in \mathbb{R}^{p \times p}$ and vector $\mathbf{b} \in \mathbb{R}^p$ are defined as

$$t_{ij} = \mathbb{P}((x_i - x'_i)(x_j - x'_j) > 0) - \mathbb{P}((x_i - x'_i)(x_j - x'_j) < 0), \quad (11)$$

$$b_j = \mathbb{P}((x_j - x'_j)(y - y') > 0) - \mathbb{P}((x_j - x'_j)(y - y') < 0), \quad (12)$$

and their transformed versions $\boldsymbol{\Sigma} = [\sigma_{ij}] \in \mathbb{R}^{p \times p}$ and $\boldsymbol{\beta} \in \mathbb{R}^p$ are given by

$$\sigma_{ij} = \sin\left(\frac{\pi}{2} t_{ij}\right), \quad \beta_j = \sin\left(\frac{\pi}{2} b_j\right). \quad (13)$$

To specify the statistical assumptions, we first introduce two family of distributions, *elliptical* and *transelliptical*. The transelliptical distribution is defined based on the elliptical distribution given as follows.

Definition 2 (Elliptical distribution) A random vector $\mathbf{z} \in \mathbb{R}^p$ follows an elliptical distribution $EC(\boldsymbol{\mu}, \tilde{\boldsymbol{\Sigma}}, \xi)$ iff \mathbf{z} has a stochastic representation:

$$\mathbf{z} \sim \boldsymbol{\mu} + \xi \mathbf{A} \mathbf{u}. \quad (14)$$

Here $\boldsymbol{\mu} \in \mathbb{R}^p$, $q \triangleq \text{rank}(\mathbf{A})$, $\mathbf{A} \in \mathbb{R}^{p \times q}$, $\xi \geq 0$ is a random variable independent of \mathbf{u} , $\mathbf{u} \in \mathbb{S}^{q-1}$ is uniformly distributed on the unit sphere in \mathbb{R}^q , and $\mathbf{A} \mathbf{A}^T = \tilde{\boldsymbol{\Sigma}}$. Note that

$$\mathbb{E}[\mathbf{z}] = \boldsymbol{\mu}, \quad \text{Cov}[\mathbf{z}] = \frac{\mathbb{E}[\xi^2]}{q} \tilde{\boldsymbol{\Sigma}}. \quad (15)$$

This family of distribution contains the Gaussian distribution as a special case, and more details can be found in [9]. The extension from elliptical to transelliptical distribution parallels that from normal to nonparanormal distribution.

Definition 3 (Transelliptical distribution)

A random vector $\mathbf{x} \in \mathbb{R}^p$ follows the transelliptical distribution $TE(\tilde{\boldsymbol{\Sigma}}, \xi, \mathbf{f})$ if $\mathbf{f}(\mathbf{x}) = [f_1(x_1), f_2(x_2), \dots, f_p(x_p)]^T \sim EC(\boldsymbol{\mu}, \tilde{\boldsymbol{\Sigma}}, \xi)$, where f_1, f_2, \dots, f_p are all strictly increasing functions, $\boldsymbol{\mu} = \mathbf{0}$, $\text{diag}(\tilde{\boldsymbol{\Sigma}}) = \mathbf{I}$, and $\mathbb{P}(\xi = 0) = 0$.

The conditions on $\boldsymbol{\mu}$ and $\text{diag}(\tilde{\boldsymbol{\Sigma}})$ are imposed for identifiability. If the underlying elliptical distribution is multivariate Gaussian, then the transelliptical family is reduced to the nonparanormal. We refer the readers to [20] for more discussions on transelliptical distribution. Based on the elliptical and transelliptical family, we introduce our assumptions on distribution of (\mathbf{x}, y) :

- $\mathbf{x} \in \mathbb{R}^p$ follows transelliptical distribution $TE(\tilde{\boldsymbol{\Sigma}}, \xi, \mathbf{f})$, or equivalently $\tilde{\mathbf{x}}$ follows elliptical distribution $EC(\mathbf{0}, \tilde{\boldsymbol{\Sigma}}, \xi)$, where $\mathbb{E}[\xi^2] = p$.
- The smallest eigenvalue λ_{\min} of $\tilde{\boldsymbol{\Sigma}}$ is strictly positive.
- $\tilde{\mathbf{x}}$ and y are jointly elliptically distributed.

The assumption $\mathbb{E}[\xi^2] = p$ is also out of the consideration of identifiability. The last assumption on the joint distribution of $(\tilde{\mathbf{x}}, y)$ may seem obscure. But it can be satisfied, for example, when \mathbf{x} is nonparanormal and y is a noisy observation of $\langle \tilde{\boldsymbol{\theta}}, \mathbf{f}(\mathbf{x}) \rangle$ perturbed by an additive zero-mean Gaussian noise. Under these assumptions, we have the following recovery guarantee for the KDS $\hat{\boldsymbol{\theta}}$.

