Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(516): 1726–1735. doi:10.1080/01621459.2015.1113973.

Multitask Quantile Regression under the Transnormal Model

Jianqing Fan, Lingzhou Xue, and Hui Zou

Princeton University, Pennsylvania State University and University of Minnesota

Abstract

We consider estimating multi-task quantile regression under the transnormal model, with focus on high-dimensional setting. We derive a surprisingly simple closed-form solution through rank-based covariance regularization. In particular, we propose the rank-based ℓ_1 penalization with positive definite constraints for estimating sparse covariance matrices, and the rank-based banded Cholesky decomposition regularization for estimating banded precision matrices. By taking advantage of alternating direction method of multipliers, nearest correlation matrix projection is introduced that inherits sampling properties of the unprojected one. Our work combines strengths of quantile regression and rank-based covariance regularization to simultaneously deal with nonlinearity and nonnormality for high-dimensional regression. Furthermore, the proposed method strikes a good balance between robustness and efficiency, achieves the “oracle”-like convergence rate, and provides the provable prediction interval under the high-dimensional setting. The finite-sample performance of the proposed method is also examined. The performance of our proposed rank-based method is demonstrated in a real application to analyze the protein mass spectroscopy data.

Keywords

Copula model; Optimal transformation; Rank correlation; Cholesky decomposition; Quantile regression; Prediction interval; Alternating direction method of multipliers

1 Introduction

Consider a multi-task high-dimensional learning paradigm that independent variables $\mathbf{z} = (z_1, \dots, z_{p_0})'$ are used to simultaneously predict multiple response variables $\mathbf{y} = (y_1, \dots, y_{q_0})'$. Here, dimensions p_0 and q_0 can be of larger order of magnitude than sample size n . This work is motivated by obtaining the optimal prediction $\mathbf{t}(\mathbf{z}) = (t_1(\mathbf{z}), \dots, t_{q_0}(\mathbf{z}))'$ for $\mathbf{y}$. To this end, we shall solve optimal transformations $\mathbf{t}$ from

$$\min_{\mathbf{t}: \mathbb{R}^{p_0} \mapsto \mathbb{R}^{q_0}} \sum_{j=1}^{q_0} E[L(y_j - t_j(\mathbf{z}))],$$

where $L(\cdot)$ is some convex loss function. If $\mathbf{z}$ and $\mathbf{y}$ have a joint normal distribution, nice properties such as linearity and homoscedasticity hold for $\mathbf{t}$. Under normality, it is appropriate to specify $L(\cdot)$ as the squared loss, i.e., $L(u) = u^2$. Thus, normality makes the neat connection between optimal prediction and ordinary least squares, and optimal transformations can be easily obtained by solving ordinary least squares.

However, observed data are often skewed or heavy-tailed, and rarely normally distributed in real-world applications. Transformations are usually used to achieve normality in regression analysis. Under the traditional low-dimensional setting, estimating transformations for regression has received considerable attention. On one hand, parametric methods focus on the parametric families of transformations, such as the celebrated Box-Cox transformation (Box & Cox 1964). On the other hand, nonparametric estimation of regression transformations are also studied, for instance, projection pursuit regression (Friedman & Stuetzle 1981), alternating conditional expectation (Breiman & Friedman 1985), additivity and variance stabilization (Stone 1985, Tibshirani 1988), among others. However, it is nontrivial to extend them to the high-dimensional setting, since these methods would suffer from the curse of dimensionality. Without requiring the restricted normality, there are significant demands for estimating optimal transformations for high-dimensional regression.

As a nice combination of flexibility and interpretability, the Gaussian copulas have generated a lot of interests in statistics and econometrics, and are deemed as a favorable alternative to the Gaussian model in high-dimensional statistical problems, including linear or quadratic discriminant analysis (Lin & Jeon 2003, Mai & Zou 2012, Fan et al. 2013), graphical modeling (Liu et al. 2009, 2012, Xue & Zou 2012), covariance matrix estimation (Xue & Zou 2014, Wegkamp & Zhao 2015), and principal component analysis (Han & Liu 2012). The semiparametric Gaussian copula provides a semiparametric generalization of the Gaussian model by assuming the existence of univariate monotone transformations f for $\mathbf{x} \in \mathbb{R}^p$ such that $f(\mathbf{x}) \sim N_p(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$. Throughout this paper, we follow Lin & Jeon (2003) to call this copula model as the transnormal model, which is also called the nonparanormal model in Liu et al. (2009). We will show the power of the transnormal model in estimating optimal transformations for high-dimensional regression. We suppose that $\mathbf{x} = (\mathbf{z}', \mathbf{y}')'$ consists of $\mathbf{z} = (x_1, \dots, x_{p_0})'$ and $\mathbf{y} = (x_{p_0+1}, \dots, x_{p_0+q_0=p})'$. Denote by F_j the marginal cumulative distribution function (CDF) of x_j and $f_j = \Phi^{-1} \cdot F_j$ for $j = 1, \dots, p_0 + q_0$. Let $\mathbf{g} = (f_1, \dots, f_{p_0})'$ and $\mathbf{h} = (f_{p_0+1}, \dots, f_{p_0+q_0})'$. The transnormal model entails the existence of monotone transformations $\mathbf{f} = (\mathbf{g}', \mathbf{h}')$ such that

$$\mathbf{f}(\mathbf{x}) = \begin{pmatrix} \mathbf{g}(\mathbf{z}) \\ \mathbf{h}(\mathbf{y}) \end{pmatrix} \sim N_p \left(\boldsymbol{\mu}^* = \begin{pmatrix} \boldsymbol{\mu}_z^* \\ \boldsymbol{\mu}_y^* \end{pmatrix}, \boldsymbol{\Sigma}^* = \begin{pmatrix} \boldsymbol{\Sigma}_{zz}^* & \boldsymbol{\Sigma}_{zy}^* \\ \boldsymbol{\Sigma}_{yz}^* & \boldsymbol{\Sigma}_{yy}^* \end{pmatrix} \right) \quad (1)$$

where we may assume that $\boldsymbol{\mu}^* = \mathbf{0}$ and $\boldsymbol{\Sigma}^*$ is a correlation matrix. The transnormal model essentially assumes the joint normality of marginally normal-transformed variables, and it strikes a balance between model robustness and interpretability. This work aims to estimate transformations of predictors $\mathbf{t}(\mathbf{z})$ to optimally predict $\mathbf{y}$ under the transnormal model (1).

Given transformations $\mathbf{g}$ and $\mathbf{h}$, we derive the conditional distribution of $\mathbf{h}(\mathbf{y})$ given $\mathbf{g}(\mathbf{z})$ as

$$\mathbf{h}(\mathbf{y})|\mathbf{g}(\mathbf{z}) \sim N_{q_0}(\boldsymbol{\Sigma}_{yz}^* (\boldsymbol{\Sigma}_{zz}^*)^{-1} \cdot \mathbf{g}(\mathbf{z}), \boldsymbol{\Sigma}_{yy}^* - \boldsymbol{\Sigma}_{yz}^* (\boldsymbol{\Sigma}_{zz}^*)^{-1} \boldsymbol{\Sigma}_{zy}^*). \quad (2)$$

But, since both $\mathbf{g}$ and $\mathbf{h}$ are unknown, it is challenging to obtain the explicit conditional distribution of $\mathbf{y}$ given $\mathbf{z}$. Moreover, linearity and homoscedasticity do not hold for the optimal transformations under the transnormal model. In the presence of nonlinear $\mathbf{g}$ and $\mathbf{h}$, we should take into account the coordinate-wise heterogeneity among the conditional distributions of $\mathbf{y}$ given $\mathbf{z}$. To this end, we employ quantile regression to handle such heterogeneity.

Quantile regression is introduced by the seminal paper of Koenker & Bassett (1978), and it has received much attention in various topics such as survival analysis (Koenker & Geling 2001), time series analysis (Koenker & Xiao 2006), microarray analysis (Wang & He 2007), and variable selection (Zou & Yuan 2008, Bradic et al. 2011). Let $\rho_{\tau}(u) = u(\tau - I_{\{u \leq 0\}})$ be the check loss function. Then we consider the multi-task quantile regression:

$$\min_{\mathbf{t}: \mathbb{R}^{p_0} \rightarrow \mathbb{R}^{q_0}} \sum_{j=1}^{q_0} E[\rho_{\tau}(y_j - t_j(\mathbf{z}))]. \quad (3)$$

When $\tau = \frac{1}{2}$, $\rho_{\tau=\frac{1}{2}}(u) = |u|$, and (3) reduces to multi-task median regression, namely,

$$\min_{\mathbf{t}: \mathbb{R}^{p_0} \rightarrow \mathbb{R}^{q_0}} \sum_{j=1}^{q_0} E[\rho_{\frac{1}{2}}(y_j - t_j(\mathbf{z}))] = \frac{1}{2} \min_{\mathbf{t}: \mathbb{R}^{p_0} \rightarrow \mathbb{R}^{q_0}} E\|\mathbf{y} - \mathbf{t}(\mathbf{z})\|_{\ell_1}.$$

The use of the ℓ_1 loss in prediction was recommended in Friedman (2001, 2002). Our work shares the similar philosophy, and includes multi-task median regression as a special case.

In this paper, we will show that optimal transformations in multi-task quantile regression can be efficiently estimated under the transnormal model. Although estimating optimal transformation is difficult, we derive the surprisingly closed-form optimal prediction without using any smoothing techniques as in Friedman & Stuetzle (1981), Breiman & Friedman (1985), Stone (1985) or Tibshirani (1988). The key ingredient of our proposed method is the positive definite regularized estimation of large covariance matrices under the transnormal model. We propose two novel rank-based covariance regularization methods: the *rank-based positive definite ℓ_1 penalization* for estimating sparse covariance matrices, and the *rank-based banded Cholesky decomposition regularization* for estimating banded inverse covariance matrices. Our proposed rank-based covariance regularization critically depends on a positive definite correlation matrix that retains the desired sampling properties of the adjusted Spearman's or Kendall's rank correlation matrix (Kendall 1948, Liu et al. 2012, Xue & Zou 2012), which can be solved by an efficient alternating direction method of multipliers. By combining both strengths of quantile regression and rank-based covariance regularization, we simultaneously address issues of nonlinearity, non-normality and high dimensionality in estimating optimal transformations for high-dimensional regression. Especially, our proposed method achieves the “oracle”-like convergence rate and provides a provable prediction interval in high dimensions, where dimension is on the nearly exponential order of sample size.

The rest of this paper is organized as follows. We first present the methodological details of optimal prediction in multi-task quantile regression in Section 2. Section 3 establishes the theoretical properties of our proposed method under the transnormal model. Section 4 contains simulation results and a real application to analyze the protein mass spectroscopy data. Technical proofs are presented in the appendices.

2 Multi-task quantile regression

This section presents the optimal prediction in multi-task quantile regression under the transnormal model. The transnormal family of distributions retains the nice interpretation of the normal model, and enables us to make a good use of normal model and theory. In particular, it allows us to obtain the closed-form optimal prediction, which is very appealing to deal with non-normality and high dimensionality. Moreover, the monotone transformation is easy to handle for quantile estimation.

2.1 The closed-form solution

Let $Q_\tau(y_j|z)$ be the τ -th conditional quantile of y_j given z . Then, it immediately implies that $Q_\tau(y_j|z) = \min_{t_j} E[\rho_\tau(y_j - t_j(z))]$. Define $\mathbf{Q}_\tau(\mathbf{y}|z) = (Q_\tau(y_1|z), \dots, Q_\tau(y_p|z))'$. The τ -th equicoordinate conditional quantile $\mathbf{Q}_\tau(\mathbf{y}|z)$ is the analytical solution to (3), i.e.

$$\mathbf{Q}_\tau(\mathbf{y}|z) = \arg \min_{\mathbf{t}: \mathbb{R}^{p\alpha} \mapsto \mathbb{R}^{q\alpha}} \sum_{j=1}^{q\alpha} E[\rho_\tau(y_j - t_j(z))],$$

By using the fact that $\mathbf{h}$ and $\mathbf{g}$ are monotone under the transnormal model, we have

$$\mathbf{Q}_\tau(\mathbf{y}|z) = \mathbf{Q}_\tau(\mathbf{y}|\mathbf{g}(z)) = \mathbf{h}^{-1}(\mathbf{Q}_\tau(\mathbf{h}(\mathbf{y})|\mathbf{g}(z))).$$

Since $\mathbf{h}$ is monotonically nondecreasing, it now follows from (2) that

$$\mathbf{Q}_\tau(\mathbf{y}|z) = \mathbf{h}^{-1} \left(\left(\sum_{yy}^* \left(\sum_{zz}^* \right)^{-1} \cdot \mathbf{g}(z) + \Phi^{-1}(\tau) \text{vdiag} \left(\sum_{yy}^* - \sum_{yz}^* \left(\sum_{zz}^* \right)^{-1} \sum_{zy}^* \right) \right) \right) \quad (4)$$

where $\Phi(\cdot)$ is the CDF of $N(0, 1)$, and $\text{vdiag}(\mathbf{A})$ denotes the vector formed by the diagonal element of $\mathbf{A}$. Therefore, the transnormal model enables us to obtain the closed-form solution to multi-task quantile regression.

Moreover, the closed-form solution (4) can be used to construct prediction intervals for predicting $\mathbf{y}$ given $\mathbf{z}$ in high dimensions. To be more specific, we can obtain the closed-form $100(1 - \tau)\%$ prediction interval as

$$[\mathbf{Q}_{\frac{\tau}{2}}(\mathbf{y}|z), \mathbf{Q}_{1-\frac{\tau}{2}}(\mathbf{y}|z)].$$

For different values of τ , we only need to adjust the value $\Phi(\tau)$. This is also an appealing feature of the transnormal model, especially under the high-dimensional setting.

The solution $\hat{Q}_\tau(y|z)$ uses true covariance matrices and transformations, and it is not a feasible estimator. To utilize (4), we need to estimate the covariance matrix Σ^* and transformation f . It turns out that these two tasks are relatively easy under the transnormal model. Section 2.2 describes how to estimate transformation f , and then Section 3 shows how to estimate the large structured covariance matrix Σ^* . With estimated transformation $\hat{f}=(\hat{g}, \hat{h})$ and covariance $\hat{\Sigma}$, it is straightforward to obtain the following plug-in estimator

$$\hat{Q}_\tau(y|z)=\hat{h}^{-1}\left(\hat{\Sigma}_{yz}(\hat{\Sigma}_{zz})^{-1}\cdot\hat{g}(z)+\Phi^{-1}(\tau)\sqrt{\text{diag}(\hat{\Sigma}_{yy}-\hat{\Sigma}_{yz}(\hat{\Sigma}_{zz})^{-1}\hat{\Sigma}_{zy})}\right) \quad (5)$$

Given the plug-in estimator (5), we further estimate the $100(1-\tau)\%$ prediction interval as

$$[\hat{Q}_{\frac{\tau}{2}}(y|z), \hat{Q}_{1-\frac{\tau}{2}}(y|z)].$$

Remark 1—When $\tau=\frac{1}{2}$, it reduces to optimal prediction in multi-task median regression.

Since $\Phi^{-1}\left(\frac{1}{2}\right)=0$, the optimal prediction can be further simplified as

$$\hat{Q}_{\frac{1}{2}}(y|z)=\hat{h}^{-1}\left(\hat{\Sigma}_{yz}^*(\hat{\Sigma}_{zz}^*)^{-1}\cdot\hat{g}(z)\right) \quad \text{and} \quad \hat{Q}_{\frac{1}{2}}(y|z)=\hat{h}^{-1}\left(\hat{\Sigma}_{yz}^*(\hat{\Sigma}_{zz}^*)^{-1}\cdot\hat{g}(z)\right).$$

Remark 2—Compared to multi-task mean regression, multi-task quantile regression (including median regression) is more robust against outliers in measurements, in addition to delivering the closed-form solution (4). In contrast, by solving ordinary least squares, multi-task mean regression is

$$E(y|z)=E(y|g(z))=E(h^{-1}(h(y))|g(z)).$$

But unlike the quantile regression, this can not be simplified further unless h is linear.

Remark 3—Often $h(\cdot)$ in (4) is nonlinear. The nonlinearity of $h(\cdot)$ makes the difference of conditional quantiles at different values of τ depend on z and thus models the effect of heteroscedasticity. In contrast, Wu et al. (2010), Zhu et al. (2012) and Fan et al. (2013) used the single-index or semiparametric quantile regression to model the heteroscedastic effect, by imposing the model $y=\mu(z'\beta)+\sigma(z'\beta)\varepsilon$, where $\sigma(z'\beta)\varepsilon$ is a heteroscedastic error and ε is independent of z . Unlike the closed-form expression (4), it would be difficult to employ semiparametric quantile regression to handle nonlinearity and high dimensionality.

2.2 Estimation of transformation

Note that $f_j(x_j) \sim \mathcal{N}(0, 1)$ under the transnormal model for any j . Hence, the cumulative distribution function of X_j admits the form

$$F_j(x_j) = \Phi(f_j(x_j)), \text{ or } f_j(x_j) = \Phi^{-1}(F_j(x_j)). \quad (6)$$

Equation (6) provides an intuitive way to estimate the transformation function f_j . Define

$\tilde{F}_j(u) = \frac{1}{n} \sum_{i=1}^n I_{\{x_{ij} \leq u\}}$ as the empirical distribution function. We consider the Winsorized empirical distribution function $\hat{F}_j(\cdot)$ as follows:

$$\tilde{F}_j(u) = \delta_n \cdot I_{\{\tilde{F}_j(u) < \delta_n\}} + \tilde{F}_j(u) \cdot I_{\{\delta_n \leq \tilde{F}_j(u) \leq 1 - \delta_n\}} + (1 - \delta_n) \cdot I_{\{\tilde{F}_j(u) > 1 - \delta_n\}}, \quad (7)$$

where δ_n is the Winsorization parameter to avoid infinity values and to achieve better bias-variance tradeoff. Following Liu et al. (2009) and Mai & Zou (2012), we specify the

Winsorization parameter as $\delta_n = \frac{1}{n^2}$, which facilitates both theoretical analysis and practical performance. Next, we estimate the transformation f in the transnormal model as follows:

$$\hat{f} = (\hat{f}_1, \dots, \hat{f}_p) = (\Phi^{-1} \circ \hat{F}_1, \dots, \Phi^{-1} \circ \hat{F}_p).$$

Remark that $\hat{f}_1, \dots, \hat{f}_p$ are nondecreasing, and they can be substituted into (5) to estimate optimal transformations.

2.3 Estimation of correlation matrix

This section presents two covariance regularization methods based on the *i.i.d.* transnormal data $\mathbf{x}_1, \dots, \mathbf{x}_n$: the *rank-based positive definite ℓ_1 penalization* for estimating sparse correlation matrices, and the *rank-based banded Cholesky decomposition regularization* for estimating banded precision matrices. Both estimates achieve the critical positive definiteness, and they can be used in (5) to obtain the optimal prediction.