Theorem 1 For any s -sparse $\tilde{\boldsymbol{\theta}}$, if we choose $\gamma_n = \frac{5\pi}{\sqrt{\lambda_{\min}}} \sqrt{\frac{s \log p}{n}}$ and $n \geq \left(\frac{24\pi}{\lambda_{\min}}\right)^2 s^2 \log p$, with probability at least $1 - \frac{2}{p} - \frac{1}{p^{2.5}}$, $\hat{\boldsymbol{\theta}}$ given by (8) satisfies

$$\left\| \hat{\boldsymbol{\theta}} - \frac{\tilde{\boldsymbol{\theta}}}{\sigma_y} \right\|_2 \leq \frac{40\pi}{\lambda_{\min}^{3/2}} \sqrt{\frac{s^2 \log p}{n}}, \quad (16)$$

In the theorem above, though KDS only approximates a normalized version of $\tilde{\boldsymbol{\theta}}$, the scale σ_y can be estimated by computing the sample variance $\hat{\sigma}_y^2$ of \mathbf{y} , and the final estimate of $\tilde{\boldsymbol{\theta}}$ is $\hat{\boldsymbol{\theta}} = \hat{\sigma}_y \tilde{\boldsymbol{\theta}}$ as given in Algorithm 1.

To prove the theorem above, we need to characterize certain properties of $\tilde{\boldsymbol{\Sigma}}$ and $\tilde{\boldsymbol{\beta}}$. One notable result that has been shown for $\boldsymbol{\Sigma}$, $\tilde{\boldsymbol{\Sigma}}$ and $\tilde{\boldsymbol{\beta}}$ is given in the following lemma.

Lemma 1 For $\mathbf{x} \sim TE(\tilde{\Sigma}, \xi, \mathbf{f})$, the transformed population Kendall's tau correlation matrix Σ satisfies

$$\Sigma = \tilde{\Sigma}, \quad (17)$$

and the sample version $\hat{\Sigma}$ for Σ defined in (5), with probability at least $1 - p^{-2.5}$, satisfies

$$\|\hat{\Sigma} - \tilde{\Sigma}\|_{\max} \leq 3\pi \sqrt{\frac{\log p}{n}} \quad (18)$$

The lemma is essentially Theorem 3.2 and 4.1 in [13]. Similarly we have the following lemma for β , $\hat{\beta}$ and $\tilde{\beta}$.

Lemma 2 The transformed population Kendall's tau correlation vector β satisfies

$$\beta = \frac{\tilde{\beta}}{\sigma_y} = \frac{\tilde{\Sigma}\tilde{\theta}}{\sigma_y}, \quad (19)$$

where σ_y^2 is the variance of y . The transformed sample Kendall's tau correlation vector $\hat{\beta}$, with probability at least $1 - \frac{2}{p}$, satisfies

$$\|\hat{\beta} - \beta\|_{\infty} \leq 2\pi \sqrt{\frac{\log p}{n}} \quad (20)$$

The proof is included in the supplementary material. In the light of Lemma 1 and 2, it becomes clear that $\frac{\tilde{\mathbf{x}}^T \tilde{\mathbf{x}}}{n}$ and $\frac{\tilde{\mathbf{x}}^T \tilde{\mathbf{y}}}{n}$ in (3) are replaced by $\hat{\Sigma}$ and $\hat{\beta}$ in (8). The population counterpart of $\hat{\Sigma}$ is $\Sigma = \tilde{\Sigma} = \mathbb{E}[\frac{\tilde{\mathbf{x}}^T \tilde{\mathbf{x}}}{n}]$. Unfortunately, the population version β of $\hat{\beta}$ is not equal to $\tilde{\beta} = \mathbb{E}[\frac{\tilde{\mathbf{x}}^T \tilde{\mathbf{y}}}{n}]$, which is additionally normalized by σ_y . Therefore we will see later that KDS recovers a scaled $\tilde{\theta}$.

In order to bound the estimation error, we additionally need to show that the transformed sample Kendall's tau correlation matrix $\hat{\Sigma}$ satisfies the following *restricted eigenvalue* (RE) condition [5, 35, 27, 26, 2], which is critical in the recovery analysis.

Lemma 3 Define the descent cone for any s -sparse vector $\theta^* \in \mathbb{R}^p$,

$$\mathcal{C} = \{\mathbf{v} \in \mathbb{R}^p \mid \|\theta^* + \mathbf{v}\|_1 \leq \|\theta^*\|_1\}. \quad (21)$$

If $\mathbf{x} \sim TE(\tilde{\Sigma}, \xi, \mathbf{f})$ and $n \geq \left(\frac{24\pi}{\lambda_{\min}}\right)^2 s^2 \log p = O(s^2 \log p)$, with probability at least $1 - p^{-2.5}$, the following RE condition holds for $\hat{\Sigma}$ in \mathcal{C} ,

$$\inf_{\mathbf{v} \in \mathcal{C} \cap \mathbb{S}^{p-1}} \mathbf{v}^T \hat{\Sigma} \mathbf{v} \geq \frac{\lambda_{\min}}{2}, \quad (22)$$

where λ_{\min} is the smallest eigenvalue of $\tilde{\Sigma}$.