2.3.1 Positive definite sparse correlation matrix—Sparse correlation matrices are widely used in many applications where variables are permutation invariant. By truncating small entries to zero, thresholding (Bickel & Levina 2008b, Fan et al. 2013) is a powerful approach to encourage (conditional) sparsity when estimating large correlation matrices. However, the resulting sparse estimator may not be positive definite in practice. Xue et al. (2012) proposed an efficient positive-definite ℓ_1 -penalized estimator to address the indefiniteness issue of thresholding. Now we extend the positive-definite ℓ_1 -penalized estimator to the transnormal model.

First, we introduce the “oracle” positive-definite ℓ_1 -penalized estimator using the “oracle” transformations f , i.e.,

$$\hat{\Sigma}_{\ell_1}^{\circ} = \arg \min_{\Sigma} \frac{1}{2} \|\Sigma - \hat{R}^{\circ}\|_F^2 + \lambda \|\Sigma\|_{1, \text{off}} \text{ subject to } \text{diag}(\Sigma) = \mathbf{1}; \Sigma \succcurlyeq \mathbf{I}$$

where $\hat{\mathbf{R}}^o$ is the sample correlation matrix of “oracle” data $\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_n)$, $\|\mathbf{\Sigma}\|_{1,\text{off}}$ is the ℓ_1 norm of off-diagonal elements in $\mathbf{\Sigma}$, and ε is some arbitrarily small constant (say, $\varepsilon = 10^{-4}$) satisfying that $\lambda_{\min}(\mathbf{\Sigma}^*) \geq \varepsilon$.

Motivated by $\sum_{\ell=1}^o$, we only need to derive a comparable alternative to $\hat{\mathbf{R}}^o$. For space consideration, we focus on the adjusted Spearman’s rank correlation throughout this paper, since the same analysis can be easily adapted to the adjusted Kendall’s rank correlation. Let $\mathbf{s}_j = (s_{1j}, s_{2j}, \dots, s_{nj})'$ be the ranks for $(x_{1j}, x_{2j}, \dots, x_{nj})'$. Denote by $\hat{r}_{jl} = \text{corr}(\mathbf{s}_j, \mathbf{s}_l)$ the

Spearman’s rank correlation, and by $\hat{r}_{jl}^s = 2 \sin\left(\frac{\pi}{6} \hat{r}_{jl}\right)$ the adjusted Spearman’s rank correlation. It is well-known that $\hat{r}_{jl}^s$ corrects the bias of $\hat{r}_{jl}$ (Kendall 1948). Now we

consider the rank correlation matrix $\tilde{\mathbf{R}}^s = (\hat{r}_{jl}^s)_{p \times p}$, which does not require estimating any transformation function. Thus, the rank-based positive definite ℓ_1 penalization is as follows:

$$\hat{\sum}_{\ell_1}^s = \arg \min_{\sum} \frac{1}{2} \|\sum - \tilde{\mathbf{R}}^s\|_F^2 + \lambda \|\sum\|_{1,\text{off}} \quad \text{subject to} \quad \text{diag}(\sum) = \mathbf{1}; \sum \succeq \mathbf{I}. \quad (8)$$

Remark that $\hat{\sum}_{\ell_1}^s$ will guarantee the critical positive definiteness, and it can be efficiently solved by the alternating direction method of multiplier in Xue et al. (2012).

2.3.2 Banded Cholesky decomposition regularization—When an ordering exists among variables, the bandable structure is used in estimating large covariance matrices. With guaranteed positive definiteness, banded Cholesky decomposition regularization receives much attention, for example, Wu & Pourahmadi (2003), Huang et al. (2006), Bickel & Levina (2008a) and Levina et al. (2008).

To motivate our proposal, we first present the “oracle” estimator under the transnormal model. Based on “oracle” data $\mathbf{f}(\mathbf{x}_1), \dots, \mathbf{f}(\mathbf{x}_n)$, the “oracle” banded Cholesky decomposition regularization estimates $\mathbf{\Sigma}^*$ through banding the Cholesky factor of its inverse $\mathbf{\Theta}^*$. Suppose $\mathbf{\Theta}^*$ has the decomposition $\mathbf{\Theta}^* = (\mathbf{I} - \mathbf{A})' \mathbf{D}^{-1} (\mathbf{I} - \mathbf{A})$, where $\mathbf{D}$ is a diagonal matrix and $\mathbf{A} = (a_{jl})_{p \times p}$ is a lower triangular matrix with zero diagonal. Let $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_p)'$. Since $\mathbf{f}(\mathbf{x}) \sim N_p(\mathbf{0}, \mathbf{\Sigma}^*)$, we have $(\mathbf{I} - \mathbf{A})\mathbf{f}(\mathbf{x}) \sim N_p(\mathbf{0}, \mathbf{D})$. As in Bickel & Levina (2008a), the “oracle” estimator is derived by regressing $f(x_j)$ on its closest $\min\{k, j-1\}$ predecessors, i.e. $\hat{a}_1^o = 0$

and for $j = 2, \dots, p$, $\hat{a}_j^o = \arg \min_{\mathbf{a}_j \in \mathcal{A}_j(k)} \frac{1}{n} \sum_{i=1}^n (f_j(x_{ij}) - \mathbf{a}_j' \mathbf{f}(x_i))^2$, where $\mathcal{A}_j(k) = \{(\alpha_1, \dots, \alpha_p)': \alpha_l = 0 \text{ if } l < j - k \text{ or } l \geq j\}$. Then $\mathbf{A}$ is estimated by the k -banded lower triangular matrix, $\hat{\mathbf{A}}^o = (\hat{a}_1^o, \dots, \hat{a}_p^o)'$ and $\mathbf{D}$ is estimated by $\hat{\mathbf{D}}^o = \text{diag}(\hat{d}_1^o, \dots, \hat{d}_p^o)$ with

the residual sum of squares $\hat{d}_j^o = \frac{1}{n} \sum_{i=1}^n (f_j(x_{ij}) - (\hat{\mathbf{a}}_j^o)' \mathbf{f}(x_i))^2$. Hence, the “oracle” estimator ends up with a positive-definite estimator

$\hat{\Sigma}_{chol}^o = (\mathbf{I} - \hat{\mathbf{A}}^o)^{-1} (\hat{\mathbf{D}}^o)^{-1} [(\mathbf{I} - \hat{\mathbf{A}}^o)']^{-1}$, which has the k -banded precision matrix $\hat{\Theta}_{chol}^o = (\mathbf{I} - \hat{\mathbf{A}}^o)' (\hat{\mathbf{D}}^o)^{-1} (\mathbf{I} - \hat{\mathbf{A}}^o)$.

It is important to observe that $\hat{\mathbf{R}}^o = \frac{1}{n} \sum_{i=1}^n \mathbf{f}(x_i) (\mathbf{f}(x_i))'$ plays the central role there. To see this point, we notice that estimating $\hat{a}_j^o$ and $\hat{d}_j^o$ only depends on the quadratic term

$$\frac{1}{n} \sum_{i=1}^n (f_j(x_{ij}) - \mathbf{a}_j' \mathbf{f}(x_i))^2 = \mathbf{a}_j' \hat{\mathbf{R}}^o \mathbf{a}_j - 2 \mathbf{a}_j' \hat{\mathbf{r}}_j^o + \hat{r}_{jj}^o, \quad (9)$$

where $\hat{\mathbf{R}}^o = (\hat{r}_{jl}^o)_{p \times p}$ and $\hat{\mathbf{r}}_j^o = (\hat{r}_{j1}^o, \dots, \hat{r}_{jp}^o)'$ is its j -th row. Thus, we only need a positive definite comparable alternative to $\hat{\mathbf{R}}^o$. The adjusted rank correlation matrix $\hat{\mathbf{R}}^s$ achieves the “oracle”-like rate of convergence, but it is not guaranteed to be positive definite (Devlin et al. 1975). We employ the nearest correlation matrix projection to inherit the desired sampling properties of $\hat{\mathbf{R}}^s$ by solving the following optimization problem:

$$\hat{\mathbf{R}}^s = \arg \min_{\mathbf{R}} \|\mathbf{R} - \tilde{\mathbf{R}}^s\|_{\max} \text{ subject to } \mathbf{R} \succeq \mathbf{I}; \text{diag}(\mathbf{R}) = 1, \quad (10)$$

where $\|\cdot\|_{\max}$ is the entrywise ℓ_∞ norm and $\epsilon > 0$ is some arbitrarily small constant satisfying that $\lambda_{\min}(\Sigma^*) \geq \epsilon$. Zhao et al. (2012) considered a related projection, but used a smooth

surrogate function $\|\mathbf{R} - \tilde{\mathbf{R}}^s\|_{\max}^\nu = \max_{\|U\|_1 \leq 1} \langle U, \mathbf{R} - \tilde{\mathbf{R}}^s \rangle - \frac{\nu}{2} \|\mathbf{R} - \tilde{\mathbf{R}}^s\|_F^2$, which would introduce approximation error. Qi & Sun (2006) solved a nearest correlation matrix projection under the Frobenius norm. We use the entrywise ℓ_∞ -norm for theoretical considerations, as we now demonstrate. Since Σ^* is feasible to (10), we have

$\|\hat{\mathbf{R}}^s - \tilde{\mathbf{R}}^s\|_{\max} \leq \|\Sigma^* - \tilde{\mathbf{R}}^s\|_{\max}$. By using the triangular inequality, $\hat{\mathbf{R}}^s$ retains the desired sampling properties as $\tilde{\mathbf{R}}^s$, that is

$$\|\hat{\mathbf{R}}^s - \Sigma^*\|_{\max} \leq \|\hat{\mathbf{R}}^s - \tilde{\mathbf{R}}^s\|_{\max} + \|\tilde{\mathbf{R}}^s - \Sigma^*\|_{\max} \leq 2 \|\tilde{\mathbf{R}}^s - \Sigma^*\|_{\max}.$$

The details of nearest correlation matrix projection are presented in the appendix.