Remark: The proof is given in the supplementary material. Similar proof steps appear in [33] amid the analysis of rank-based neighborhood Dantzig selector, in which the concept of RE condition is not formulated. Here we single out this lemma in order for the later comparison in Section 4.2.

Now we are ready to present the proof of Theorem 1.

Proof of Theorem 1: For the sake of convenience, we denote $\theta^* = \frac{\tilde{\theta}}{\sigma_y}$, and it is easy to see that $\tilde{\Sigma}\theta^* = \beta$. We first show that θ^* is feasible when $\gamma_n = \frac{5\pi}{\sqrt{\lambda_{\min}}} \sqrt{\frac{s \log p}{n}}$, by bounding the left-hand side of the constraint for θ^* .

$$\begin{aligned} \|\hat{\Sigma}\theta^* - \hat{\beta}\|_{\infty} &= \left\| \left(\hat{\Sigma} - \tilde{\Sigma} \right) \theta^* - (\hat{\beta} - \beta) \right\|_{\infty} \\ &\leq \left\| \left(\hat{\Sigma} - \tilde{\Sigma} \right) \theta^* \right\|_{\infty} + \|\hat{\beta} - \beta\|_{\infty} \\ &\leq \|\theta^*\|_1 \|\hat{\Sigma} - \tilde{\Sigma}\|_{\max} + 2\pi \sqrt{\frac{\log p}{n}} \\ &\leq \sqrt{s} \cdot \|\theta^*\|_2 \|\hat{\Sigma} - \tilde{\Sigma}\|_{\max} + 2\pi \sqrt{\frac{\log p}{n}} \\ &\leq \frac{3\pi}{\sqrt{\lambda_{\min}}} \sqrt{\frac{s \log p}{n}} + 2\pi \sqrt{\frac{\log p}{n}} \leq \frac{5\pi}{\sqrt{\lambda_{\min}}} \sqrt{\frac{s \log p}{n}}, \end{aligned}$$

where we use Lemma 1 and 2, and thus θ^* is feasible with probability $1 - \frac{2}{p} - \frac{1}{p^{2.5}}$ by union bound. On the other hand, since $\check{\theta}$ is optimal solution to (8), it satisfies

$$\|\check{\theta}\|_1 \leq \|\theta^*\|_1 \quad \text{and} \quad \left\| \hat{\Sigma}\check{\theta} - \hat{\beta} \right\|_{\infty} \leq \gamma_n.$$

Letting $\mathbf{z} = \check{\theta} - \theta^*$, we thus have

$$\begin{aligned} \left\| \hat{\Sigma} \mathbf{z} \right\|_{\infty} &\leq \left\| \hat{\Sigma} \check{\theta} - \hat{\beta} \right\|_{\infty} + \left\| \hat{\Sigma} \theta^* - \hat{\beta} \right\|_{\infty} \leq 2\gamma_n \implies \\ \mathbf{z}^T \hat{\Sigma} \mathbf{z} &= \left\langle \mathbf{z}, \hat{\Sigma} \mathbf{z} \right\rangle \leq \|\mathbf{z}\|_1 \left\| \hat{\Sigma} \mathbf{z} \right\|_{\infty} \leq 2\gamma_n \|\mathbf{z}\|_1 \end{aligned}$$

Using Lemma 3 combined with the inequality above, with probability at least $1 - \frac{2}{p} - \frac{1}{p^{2.5}}$, we get

$$\begin{aligned} \frac{\lambda_{\min}}{2} \|\mathbf{z}\|_2^2 &\leq \mathbf{z}^T \hat{\Sigma} \mathbf{z} \leq 2\gamma_n \|\mathbf{z}\|_1 \implies \\ \|\mathbf{z}\|_2 &\leq \frac{4\gamma_n}{\lambda_{\min}} \frac{\|\mathbf{z}\|_1}{\|\mathbf{z}\|_2} \leq \frac{40\pi}{\lambda_{\min}^{3/2}} \sqrt{\frac{s^2 \log p}{n}}, \end{aligned}$$

where we use the fact that $\sup_{\mathbf{z} \in \mathcal{C}} \frac{\|\mathbf{z}\|_1}{\|\mathbf{z}\|_2} \leq 2\sqrt{s}$. \blacksquare

4.2 Improved RE condition

From the result stated in Lemma 3, we see that the $O(s^2 \log p)$ sample complexity for RE condition of $\hat{\Sigma}$ is worse than that of $\frac{\tilde{\mathbf{x}}^T \tilde{\mathbf{x}}}{n}$, which is $O(s \log p)$ [5, 26].

Next we show that this sharper bound (see Theorem 2) can be obtained for $\hat{\Sigma}$ if the distribution of \mathbf{x} further satisfies the *sign sub-Gaussian condition* [12]. This result may be of independent interest.

Definition 4 (Sign sub-Gaussian condition)

For a random variable x , the operator $\psi : \mathbb{R} \mapsto \mathbb{R}$ is defined as

$$\begin{aligned} \psi(x; \alpha, t_0) &\triangleq \inf \left\{ c > 0 : \mathbb{E} \exp\{t(x^\alpha - \mathbb{E}x^\alpha)\} \right. \\ &\quad \left. \leq \exp(ct^2), \text{ for } |t| < t_0 \right\}. \end{aligned} \quad (23)$$

The random vector $\mathbf{x} \in \mathbb{R}^p$ satisfies the sign sub-Gaussian condition iff

$$\sup_{\mathbf{v} \in \mathbb{S}^{p-1}} \psi(\langle \text{sign}(\mathbf{x} - \mathbf{x}'), \mathbf{v} \rangle; 2, t_0) \leq \kappa \|\mathbf{T}\|_2^2, \quad (24)$$

for a fixed constant κ and some $t_0 > 0$ such that $t_0 \kappa \|\mathbf{T}\|_2^2$ is lower bounded by a fixed constant, where \mathbf{x}' is an independent copy of \mathbf{x} and \mathbf{T} is the population Kendall's tau correlation matrix defined in (11).

Detailed discussions on the sign sub-Gaussian condition can be found in [12], which is out of the scope of this paper. In particular, [12] show that if sign sub-Gaussian condition for transelliptical \mathbf{x} , the $\hat{\Sigma}$ will converge with high probability to $\tilde{\Sigma}$ at rate $O(\sqrt{\frac{s \log p}{n}})$ in terms of *restricted spectral norm*,

$$\|\hat{\Sigma} - \tilde{\Sigma}\|_{2,s} \triangleq \sup_{\substack{\mathbf{v} \in \mathbb{S}^{p-1} \\ \|\mathbf{v}\|_0 \leq s}} \left| \mathbf{v}^T (\hat{\Sigma} - \tilde{\Sigma}) \mathbf{v} \right| = O\left(\sqrt{\frac{s \log p}{n}}\right). \quad (25)$$

Starting from this result, we show that with high probability the RE condition will hold for $\hat{\Sigma}$ with $O(s \log p)$ samples.

Theorem 2 *Let $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^T$ be i.i.d. samples of $\mathbf{x} \sim TE(\tilde{\Sigma}, \xi, \mathbf{f})$ for which the sign sub-Gaussian condition holds with constant κ . Define the constant*

$$c_0 = \max \left\{ \frac{320\kappa\pi^4 \|\tilde{\Sigma}\|_2^2}{\lambda_{\min}^2}, \frac{\pi^2}{\lambda_{\min}} \right\},$$

in which λ_{\min} is the smallest eigenvalue $\tilde{\Sigma}$. If $n \geq \frac{128c_0}{\lambda_{\min}} s \log p = O(s \log p)$, with probability at least $1 - \frac{2}{p} - \frac{1}{p^2}$, $\hat{\Sigma}$ satisfies the following RE condition,

$$\inf_{\mathbf{v} \in \mathcal{C} \cap \mathbb{S}^{p-1}} \mathbf{v}^T \hat{\Sigma} \mathbf{v} \geq \frac{\lambda_{\min}}{2}, \quad (26)$$

where \mathcal{C} is defined in (21).

Remark: The proof of Theorem 2 is deferred to the supplementary material. Note that Theorem 1 relies on the RE condition described in Lemma 3, but we emphasize that if sign sub-Gaussian condition holds we can obtain similar result as long as n attains the bound in Theorem 2, which is smaller than the one required in Lemma 3.

4.3 Computation of \mathcal{F}

After $\hat{\theta}$ is obtained, we can turn to the estimation of transformations \mathcal{F} . As we only have access to a finite number of samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, it is impossible to know the exact function. Hence we use the simple nearest-neighbor interpolation to approximate the f_j as mentioned in (10). By leveraging the monotonicity of f_j , we can estimate $\hat{\mathbf{X}}$ via solving the constrained least squares problem below,

$$\begin{aligned} \hat{\mathbf{X}} &= \underset{\mathbf{Z} \in \mathbb{R}^{n \times p}}{\text{argmin}} \ell(\mathbf{Z}) = \frac{1}{2} \|\mathbf{Z}\hat{\theta} - \mathbf{y}\|_2^2 \\ \text{s.t. } &\mathbf{z}^j \in \mathcal{M}(\mathbf{x}^j), \forall 1 \leq j \leq p, \end{aligned} \quad (27)$$

where the set $\mathcal{M}(\mathbf{x})$ denotes the *monotone cone* induced by vector \mathbf{x} , i.e.,

$$\mathcal{M}(\mathbf{x}) = \{\mathbf{v} \mid v_i \geq v_j \text{ iff } x_i \geq x_j, \forall 1 \leq i, j \leq p\}. \quad (28)$$