Now we propose the rank-based estimator on the basis of $\hat{\mathbf{R}}^s$. Let $\hat{\mathbf{R}}^s = (\hat{r}_{jl}^s)_{p \times p}$ and $\hat{\mathbf{r}}_j^s = (\hat{r}_{j1}^s, \dots, \hat{r}_{jp}^s)'$. We substitute $\hat{\mathbf{R}}^s$ into the quadratic term (9) as

$$\mathbf{a}_j' \hat{\mathbf{R}}^s \mathbf{a}_j - 2 \mathbf{a}_j' \hat{\mathbf{r}}_j^s + \hat{r}_{jj}^s = \mathbf{a}_j' \hat{\mathbf{R}}^s \mathbf{a}_j - 2 \mathbf{a}_j' \hat{\mathbf{r}}_j^s + 1.$$

Accordingly, we estimate $\mathbf{A}$ by

$$\hat{\mathbf{A}}^s = (\hat{\mathbf{a}}_1^s, \dots, \hat{\mathbf{a}}_p^s)',$$

with $\hat{\mathbf{a}}_1^s = \mathbf{0}$, and for $j = 2, \dots, p$,

$$\hat{\mathbf{a}}_j^s = \arg \min_{\mathbf{a}_j \in \mathcal{A}_j(k)} \mathbf{a}_j' \hat{\mathbf{R}}^s \mathbf{a}_j - 2\mathbf{a}_j' \hat{\mathbf{r}}_j^s + 1,$$

where $\mathcal{A}_j(k) = \{(\alpha_1 \dots \alpha_p) : \alpha_i = 0 \text{ if } i < j - k \text{ or } i \geq j\}$. In addition, we estimate $\mathbf{D}$ by with

$$\hat{\mathbf{D}}^s = \text{diag}(\hat{d}_1^s, \dots, \hat{d}_p^s),$$

with

$$\hat{d}_j^s = (\hat{\mathbf{a}}_j^s)' \hat{\mathbf{R}}^s \hat{\mathbf{a}}_j^s - 2(\hat{\mathbf{a}}_j^s)' \hat{\mathbf{r}}_j^s + 1.$$

Thus, the rank-based banded Cholesky decomposition regularization yields

$$\hat{\Sigma}_{chol}^s = (\mathbf{I} - \hat{\mathbf{A}}^s)^{-1} \hat{\mathbf{D}}^s (\mathbf{I} - \hat{\mathbf{A}}^s)'^{-1},$$

which has the k -banded precision matrix $\hat{\Theta}_{chol}^s = (\mathbf{I} - \hat{\mathbf{A}}^s)' (\hat{\mathbf{D}}^s)^{-1} (\mathbf{I} - \hat{\mathbf{A}}^s)$.

3 Theoretical properties

This section presents theoretical properties of our proposed methods. We use several matrix norms: the ℓ_1 norm $\|\mathbf{U}\|_{\ell_1} = \max_j \sum_i |u_{ij}|$, the spectral norm $\|\mathbf{U}\|_{\ell_2} = \lambda_{\max}^{1/2}(\mathbf{U}'\mathbf{U})$ and the ℓ_∞ norm $\|\mathbf{U}\|_{\ell_\infty} = \max_i \sum_j |u_{ij}|$. For a symmetric matrix, its matrix ℓ_1 norm coincides with its matrix ℓ_∞ norm. We use c or C to denote constants that do not depend on n or p .

Throughout this section, we follow Bickel & Levina (2008a,b) to assume that

$$\varepsilon_0 \leq \lambda_{\min}(\Sigma^*) \leq \lambda_{\max}(\Sigma^*) \leq \frac{1}{\varepsilon_0}$$

3.1 A general theory

First of all, we consider any regularized estimator $\hat{\Sigma}$ satisfying the following condition:

(C1) There exists a regularized estimator $\hat{\Sigma} = \begin{pmatrix} \hat{\Sigma}^{zz} & \hat{\Sigma}^{zy} \\ \hat{\Sigma}^{yz} & \hat{\Sigma}^{yy} \end{pmatrix}$ satisfying that

$$\|\hat{\Sigma} - \Sigma^*\|_{\ell_1} = O_p(\xi_{n,p}), \text{ with } \xi_{n,p} = o((\log n)^{-1/2}).$$

In the following theorem, we show that the conditional distribution of $\mathbf{h}(\mathbf{y})$ given $\mathbf{g}(\mathbf{z})$ can be well estimated under Condition (C1).

Theorem 1—Assume the data follow the transnormal model. Suppose that there exists $\kappa \in (0, 1)$ such that $n^\kappa \gg \log n \log p_0$. Given Condition (C1), we have the error bounds concerning the estimated conditional distribution of $\mathbf{h}(\mathbf{y})$ given $\mathbf{g}(\mathbf{z})$ in (2) as follows:

i. on the conditional mean vector:

$$\|\sum_{yz} \hat{\sum}_{zz}^{-1} \cdot \hat{\mathbf{g}}(\mathbf{z}) - \sum_{yz}^* (\sum_{zz}^*)^{-1} \cdot \mathbf{g}(\mathbf{z})\|_{\max} = O_P \left(\sqrt{\frac{\log n \log p_0}{n^k}} + \xi_{n,p} \sqrt{\log n} \right).$$

ii. on the conditional variance matrix:

$$\|(\sum_{yy} - \sum_{yz} \hat{\sum}_{zz}^{-1} \hat{\sum}_{zy}) - (\sum_{yy}^* - \sum_{yz}^* (\sum_{zz}^*)^{-1} \sum_{zy}^*)\|_{\ell_2} = O_P(\xi_{n,p}).$$

Next, we show that the plug-in estimator $\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z})$ is asymptotically as good as $\mathbf{Q}_\tau(\mathbf{y}|\mathbf{z})$ under mild regularity conditions. Let $\psi_f(\cdot)$ be the probability density function of X_j . Define $\mathbf{L} = \sum_{yz}^* (\sum_{zz}^*)^{-1} \cdot \mathbf{g}(\mathbf{z}) + \Phi^{-1}(\tau) \text{vdiag}(\sum_{yy}^* - \sum_{yz}^* (\sum_{zz}^*)^{-1} \sum_{zy}^*)$, and let $\mathbf{L} = (L_1, \dots, L_{q_0})'$.

Theorem 2—Under the conditions of Theorem 1, we have the error bound concerning the plug-in estimator in (5) as follows:

$$\|\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z}) - \mathbf{Q}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} = O_P \left(\frac{1}{M} \sqrt{\frac{\log q_0}{n}} + \frac{1}{M} \sqrt{\frac{\log n \log p_0}{n^k}} + \frac{\xi_{n,p}}{M} \sqrt{\log n} \right) \quad (11)$$

where M is the minimum of $\psi_{p_0+j}(x)$

over $x \in \mathcal{J}_j = [-|2L_j|, |2L_j|] \cup [-\int_{p_0+j}^{-1}(|2L_j|), \int_{p_0+j}^{-1}(|2L_j|)]$ for any $j = 1, \dots, q_0$, i.e.,

$$M = \min_{j=1, \dots, q_0} \min_{x \in \mathcal{J}_j} \psi_{p_0+j}(x)$$

Theorem 2 immediately implies that the plug-in estimator $\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z})$ can be used to construct provable prediction intervals in high dimensions. Therefore, we use

$$[\hat{\mathbf{Q}}_{\frac{\tau}{2}}(\mathbf{y}|\mathbf{z}), \hat{\mathbf{Q}}_{1-\frac{\tau}{2}}(\mathbf{y}|\mathbf{z})]$$

to construct the $100(1 - \tau)\%$ prediction interval for predicting $\mathbf{y}$ given $\mathbf{z}$.

In what follows, we show how Condition (C1) is satisfied. Specifically, we consider two parameter spaces for Σ^* :

$$\mathcal{G}_q = \left\{ \sum_j : \max_j \sum_{j \neq i} |a_{ij}|^q \leq s_0 \right\}$$

$$\mathcal{H}_\alpha = \left\{ \sum_j : \max_j \sum_{j < i-k} |a_{ij}| \leq c_0 k^{-\alpha}, \forall k \right\}.$$

These two parameter spaces were previously studied by Bickel & Levina (2008a,b) and Cai & Zhou (2012). We will show that the proposed estimators achieve the “oracle”-like convergence rate over $\mathcal{G}_q$ and $\mathcal{H}_\alpha$ respectively. As a result, we obtain the optimal prediction result.

3.2 Positive definite sparse correlation matrix

First we derive the estimation bound for the rank-based positive-definite ℓ_1 penalization and show its application to estimate the conditional distribution of $\mathbf{h}(\mathbf{y})$ given $\mathbf{g}(\mathbf{z})$.

Theorem 3—Assume the data follow the transnormal model with $\sum^* \in \mathcal{G}_q$, and also assume that $n \gg \log p$. Suppose that $\lambda_{\min}(\sum^*) \geq \varepsilon_0 \gg s_0 (\log p/n)^{\frac{1-q}{2}} + \varepsilon$. With probability tending to 1, the rank-based positive-definite ℓ_1 penalization with $\lambda = c(\log p/n)^{1/2}$ achieves the following upper bound under the matrix ℓ_1 norm,

$$\sup_{\sum^* \in \mathcal{G}_q} \|\hat{\sum}_{\ell_1}^s - \sum^*\|_{\ell_1} \leq C \cdot s_0 \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}}.$$

Remark 5—Both the rank-based estimator $\hat{\sum}_{\ell_1}^s$ and its “oracle” counterpart achieve the minimax optimal convergence rate under the matrix ℓ_1 norm (Cai & Zhou 2012).