The problem (27) is convex w.r.t. \mathbf{Z} . Note that if $\hat{\theta} = \mathbf{1}$, the problem (27) is reduced to the estimation of \mathcal{F} in AIM, which can be solved by the CPAV algorithm. Hence similar CPAV-type algorithm applies here, which is essentially a procedure of cyclic block coordinate descent (BCD) with exact minimization (i.e., minimizing $\ell(\mathbf{Z})$ w.r.t. each \mathbf{z}^j cyclically while keeping other blocks fixed). In this scheme, each subproblem turns out to be an isotonic regression [3]. To be specific, we let $\hat{\mathbf{X}}_{(k)}$ be the iterate of the k -th round update, and define the residue for the j -th block as

$$\mathbf{r}_{(k)}^j = \mathbf{y} - \sum_{i < j} \hat{\theta}_i \hat{\mathbf{x}}_{(k)}^i - \sum_{i > j} \hat{\theta}_i \hat{\mathbf{x}}_{(k-1)}^i. \quad (29)$$

Then each $\hat{\mathbf{x}}_{(k)}^j$ is obtained by solving

$$\hat{\mathbf{x}}_{(k)}^j = \underset{\mathbf{z}^j \in \mathcal{M}(\mathbf{x}^j)}{\text{argmin}} \frac{1}{2} \left\| \mathbf{z}^j - \frac{\mathbf{r}_{(k)}^j}{\hat{\theta}_j} \right\|_2^2, \quad (30)$$

which can be efficiently computed in $O(n)$ time using a skillful implementation of PAVA [11]. If we define for a set \mathcal{A} the projection operator as $P_{\mathcal{A}}(\mathbf{z}) = \underset{\mathbf{x} \in \mathcal{A}}{\text{argmin}} \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2$, the isotonic regression (30) is simply the projection of $\mathbf{r}_{(k)}^j / \hat{\theta}_j$ onto the monotone cone $\mathcal{M}(\mathbf{x}^j)$. Note that $\ell(\cdot)$ is a function of the design \mathbf{Z} instead of the coefficient vector $\hat{\theta}$. Though being convex, the problem (27) can have infinitely many

solutions, some of which can be far from the original $\tilde{\mathbf{X}}$. For example, given any $\tilde{\mathbf{X}}$, we can construct another optimum via shifting two columns $\hat{\mathbf{x}}^i$ and $\hat{\mathbf{x}}^j$ by μ_i and μ_j respectively, such that $\hat{\theta}_i\mu_i + \hat{\theta}_j\mu_j = 0$. To avoid these “bad” solutions, we further impose on each $\hat{\mathbf{x}}^j$ the constraints $\mathbf{1}^T\hat{\mathbf{x}}^j = 0$ and $\|\hat{\mathbf{x}}^j\|_2 \leq \sqrt{n}$, as the marginal distribution of \tilde{x}_{ij} is zero-mean and unit-variance. With additional constraints, the new problem is given by

$$\hat{\mathbf{X}} = \underset{\mathbf{Z} \in \mathbb{R}^{n \times p}}{\operatorname{argmin}} \ell(\mathbf{Z}) \quad (31)$$

s.t. $\mathbf{z}^j \in \mathcal{M}(\mathbf{x}^j)$, $\mathbf{1}^T\mathbf{z}^j = 0$, $\|\mathbf{z}^j\|_2 \leq \sqrt{n}$, $\forall 1 \leq j \leq p$,

and the subproblem for each block boils down to

$$\begin{aligned} \hat{\mathbf{x}}_{(k)}^j &= \underset{\mathbf{z}^j \in \mathcal{M}(\mathbf{x}^j)}{\operatorname{argmin}} \frac{1}{2} \left\| \mathbf{z}^j - \frac{\mathbf{r}_{(k)}^j}{\hat{\theta}_j} \right\|_2^2 \\ \text{s.t. } &\mathbf{1}^T\mathbf{z}^j = 0, \quad \|\mathbf{z}^j\|_2 \leq \sqrt{n}, \end{aligned} \quad (32)$$

which we name *standardized isotonic regression*. The solution to (32) can be viewed as the projection onto the intersection of monotone cone $\mathcal{M}(\mathbf{x}^j)$, hyperplane $\mathcal{L} = \{\mathbf{z} \mid \mathbf{1}^T\mathbf{z} = 0\}$, and scaled L_2 -norm ball $\mathcal{B} = \{\mathbf{z} \mid \|\mathbf{z}\|_2 \leq \sqrt{n}\}$. The next theorem show that the standardized isotonic regression is equivalent to the ordinary isotonic regression followed by successive projection on \mathcal{L} and \mathcal{B} .