Corollary 1—Under the conditions of Theorems 1, 2 and 3, with $\hat{\sum} = \hat{\sum}_{\ell_1}^s$ we have

$$\sup_{\sum^* \in \mathcal{G}_q} \left\| \hat{\sum}_{yz} \hat{\sum}_{zz}^{-1} \cdot \hat{\mathbf{g}}(\mathbf{z}) - \sum_{yz}^* (\sum_{zz}^*)^{-1} \cdot \mathbf{g}(\mathbf{z}) \right\|_{\max} = O_P \left(\sqrt{\frac{\log n \log p_0}{n^k}} + s_0 \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}} \sqrt{\log n} \right),$$

and

$$\sup_{\sum^* \in \mathcal{G}_q} \left\| (\hat{\sum}_{yy} - \hat{\sum}_{yz} \hat{\sum}_{zz}^{-1} \hat{\sum}_{zy}) - (\sum_{yy}^* - \sum_{yz}^* (\sum_{zz}^*)^{-1} \sum_{zy}^*) \right\|_{\ell_2} = O_P \left(s_0 \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}} \right).$$

In light of Theorem 3 and Corollary 1, we can derive the convergence rate for the proposed optimal prediction in the following theorem.

Theorem 4—Under same conditions of Theorems 1, 2 and 3, with $\hat{\Sigma} = \hat{\Sigma}_{\ell_1}^s$ we have

$$\sup_{\sum^* \in \mathcal{G}_q} \|\hat{Q}_\tau(\mathbf{y}|\mathbf{z}) - Q_\tau(\mathbf{y}|\mathbf{z})\|_{\max} = O_P \left(\frac{1}{M} \sqrt{\frac{\log q_0}{n}} + \frac{1}{M} \sqrt{\frac{\log n \log p_0}{n^k}} + \frac{s_0}{M} \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}} \sqrt{\log n} \right).$$

3.3 Banded Cholesky decomposition regularization

Next we derive the estimation bound for the rank-based banded Cholesky decomposition regularization with applications to estimate the conditional distribution of $\mathbf{h}(\mathbf{y})$ given $\mathbf{g}(\mathbf{z})$.

Theorem 5—Assume the data follow the transnormal model with $\sum^* \in \mathcal{H}_\alpha$, and also assume that $n \gg k^2 \log p$. With probability tending to 1, the rank-based banded Cholesky decomposition regularization achieves the following upper bound under the matrix ℓ_1 norm,

$$\sup_{\sum^* \in \mathcal{H}_\alpha} \|\hat{\Sigma}_{chol}^s - \sum^*\|_{\ell_1} \leq Ck \left(\frac{\log p}{n} \right)^{1/2} + Ck^{-\alpha}.$$

If $k = c \cdot \left(\frac{\log p}{n} \right)^{\frac{-1}{2(\alpha+1)}}$ in the rank-based banded Cholesky decomposition regularization, we have

$$\sup_{\sum^* \in \mathcal{H}_\alpha} \|\hat{\Sigma}_{chol}^s - \sum^*\|_{\ell_1} = O_P \left(\left(\frac{\log p}{n} \right)^{\frac{\alpha}{2(\alpha+1)}} \right).$$

Remark 6—Bickel & Levina (2008a) studied the banded Cholesky decomposition regularization under the normal model. Their analysis directly applies to the “oracle” banded Cholesky decomposition regularization. By Theorem 3 of Bickel & Levina (2008a), we have

$$\sup_{\sum^* \in \mathcal{H}_\alpha} \|\hat{\Sigma}_{chol}^o - \sum^*\|_{\ell_1} = O_P \left(\left(\frac{\log p}{n} \right)^{\frac{\alpha}{2(\alpha+1)}} \right).$$

Therefore, the rank-based banded Cholesky decomposition regularization achieves the same convergence rate as the “oracle” counterpart.

Corollary 2—Under the conditions of Theorems 1, 2 and 5, with $\hat{\Sigma} = \hat{\Sigma}_{chol}^s$ we have

$$\sup_{\sum^* \in \mathcal{H}_\alpha} \left\| \hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \cdot \hat{\mathbf{g}}(\mathbf{z}) - \sum_{yz}^* (\sum_{zz}^*)^{-1} \cdot \mathbf{g}(\mathbf{z}) \right\|_{\max} = O_P \left(\sqrt{\frac{\log n \log p_0}{n^k}} + \left(\frac{\log p}{n} \right)^{\frac{\alpha}{2(\alpha+1)}} \sqrt{\log n} \right),$$

and

$$\sup_{\sum^* \in \mathcal{H}_\alpha} \|(\hat{\sum}_{yy} - \hat{\sum}_{yz} \hat{\sum}_{zz}^{-1} \hat{\sum}_{zy}) - (\sum_{yy}^* - \sum_{yz}^* (\sum_{zz}^*)^{-1} \sum_{zy}^*)\|_{\ell_2} = O_P\left(\left(\frac{\log p}{n}\right)^{\frac{\alpha}{2(\alpha+1)}}\right).$$

In light of Theorem 5 and Corollary 2, we can derive the convergence rate for the proposed optimal prediction in the following theorem.

Theorem 6—Under the conditions of Theorems 1, 2 and 5, with $\hat{\sum} = \hat{\sum}_{chol}^s$ we have

$$\sup_{\sum^* \in \mathcal{G}_q} \|\hat{Q}_\tau(y|z) - Q_\tau(y|z)\|_{\max} = O_P\left(\frac{1}{M} \sqrt{\frac{\log q_0}{n}} + \frac{1}{M} \sqrt{\frac{\log n \log p_0}{n^k}} + \frac{\sqrt{\log n}}{M} \left(\frac{\log p}{n}\right)^{\frac{\alpha}{2(\alpha+1)}}\right).$$

4 Numerical properties

This section demonstrates the numerical performance of the proposed methods. For space consideration, we only focus on multi-task median regression and rank-based banded Cholesky decomposition regularization.

4.1 Simulation studies

Simulation studies investigate the “oracle” estimator, the proposed rank-based estimator and the “naïve” estimator. The “oracle” estimator serves as a benchmark in the numerical comparison, and the “naïve” estimator directly regularizes on the covariance of the original data. We summarize notation and details of three regularized estimator in Table 1.

In Models 1–3, we consider three different designs for the inverse covariance matrix Ω^* :

Model 1: $\Omega^* = (\mathbf{I} - \mathbf{A})' \mathbf{D}^{-1} (\mathbf{I} - \mathbf{A})$: $d_{ii} = 0.01$, $a_{i+1,i} = 0.8$ and $d_{ij} = a_{ij} = 0$ otherwise;

Model 2: Ω^* : $\omega_{ii}^* = 1$, $\omega_{i,i\pm 1}^* = 0.5$, $\omega_{i,i\pm 2}^* = 0.25$ and $\omega_{ij}^* = 0$ otherwise;

Model 3: Ω^* : $\omega_{ii}^* = 1$, $\omega_{i,i\pm 1}^* = 0.4$, $\omega_{i,i\pm 2}^* = \omega_{i,i\pm 3}^* = 0.2$, $\omega_{i,i\pm 4}^* = 0.1$ and $\omega_{ij}^* = 0$ otherwise.

Models 1–3 are autoregressive models considered by Huang et al. (2006) and Xue & Zou

(2012). Based on $\Gamma^* = (\Omega^*)^{-1} = (\gamma_{ij}^*)_{p \times p}$, we calculate the true correlation matrix

$\sum^* = (\sigma_{ij}^*)_{p \times p} = \left(\frac{\gamma_{ij}^*}{\sqrt{\gamma_{ii}^* \gamma_{jj}^*}}\right)_{p \times p}$ and also its inverse $\Theta^* = (\Sigma^*)^{-1}$. Then, we generate n independent normal data from $N_p(\mathbf{0}, \Sigma^*)$ and transfer the normal data to the desired transnormal data $(\mathbf{x}_1; \dots; \mathbf{x}_n)$ using

$$[f_1^{-1}, f_2^{-1}, f_3^{-1}, f_4^{-1}, f_1^{-1}, f_2^{-1}, f_3^{-1}, f_4^{-1}, \dots], \quad (12)$$

where $f_1(x) = x^{\frac{1}{3}}$ (power transformation), $f_2(x) = \log\left(\frac{x}{1-x}\right)$ (logit transformation), $f_3(x) = \log(x)$ (log transformation) and $f_4(x) = f_1(x)I_{\{x < -1\}} + f_2(x)I_{\{-1 \leq x \leq 1\}} + (f_3(x-1)+1)I_{\{x > 1\}}$.

In all cases we let $n = 200$ and $p = 100, 200$ & 500 . For each estimator, tuning parameter is chosen by cross-validation. Estimation accuracy is measured by the matrix ℓ_1 -norm and ℓ_2 -norm averaged over 100 independent replications.

Tables 2–3 summarize the numerical performance of estimating the correlation matrix and its inverse for three estimators. We see that the “naïve” estimator performs the worst in the presence of non-normality. The rank-based estimator effectively deals with the transnormal data, and performs comparably with the “oracle” estimator. The numeric evidence is consistent with theoretical results presented in Section 3.

Next, we compare the performance of estimating optimal transformations for multi-task median regression. In each of the 100 replications, we simulate another $4 \times n$ independent transnormal data $(\mathbf{x}_{n+1}, \dots, \mathbf{x}_{5n})$ as the testing data. Now we take the first $p_0 = p/2$ variables in $\mathbf{x}$ as response $\mathbf{y}$ and take the second $q_0 = p/2$ variables in $\mathbf{x}$ as predictor $\mathbf{z}$. Now, given three

estimators $\hat{\Sigma}_{chol}^o$, $\hat{\Sigma}_{chol}^s$ and $\hat{\Sigma}_{chol}^n$, we examine their performance in predicting $\mathbf{y}_i$ based on $\mathbf{z}_i$ for $i = n + 1, \dots, 5n$. Recall that $\hat{\mathbf{t}}_1(\mathbf{z}) = \hat{\mathbf{h}}^{-1}(\hat{\Sigma}_{yz}(\hat{\Sigma}_{zz})^{-1} \cdot \hat{\mathbf{g}}(\mathbf{z}))$ is the closed-form solution for multi-task median regression. In what follows, we consider three optimal prediction methods using $\hat{\Sigma}_{chol}^o$, $\hat{\Sigma}_{chol}^s$ and $\hat{\Sigma}_{chol}^n$ respectively:

- the “oracle” estimator: $\hat{\mathbf{t}}_1^o(\mathbf{z}) = \hat{\mathbf{h}}^{-1}(\hat{\Sigma}_{yz}^o(\hat{\Sigma}_{zz}^o)^{-1} \cdot \mathbf{g}(\mathbf{z}))$
- the rank-based estimator: $\hat{\mathbf{t}}_1^s(\mathbf{z}) = \hat{\mathbf{h}}^{-1}(\hat{\Sigma}_{yz}^s(\hat{\Sigma}_{zz}^s)^{-1} \cdot \hat{\mathbf{g}}(\mathbf{z}))$
- the “naïve” estimator: $\hat{\mathbf{t}}_1^n(\mathbf{z}) = \hat{\boldsymbol{\mu}}_y + \hat{\Sigma}_{yz}^n(\hat{\Sigma}_{zz}^n)^{-1} \cdot (\mathbf{z} - \hat{\boldsymbol{\mu}}_z)$

Tables 4 shows their performances. Prediction accuracy is measured by the difference of prediction errors, which is defined by subtracting the “oracle” prediction error, i.e.

$$\text{DPE}(\hat{\mathbf{t}}_1(\mathbf{z})) = \frac{1}{4n} \sum_{i=n+1}^{5n} \|\hat{\mathbf{t}}_1(\mathbf{z}_i) - \mathbf{y}_i\|_{\ell_1} - \frac{1}{4n} \sum_{i=n+1}^{5n} \|\hat{\mathbf{t}}_1^o(\mathbf{z}_i) - \mathbf{y}_i\|_{\ell_1}.$$

The relative prediction error $\text{DPE}(\hat{\mathbf{t}}_1(\mathbf{z}))$ is averaged over 100 independent replications.

As shown in Table 4, we see that the rank-based estimator performs similarly to its oracle counterpart, and outperforms the “naïve” estimator. The numeric evidence is also consistent with theoretical results presented in Section 3.

4.2 Application to the protein mass spectroscopy data

Recent advances in high-throughput mass spectroscopy technology has enabled biomedical researchers to simultaneously analyze thousands of proteins. This subsection illustrates the power of the proposed rank-based method in an application to study the prostate cancer using the protein mass spectroscopy data (Adam et al. 2002), which was previously analyzed

by Levina et al. (2008). This dataset consists of the protein mass spectroscopy measurements for the blood serum samples of 157 healthy people and 167 prostate cancer patients. In each blood serum sample, the protein mass spectroscopy measures the intensity for the ordered time-of-flight values, which are related to the mass over charge ratio of proteins. In our analysis, we exclude the measurements with mass over charge ratio below 2000 to avoid chemical artifacts, and perform the same preprocessing as in Levina et al. (2008) to smooth the intensity profile. This gives a total of $p = 218$ ordered mass over charge ratio indices for each sample. Then, we have the control data $x_i^{co} = (x_{i,1}^{co}, \dots, x_{i,218}^{co})$ for $i = 1, \dots, 157$, and the cancer data $x_j^{ca} = (x_{j,1}^{ca}, \dots, x_{j,218}^{ca})$ for $j = 1, \dots, 167$. We refer the readers to Adam et al. (2002) and Levina et al. (2008) for more details about this dataset and also the preprocessing procedure. We use the more stable intensity measurements in the latter 168 mass over charge ratio indices to predict the more volatile intensity measurements in the first 50 indices.

The analysis of Levina et al. (2008) is based on the multivariate normal assumption. We perform normality tests on both control and cancer data. Table 5 shows the number of rejecting null hypotheses among 218 normality tests. At least 50% of mass spectroscopy measurements appear to be non-normal. Even after Bonferroni correction, there are at least 30 indices in the control data and 116 indices in the cancer data rejecting all null hypotheses. Figure 1 illustrates non-normality (e.g. heavy tails, skewness) in two indices (109, 218).

The nonnormality of measurements suggests us to try the transnormal model. Let z^{co} (z^{ca}) and y^{co} (y^{ca}) denote the intensity measurements in the latter 168 and first 50 indices of the control (cancer) data. We divide this dataset into training sets $(x_1^{co}, \dots, x_{120}^{co})$, $(x_1^{ca}, \dots, x_{120}^{ca})$ and testing sets $(x_{121}^{co}, \dots, x_{157}^{co})$, $(x_{121}^{ca}, \dots, x_{167}^{ca})$. Note that the sample estimator $\hat{t}_1^{\text{sample}}(z) = \hat{\mu}_y + \sum_{yz}^{\text{sample}} \left(\sum_{zz}^{\text{sample}} \right)^{-1} \cdot (z - \hat{\mu}_z)$ is infeasible since the usual sample correlation matrix $\sum_{zz}^{\text{sample}}$ is not invertible. We consider two different methods to predict y : the proposed rank-based estimator $\hat{t}_1^s(z) = \hat{h}^{-1} \left(\sum_{yz}^s \left(\sum_{zz}^s \right)^{-1} \cdot \hat{g}(z) \right)$ by using the proposed rank-based banded Cholesky decomposition regularization and the “naïve” estimator $\hat{t}_1^n(z) = \hat{\mu}_y + \sum_{yz}^n \left(\sum_{zz}^n \right)^{-1} \cdot (z - \hat{\mu}_z)$ by using the banded Cholesky decomposition regularization (Bickel & Levina 2008a). Following Levina et al. (2008), tuning parameters are chosen via cross-validation. The difference of prediction error at j -th index is computed as

$$\text{DPE}_j(\hat{t}_1(z)) = \text{PE}_j(\hat{t}_1^s(z)) - \text{PE}_j(\hat{t}_1^n(z)).$$

where $\text{PE}_j(\hat{t}_1(z))$ is prediction error at the j -th index computed by averaging the absolute prediction error $|\hat{t}_1(z_i)_j - y_{ij}|$ over $i = 121, \dots, 157$ for the control data and over $i = 121, \dots, 167$ for the cancer data. The differences of prediction errors are shown in Figure 2. As demonstrated in Figure 2, the rank-based method outperforms the “naïve” estimator in 38 out of 50 indices for the control data and 46 out of 50 indices for the cancer data.

Moreover, we further use $[\hat{Q}_{\frac{\tau}{2}}(\mathbf{y}|\mathbf{z}), \hat{Q}_{1-\frac{\tau}{2}}(\mathbf{y}|\mathbf{z})]$ to construct the $100(1 - \tau)\%$ prediction interval. The predict target and prediction intervals are averaged over the testing set. Specifically, we estimate the $100(1 - \tau)\%$ prediction interval as follows,

$$100(1 - \tau)\%PI = [\frac{1}{34} \sum_{i=206}^{239} \hat{Q}_{\frac{\tau}{2}}(\mathbf{y}_i|\mathbf{z}_i), \frac{1}{34} \sum_{i=206}^{239} \hat{Q}_{1-\frac{\tau}{2}}(\mathbf{y}_i|\mathbf{z}_i)].$$

We plot the 95% prediction intervals for both the control data and the cancer data in Figure 3. Figure 3 shows that the estimated prediction intervals cover most of the predict target. This suggests that the transnormal model is a good fit to the data.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors sincerely thank the associate editor and referees for their help comments and suggestions. Jianqing Fan's research is supported in part by R01GM100474-04 and National Science Foundation grants DMS-1206464 and DMS-1406266. Lingzhou Xue's research is supported by the National Institutes of Health grant R01-GM072611-09 and National Science Foundation grant DMS-1505256. Hui Zou's research is supported by NSF grants DMS-0846068 and DMS-1505111.

Appendix A: ADMM for nearest matrix projection

The alternating direction method of multiplier (ADMM) has been widely applied to solving large-scale optimization problems, for example, large covariance matrix estimation (Bien & Tibshirani 2011, Xue et al. 2012) and graphical model selection (Ma et al. 2013, Danaher et al. 2013). Now we design a new ADMM to solve the nearest correlation matrix projection (10). To begin with, we introduce a new variable $\mathbf{S} = \mathbf{R} - \hat{\mathbf{R}}^s$. Then we obtain the equivalent convex optimization problem as

$$(\hat{\mathbf{R}}^s, \hat{\mathbf{S}}) = \arg \min_{(\mathbf{R}, \mathbf{S})} \|\mathbf{S}\|_{\max} \text{ subject to } \mathbf{R} \succeq \varepsilon \mathbf{I}; \mathbf{S} = \mathbf{R} - \hat{\mathbf{R}}^s; \text{diag}(\mathbf{S}) = 0; \mathbf{S} = \mathbf{S}'. \quad (13)$$

Let $\mathbf{\Lambda}$ be the Lagrange multiplier associated with $\mathbf{S} = \mathbf{R} - \hat{\mathbf{R}}^s$ in (13). Next, we introduce the augmented Lagrangian function

$$\mathcal{L}_{\rho}(\mathbf{R}, \mathbf{S}; \mathbf{\Lambda}) = \|\mathbf{S}\|_{\max} - \langle \mathbf{\Lambda}, \mathbf{R} - \mathbf{S} - \hat{\mathbf{R}}^s \rangle + \frac{1}{2\rho} \|\mathbf{R} - \mathbf{S} - \hat{\mathbf{R}}^s\|_F^2, \quad (14)$$

for any given constant $\rho > 0$. For $i = 0, 1, 2, \dots$, we iteratively solve $(\mathbf{R}^{i+1}, \mathbf{S}^{i+1})$ through alternating minimization with respect to $\mathbf{R}$ and $\mathbf{S}$ respectively, i.e.