Theorem 3 *Given any monotone cone \mathcal{M} , the following equality holds*

$$P_{\mathcal{M} \cap \mathcal{L} \cap \mathcal{B}}(\cdot) = P_{\mathcal{B}}(P_{\mathcal{L}}(P_{\mathcal{M}}(\cdot))), \quad (33)$$

where $P_{\mathcal{L}}(\mathbf{z}) = \mathbf{z} - \frac{\mathbf{1}^T\mathbf{z}}{n} \cdot \mathbf{1}$ and $P_{\mathcal{B}}(\mathbf{z}) = \min\{\frac{\sqrt{n}}{\|\mathbf{z}\|_2}, 1\} \cdot \mathbf{z}$.

The proof of Theorem 3 is given in the supplementary material. Theorem 3 indicates that the extra cost for each subproblem of our CPAV algorithm is very minimal, since the projection onto \mathcal{L} and \mathcal{B} can be done in linear time. Note that the CPAV for AIM needs to work with p blocks of variables, and pre-specifying the monotonicity for each f_j could lead to as many as 2^p different combinations, which is computationally prohibitive. In contrast, our algorithm only deals with roughly $O(s)$ blocks and need not specify the monotonicity. The details of our CPAV is given Algorithm 2. For $\hat{\theta}_j = 0$, the corresponding f_j will have no contribution to the estimated SLIM, which is thus skipped in our CPAV. The convergence of Algorithm 2 basically follows from the extensive studies on cyclic BCD type algorithms [22, 31, 4]. Recently [29] show that the convergence rate of BCD with exact minimization achieves $O(1/t)$ for a family of quadratic nonsmooth problem without linear dependency on the number of blocks, which applies to Algorithm 2 for solving (31).

Algorithm 2 Estimating $\tilde{\mathbf{X}}$

Input: Data $\mathbf{y} \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{n \times p}$, estimated $\hat{\boldsymbol{\theta}}$, number of round t

Output: Estimated hidden design $\hat{\mathbf{X}}$

- 1: Initialize $\hat{\mathbf{X}}_{(0)} = \mathbf{0}_{n \times p}$
 - 2: **for** $k := 1, 2, \dots, t$ **do**
 - 3: **for** $j := 1, 2, \dots, p$ **do**
 - 4: **if** $\hat{\theta}_j \neq 0$ **then**
 - 5: Compute $\mathbf{r}_{(k)}^j$ using (29)
 - 6: Compute $\mathbf{z}_{(k)}^j = P_{\mathcal{M}(\mathbf{x}^j)}\left(\frac{\mathbf{r}_{(k)}^j}{\hat{\theta}_j}\right)$ using PAVA
 - 7: $\hat{\mathbf{x}}_{(k)}^j := P_{\mathcal{B}}(P_{\mathcal{L}}(\mathbf{z}_{(k)}^j))$
 - 8: **end if**
 - 9: **end for**
 - 10: **end for**
 - 11: **Return** $\hat{\mathbf{X}} = \hat{\mathbf{X}}_{(t)}$
-

5 Experimental Results


In this section, we show some experimental evidence for the effectiveness of SLIM. We test our estimation algorithm on the synthetic data. Specifically we fix the problem dimension $p = 500$, the sparsity level of $\tilde{\boldsymbol{\theta}}$, $s = 10$. The distribution of \mathbf{x} is chosen as $NPN(\tilde{\boldsymbol{\Sigma}}, \mathbf{f})$, and $y \sim \langle \tilde{\boldsymbol{\theta}}, \tilde{\mathbf{x}} \rangle + \mathcal{N}(0, 0.25)$. The covariance matrix is given by $\tilde{\boldsymbol{\Sigma}} = \mathbf{A}\mathbf{A}^T$, where \mathbf{A} is a Gaussian random matrix with normalized rows. In data preparation, we first generate $\tilde{\mathbf{x}}$ from $\mathcal{N}(\mathbf{0}, \tilde{\boldsymbol{\Sigma}})$. For the ten \tilde{x}_j 's whose corresponding $\tilde{\theta}_j$'s are nonzero, we then apply ten different monotonically increasing functions to obtain x_j 's, which are basically the inverse of f_j 's. The ten inverse functions are summarized in the table below. The $\Phi(\cdot)$ in f_4^{-1} is the CDF of standard norm distribu-

$f_1^{-1}(x) = x^3$	$f_6^{-1}(x) = x \log(x + 1)$
$f_2^{-1}(x) = \operatorname{sign}(x)\sqrt{ x }$	$f_7^{-1}(x) = 1/(1 + \exp(-x))$
$f_3^{-1}(x) = \exp(x)$	$f_8^{-1}(x) = x - 1$
$f_4^{-1}(x) = \Phi(x)$	$f_9^{-1}(x) = \operatorname{sign}(x) \log(x + 1)$
$f_5^{-1}(x) = x \exp(\sqrt{ x })$	$f_{10}^{-1}(x) = \log(\exp(x) + 1)$

Table 1: Inverse of f_j for nonzero $\tilde{\theta}_j$

tion. For the rest of \tilde{x}_j , we randomly apply one of the functions above. All the results are obtained based on the average over 100 trials.