$$(\mathbf{R}^{i+1}, \mathbf{S}^{i+1}) = \arg \min_{(\mathbf{R}, \mathbf{S})} \mathcal{L}_\rho(\mathbf{R}, \mathbf{S}; \mathbf{\Lambda}^i),$$

and then update the Lagrange multiplier $\mathbf{\Lambda}$ by

$$\mathbf{\Lambda}^{i+1} = \mathbf{\Lambda}^i - \frac{1}{\rho}(\mathbf{R}^{i+1} - \mathbf{S}^{i+1} - \tilde{\mathbf{R}}^s).$$

To sum up, the entire algorithm proceeds sequentially as follows till convergence,

$$\mathbf{R} \text{ step: } \mathbf{R}^{i+1} = \arg \min_{\mathbf{R} \succeq \varepsilon \mathbf{I}} \mathcal{L}_\rho(\mathbf{R}, \mathbf{S}^i; \mathbf{\Lambda}^i); \quad (15)$$

$$\mathbf{S} \text{ step: } \mathbf{S}^{i+1} = \arg \min_{\mathbf{S} = \mathbf{S}': \text{diag}(\mathbf{S})=0} \mathcal{L}_\rho(\mathbf{R}^{i+1}, \mathbf{S}^i; \mathbf{\Lambda}^i); \quad (16)$$

$$\mathbf{\Lambda} \text{ step: } \mathbf{\Lambda}^{i+1} = \mathbf{\Lambda}^i - \frac{1}{\rho}(\mathbf{R}^{i+1} - \mathbf{S}^{i+1} - \tilde{\mathbf{R}}^s). \quad (17)$$

We note that subproblems (15) and (16) indeed have simple solutions, which we will explain in the sequel, and thus we are able to avoid solving a sequence of complex subproblems. We summarize our proposed ADMM as Algorithm 1. The global convergence of Algorithm 1 can be obtained by following Xue et al. (2012). For space consideration, we omit the technical proof in this paper.

Algorithm 1

Proposed alternating direction method of multipliers for obtaining $\hat{\mathbf{R}}^s$.

-
1. Initialization: $\rho, \mathbf{S}^0$ and $\mathbf{\Lambda}^0$.
 2. Iterative alternating direction augmented Lagrangian step: for the i -th iteration
 - 2.1 Solve $\mathbf{R}^{i+1} = (\mathbf{S}^i + \tilde{\mathbf{R}}^s + \rho \mathbf{\Lambda}^i)_+$;
 - 2.2 Solve $\mathbf{S}^{i+1} = ((m_{jl}^{i+1} - t_{jl}^{i+1}) \times I_{\{j \neq l\}})_{p \times p}$ with $\mathbf{M}^{i+1}$ and $\mathbf{T}^{i+1}$ as in (18) & (19);
 - 2.3 Update $\mathbf{\Lambda}^{i+1} = \mathbf{\Lambda}^i - \frac{1}{\rho}(\mathbf{R}^{i+1} - \mathbf{S}^{i+1} - \tilde{\mathbf{R}}^s)$.
 3. Repeat the above cycle till convergence.
-

In what follows, we point out the explicit solutions to (15) and (16). First we consider the $\mathbf{R}$ step. Denote by $(\mathbf{U})_+$ the projection of a matrix $\mathbf{U}$ onto the convex set $\{\mathbf{R} \succeq \varepsilon \mathbf{I}\}$. Suppose that

U has eigen-decomposition $\sum_{i=1}^p \lambda_i \mathbf{u}_i \mathbf{u}_i'$. Then, it is easy to see that $(U)_+ = \sum_{i=1}^p \max(\lambda_i, \varepsilon) \mathbf{u}_i \mathbf{u}_i'$. The subproblem (15) is then solved as follows,

$$\mathbf{R}^{i+1} = \arg \min_{\mathbf{R} \succeq \varepsilon \mathbf{I}} - \langle \mathbf{\Lambda}^i, \mathbf{R} \rangle + \frac{1}{2\rho} \|\mathbf{R} - \mathbf{S}^i - \tilde{\mathbf{R}}^s\|_F^2 = (\mathbf{S}^i + \tilde{\mathbf{R}}^s + \rho \mathbf{\Lambda}^i)_+.$$

We now consider the $\mathbf{S}$ step. Define two subspaces $\mathcal{S} = \{\mathbf{S} : \text{diag}(\mathbf{S}) = 0; \mathbf{S} = \mathbf{S}'\}$ and $\mathcal{T} = \{\mathbf{T} : \|\mathbf{T}\|_1 \leq \rho; \text{diag}(\mathbf{T}) = 0; \mathbf{T} = \mathbf{T}'\}$. We also define

$$\mathbf{M}^{i+1} = \mathbf{R}^{i+1} - \tilde{\mathbf{R}}^s - \rho \mathbf{\Lambda}^i. \quad (18)$$

Note that the subproblem in (16) can be equivalently written as follows,

$$\begin{aligned} \mathbf{S}^{i+1} &= \arg \min_{\mathbf{S} \in \mathcal{S}} \|\mathbf{S}\|_{\max} + \langle \mathbf{\Lambda}^i, \mathbf{S} \rangle + \frac{1}{2\rho} \|\mathbf{R}^{i+1} - \mathbf{S} - \tilde{\mathbf{R}}^s\|_F^2 \\ &= \arg \min_{\mathbf{S} \in \mathcal{S}} \|\mathbf{S}\|_{\max} + \frac{1}{2\rho} \|\mathbf{S} - \mathbf{M}^{i+1}\|_F^2. \end{aligned}$$

Lemma 1 will show that (16) can be solved by the projection onto the entrywise ℓ_1 ball.

Lemma 1

For any given symmetric matrix $\mathbf{M} = (m_{jl})_{p \times p}$, we define

$$\hat{\mathbf{S}} = \arg \min_{\mathbf{S} \in \mathcal{S}} \|\mathbf{S}\|_{\max} + \frac{1}{2\rho} \|\mathbf{S} - \mathbf{M}\|_F^2 \text{ and } \hat{\mathbf{T}} = \arg \min_{\mathbf{T} \in \mathcal{T}} \frac{1}{2} \|\mathbf{T} - \mathbf{M}\|_F^2.$$

Let $\hat{\mathbf{S}} = (\hat{s}_{jl})_{p \times p}$ and $\hat{\mathbf{T}} = (\hat{t}_{jl})_{p \times p}$. Then for any (j, l) , we always have

$$\hat{s}_{jl} = (m_{jl} - \hat{t}_{jl}) \times I_{\{j \neq l\}}.$$

Now we define

$$\mathbf{T}^{i+1} = \arg \min_{\mathbf{T} \in \mathcal{T}} \|\mathbf{T} - \mathbf{M}^{i+1}\|_F^2. \quad (19)$$

We note that $\mathbf{T}^{i+1}$ is essentially a projection of a $\frac{1}{2}p(p-1)$ -dimensional vector onto the ℓ_1 ball in $\mathbb{R}^{\frac{1}{2}p(p-1)}$, and it can be efficiently solved by applying the exact projection algorithm (Duchi et al. 2008) in the $\mathcal{O}(p^2)$ expected time. Let $\mathbf{T}^{i+1} = (t_{jl}^{i+1})_{p \times p}$ and $\mathbf{M}^{i+1} = (m_{jl}^{i+1})_{p \times p}$. Then by Lemma 1, we can obtain the desired closed-form solution for (16), i.e.

$$\mathbf{S}^{i+1} = ((m_{jl}^{i+1} - t_{jl}^{i+1}) \times I_{\{j \neq l\}})_{p \times p}.$$

Appendix B: Technical Proofs

Proof of Lemma 1

Proof

First we note that the dual norm of the entrywise ℓ_1 norm is the entrywise ℓ_∞ norm. Then it is

easy to see that $\|\mathbf{S}\|_{\max} = \frac{1}{\rho} \max_{\mathbf{T} \in \mathcal{T}} \langle \mathbf{T}, \mathbf{S} \rangle$. Now we have

$$\begin{aligned} & \min_{\mathbf{S} \in \mathcal{S}} \|\mathbf{S}\|_{\max} + \frac{1}{2\rho} \|\mathbf{S} - \mathbf{M}\|_F^2 \\ &= \min_{\mathbf{S} \in \mathcal{S}} \max_{\mathbf{T} \in \mathcal{T}} \langle \mathbf{T}, \mathbf{S} \rangle + \frac{1}{2} \|\mathbf{S} - \mathbf{M}\|_F^2 \\ &= \max_{\mathbf{T} \in \mathcal{T}} \min_{\mathbf{S} \in \mathcal{S}} \langle \mathbf{T}, \mathbf{S} \rangle + \frac{1}{2} \|\mathbf{S} - \mathbf{M}\|_F^2 \\ &= \max_{\mathbf{T} \in \mathcal{T}} \min_{\mathbf{S} \in \mathcal{S}} \frac{1}{2} \|\mathbf{S} - \mathbf{M} + \mathbf{T}\|_F^2 - \frac{1}{2} \|\mathbf{T} - \mathbf{M}\|_F^2 + \frac{1}{2} \|\mathbf{M}\|_F^2. \end{aligned}$$

Notice that the optimal solution for the inner problem with respect to $\mathbf{S}$ is given by

$$\hat{\mathbf{S}} = (\hat{s}_{jl})_{p \times p} = ((m_{jl} - t_{jl}) \cdot I_{\{j \neq l\}})_{p \times p}.$$