We plot in Figure 1(a) the normalized estimation error of $\tilde{\boldsymbol{\theta}}$ and $\tilde{\mathbf{X}}$, $\frac{\|\tilde{\boldsymbol{\theta}} - \hat{\boldsymbol{\theta}}\|_2}{\|\tilde{\boldsymbol{\theta}}\|_2}$ and $\frac{\|\tilde{\mathbf{X}} - \hat{\mathbf{X}}\|_2}{\|\tilde{\mathbf{X}}\|_2}$. As sample size n increases from 100 to 500, we can see the clear decreasing trend of error. We also compare the prediction error of SLIM with the simple linear model on 200 new data points, which is shown in Figure 1(b). The best tuning parameters for both methods are picked up via



(a) Estimation error vs. sample size (b) Prediction error vs. sample size (c) Prediction error vs. tuning parameter

Figure 1: Experimental results for SLIM

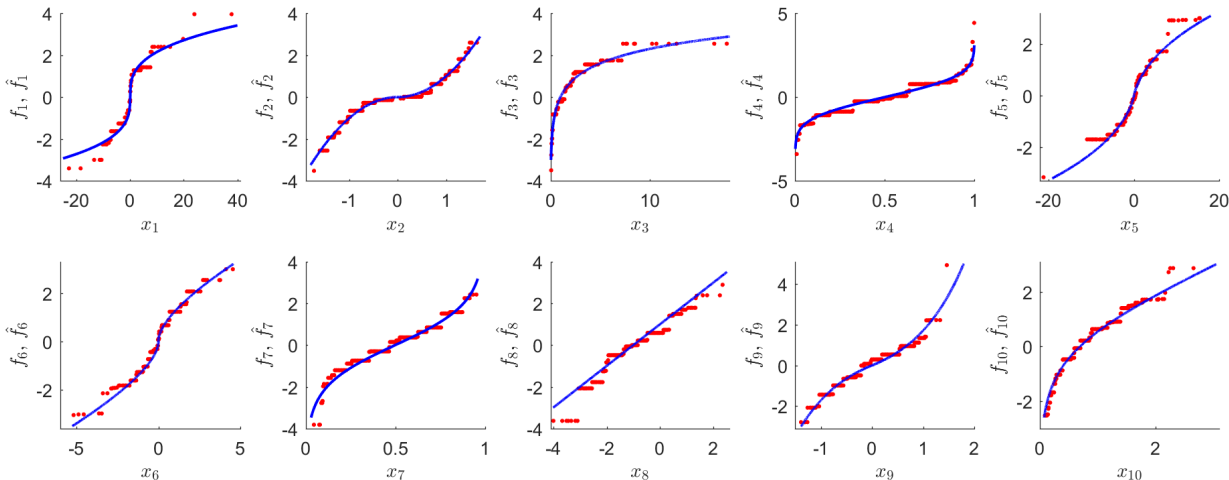


Figure 2: Function f_j (blue curves) and the corresponding estimated \hat{f}_j at observed x_j (red dots) ($n = 500$)

grid search. The simple linear model fails to capture the nonlinear correlation between \mathbf{x} and y , thus incurring large prediction errors. In contrast, SLIM better fits the data and has substantially smaller errors. In Figure 1(c), we specifically plot the prediction errors along the parameter-tuning paths when $n = 500$, and see that SLIM always outperforms the linear model (The actual parameters are different for both methods, but we keep the largest as 2^9 times the smallest). In Figure 2, we also provide the plots for f_1, f_2, \dots, f_{10} and the corresponding estimated ones at the observed x_1, x_2, \dots, x_{10} . It is not difficult to see that the red dots well capture the shape of the function plots except for some tails.

6 Conclusions

In this paper, we propose the sparse linear isotonic models (SLIMs) together with a two-step estimation algorithm. Our model enjoys a few advantages over the classical additive isotonic models. In high dimension, the proposed Kendall’s tau Dantzig selector can provably recover the sparse parameters under suitable statistical assumptions. Especially we can obtain a sharper sample complexity than previous analysis when the

so-called sign sub-Gaussian condition holds. On the optimization side, we show that as the subproblem in our backfitting algorithm for estimating monotone transformations, the standardized isotonic regression can be solved as efficiently as the ordinary isotonic regression. Some empirical evidences also demonstrate the effectiveness of SLIM.

Acknowledgements

The research was supported by NSF grants IIS-1563950, IIS-1447566, IIS-1447574, IIS-1422557, CCF-1451986, CNS- 1314560, IIS-0953274, IIS-1029711, NASA grant NNX12AQ39A, and gifts from Adobe, IBM, and Yahoo.

References

- [1] P. Bacchetti. Additive isotonic model. *Journal of the American Statistical Association*, 84(405):289–294, 1989.
- [2] A. Banerjee, S. Chen, F. Fazayeli, and V. Sivakumar. Estimation with norm regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

- [3] R. Barlow. *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression*. J. Wiley, 1972.
- [4] A. Beck and L. Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013.
- [5] P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [6] E. Candes and T Tao. The Dantzig selector: statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, 2007.
- [7] S. Chatterjee, S. Chen, and A. Banerjee. Generalized dantzig selector: Application to the k -support norm. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [8] Guang Cheng. Semiparametric additive isotonic regression. *J. Stat. Plan. Inference*, 139(6):1980–1991, 2009.
- [9] K. Fang, S. Kotz, and K. Ng. *Symmetric multivariate and related distributions*. Number 36 in Monographs on statistics and applied probability. Chapman & Hall, 1990.
- [10] Z. Fang and N. Meinshausen. LASSO isotone for high-dimensional additive isotonic regression. *Journal of Computational and Graphical Statistics*, 21(1):72–91, 2012.
- [11] S. J. Grotzinger and C. Witzgall. Projections onto order simplexes. *Appl. Math. Optim.*, 12(1):247–270, 1984.
- [12] F. Han and H. Liu. Optimal rates of convergence for latent generalized correlation matrix estimation in transelliptical distribution. *arXiv:1305.6916*, 2013.
- [13] F. Han and H. Liu. Scale-invariant sparse PCA on high-dimensional meta-elliptical data. *Journal of the American Statistical Association*, 109(505):275–287, 2014.
- [14] T. J. Hastie and R. J. Tibshirani. *Generalized additive models*. London: Chapman & Hall, 1990.
- [15] J. L. Horowitz and E. Mammen. Nonparametric estimation of an additive model with a link function. *Ann. Statist.*, 32(6):2412–2443, 2004.
- [16] M.G. Kendall. *Rank correlation methods*. C. Griffin, 1948.
- [17] S. Lee, D. Bryzski, and M. Bogdan. Fast Saddle-Point Algorithm for Generalized Dantzig Selector and FDR Control with the Ordered l_1 -Norm. In *AISTATS*, 2016.
- [18] Y. Lin and H. Zhang. Component selection and smoothing in multivariate nonparametric regression. *Ann. Statist.*, 34(5):2272–2297, 2006.
- [19] H. Liu, F. Han, M. Yuan, J. Lafferty, and L. Wasserman. High-dimensional semiparametric gaussian copula graphical models. *Ann. Statist.*, 40(4):2293–2326, 2012.
- [20] H. Liu, F. Han, and C. H. Zhang. Transelliptical graphical models. In *NIPS*, pages 809–817, 2012.
- [21] H. Liu, J. Lafferty, and L. Wasserman. The non-paranormal: Semiparametric estimation of high dimensional undirected graphs. *J. Mach. Learn. Res.*, 10:2295–2328, 2009.
- [22] Z. Q. Luo and P. Tseng. On the convergence of the coordinate descent method for convex differentiable minimization. *J. Optim. Theory. Appl.*, 72(1):7–35, 1992.
- [23] E. Mammen and K. Yu. Additive isotone regression. *IMS Lecture Notes-Monograph Series Asymptotics: Particles, Processes and Inverse Problems*, 55:179–195, 2007.
- [24] L. Meier, S. Van de Geer, and P. Bhlmann. High-dimensional additive modeling. *Ann. Statist.*, 37(6B):3779–3821, 2009.
- [25] N. Meinshausen and P. Bhlmann. High-dimensional graphs and variable selection with the lasso. *Ann. Statist.*, 34(3):1436–1462, 2006.
- [26] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for the analysis of regularized M -estimators. *Statistical Science*, 27(4):538–557, 2012.
- [27] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted Eigenvalue Properties for Correlated Gaussian Designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- [28] P. Ravikumar, J. Lafferty, H. Liu, and L. Wasserman. Sparse additive models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(5):1009–1030, 2009.
- [29] R. Sun and M. Hong. Improved iteration complexity bounds of cyclic block coordinate descent for convex problems. *arXiv:1512.04680*, 2015.

- [30] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [31] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theory. Appl.*, 109(3):475–494, 2001.
- [32] M. J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming(Lasso). *IEEE Transactions on Information Theory*, 55:2183–2202, 2009.
- [33] L. Xue and H. Zou. Regularized rank-based estimation of high-dimensional nonparanormal graphical models. *Ann. Statist.*, 40(5):2541–2571, 2012.
- [34] M. Yuan. High dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*, 11:2261–2286, 2010.
- [35] S. Zhou. Restricted eigenvalue conditions on subgaussian random matrices. Technical report, Department of Mathematics, ETH Zurich, December 2009.