After we substitute $\hat{\mathbf{S}}$ into the inner problem with respect to $\mathbf{S}$, it suffices to solve the outer problem with respect to $\mathbf{T}$, namely,

$$\max_{\mathbf{T} \in \mathcal{T}} -\frac{1}{2} \|\mathbf{T} - \mathbf{M}\|_F^2 + \frac{1}{2} \|\mathbf{M}\|_F^2 = \min_{\mathbf{T} \in \mathcal{T}} \|\mathbf{T} - \mathbf{M}\|_F^2.$$

Now by definition, $\hat{\mathbf{T}} = (\hat{t}_{jl})_{p \times p}$ gives the optimal solution to the above problem with respect to $\mathbf{T}$. Therefore, it is obvious to see that $\hat{s}_{jl} = (m_{jl} - \hat{t}_{jl}) \times I_{\{j \neq l\}}$ always holds for the optimal solutions $\hat{\mathbf{S}}$ and $\hat{\mathbf{T}}$. This completes the proof of Lemma 1. $\square$

Proof of Theorem 1

Proof

In the first part, we bound $I = \left\| \sum_{yz} \hat{\mathbf{g}}(z) - \sum_{yz}^* (\sum_{zz}^*)^{-1} \mathbf{g}(z) \right\|_{\max}$. To this end, we

bound $I_1 = \left\| \sum_{yz} \hat{\mathbf{g}}(z) - \sum_{yz}^* (\sum_{zz}^*)^{-1} \right\|_{\ell_\infty}$ and

$I_2 = \left\| \hat{\mathbf{g}}(z) - \mathbf{g}(z) \right\|_{\max} = O_P \left(\sqrt{\frac{\log n \log p_0}{n^\kappa}} \right)$ respectively. Let $\varphi_1 = \left\| (\sum_{zz}^*)^{-1} \right\|_{\ell_\infty}$ and

$\varphi_2 = \left\| \sum_{yz}^* \right\|_{\ell_\infty}$ for ease of notation. Notice that

$$\begin{aligned} \hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} - 1 &= \Sigma_{yz}^* \cdot (\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}) + (\hat{\Sigma}_{yz} - \Sigma_{yz}^*) \cdot (\Sigma_{zz}^*)^{-1} - 1 \\ &\quad + (\hat{\Sigma}_{yz} - \Sigma_{yz}^*) \cdot (\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}). \end{aligned}$$

By using the triangular inequality again, we then have

$$\begin{aligned} I_1 &\leq \|\Sigma_{yz}^*\|_{\ell_\infty} \cdot \|\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}\|_{\ell_\infty} + \|(\Sigma_{zz}^*)^{-1}\|_{\ell_\infty} \cdot \|\hat{\Sigma}_{yz} - \Sigma_{yz}^*\|_{\ell_\infty} \\ &\quad + \|\hat{\Sigma}_{yz} - \Sigma_{yz}^*\|_{\ell_\infty} \cdot \|\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}\|_{\ell_\infty} \end{aligned}$$

To bound I_1 , we need to bound $\|\hat{\Sigma}_{yz} - \Sigma_{yz}^*\|_{\ell_\infty}$ and $\|\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}\|_{\ell_\infty}$. By definition, it is easy to see that $\|\hat{\Sigma}_{yz} - \Sigma_{yz}^*\|_{\ell_\infty} \leq \|\hat{\Sigma} - \Sigma^*\|_{\ell_\infty}$. Moreover, we have

$$\|\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1}\|_{\ell_\infty} \leq C \|\hat{\Sigma}_{zz} - \Sigma_{zz}^*\|_{\ell_\infty} \cdot \|(\Sigma_{zz}^*)^{-1}\|_{\ell_\infty} \leq C \varphi_1 \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_\infty}.$$

Thus we can derive the explicit upper bound for I_1 as

$$I_1 \leq C(\varphi_1 \varphi_2 + \varphi_1) \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_\infty} = C(\varphi_1 \varphi_2 + \varphi_1) \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_1} = O_P(\xi_{n,p}). \quad (20)$$

Next, we prove that

$$I_2 = \|\hat{g}(z) - g(z)\|_{\max} = O_P\left(\sqrt{\frac{\log n \log p_0}{n^\kappa}}\right). \quad (21)$$

To this end, note that $g(z_j)$ follows the standard normal distribution. Hence we have

$$\begin{aligned} Pr(|g_j(z_j)| > \sqrt{(1-\kappa)\log n}) &= \int_{\sqrt{(1-\kappa)\log n}}^{+\infty} C \cdot \exp(-t^2/2) \cdot dt \\ &\leq \frac{C}{\sqrt{\log n}} \cdot \int_{\sqrt{(1-\kappa)\log n}}^{+\infty} \exp(-t^2/2) \cdot t dt \\ &\leq C \cdot \frac{1}{n^{\frac{1-\kappa}{2}} \sqrt{\log n}}. \end{aligned}$$

Then it immediately implies that

$$\|g(z)\|_{\max} = O_P(\sqrt{\log n}). \quad (22)$$

Let $\eta = \sqrt{(1-\kappa)\log n}$. Now by using Lemma 1 of Mai & Zou (2012), we have

$$Pr(\sup_{|g_j(z_j)| \leq \eta} |\hat{g}_j(z_j) - g_j(z_j)| \geq \eta) \leq C \exp(-\frac{cn^{\frac{1}{3}}\eta^2}{\log n}) + p_0 \exp(-\frac{cn^{\frac{1}{3}}}{\log n})$$

Thus, we apply the union bound to obtain

$$\begin{aligned} Pr(\|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_{\max} \geq \eta) &\leq \sum_{j=1}^{p_0} Pr(|g_j(z_j)| > \eta) + \sum_{j=1}^{p_0} Pr(\sup_{|g_j(z_j)| \leq \eta} |\hat{g}_j(z_j) - g_j(z_j)| \geq \eta) \\ &\leq C \cdot \frac{p_0}{n^{\frac{1-\kappa}{2}} \sqrt{\log n}} + C p_0 \exp(-\frac{cn^{\frac{1}{3}}\eta^2}{\log n}) + p_0 \exp(-\frac{cn^{\frac{1}{3}}}{\log n}), \end{aligned}$$

which immediately implies the upper bound (21). Now since

$$\begin{aligned} &\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \hat{\mathbf{g}}(\mathbf{z}) - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \mathbf{g}(\mathbf{z}) \\ &= (\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}) \cdot (\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})) + \\ &(\Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}) \cdot (\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})) + (\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}) \cdot \mathbf{g}(\mathbf{z}), \end{aligned}$$

we can combine (20), (21), (22) and also the triangle inequality to bound I as

$$\begin{aligned} I &\leq \|\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}\|_{\ell_{\infty}} \cdot \|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_{\max} \\ &\quad + \|\Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}\|_{\ell_{\infty}} \cdot \|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_{\max} \\ &\quad + \|\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1}\|_{\ell_{\infty}} \cdot \|\mathbf{g}(\mathbf{z})\|_{\max} \\ &= O_P(\sqrt{\frac{\log n \log p_0}{n^{\kappa}}} + \xi_{n,p} \sqrt{\log n}). \end{aligned}$$

In the second part, we bound

$$J = \|(\hat{\Sigma}_{yy} - \hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \hat{\Sigma}_{zy}) - (\Sigma_{yy}^* - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \Sigma_{zy}^*)\|_{\ell_2}. \text{ By using the triangular inequality, we have}$$

$$J \leq \|\hat{\Sigma}_{yy} - \Sigma_{yy}^*\|_{\ell_2} + \|\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \hat{\Sigma}_{zy} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \Sigma_{zy}^*\|_{\ell_2} \equiv J_1 + J_2.$$

Notice that

$$J_2 = \hat{\Sigma}_{yz} (\hat{\Sigma}_{zz}^*)^{-1} (\hat{\Sigma}_{zy} - \Sigma_{zy}^*) + (\hat{\Sigma}_{yz} - \Sigma_{yz}^*) (\hat{\Sigma}_{zz}^*)^{-1} \Sigma_{zy}^* + \hat{\Sigma}_{yz} (\hat{\Sigma}_{zz}^{-1} - (\hat{\Sigma}_{zz}^*)^{-1}) \hat{\Sigma}_{zy}.$$

By using the definition of spectral norm, we have

$$\begin{aligned}
\|\hat{\Sigma} - \Sigma^*\|_{\ell_2} &= \max_{u=(u'_z, u'_y)': \|u\|_{\ell_2}=1} \left\| \begin{pmatrix} \hat{\Sigma}_{zz} - \Sigma_{zz}^* & \hat{\Sigma}_{zy} - \Sigma_{zy}^* \\ \hat{\Sigma}_{yz} - \Sigma_{yz}^* & \hat{\Sigma}_{yy} - \Sigma_{yy}^* \end{pmatrix} \begin{pmatrix} u_z \\ u_y \end{pmatrix} \right\|_{\ell_2} \\
&\geq \max_{u_z: \|u_z\|_{\ell_2}=1} \left\| \begin{pmatrix} \hat{\Sigma}_{zz} - \Sigma_{zz}^* & \hat{\Sigma}_{zy} - \Sigma_{zy}^* \\ \hat{\Sigma}_{yz} - \Sigma_{yz}^* & \hat{\Sigma}_{yy} - \Sigma_{yy}^* \end{pmatrix} \begin{pmatrix} u_z \\ 0 \end{pmatrix} \right\|_{\ell_2} \\
&\geq \max_{u_z: \|u_z\|_{\ell_2}=1} \|(\hat{\Sigma}_{yz} - \Sigma_{yz}^*)u_z\|_{\ell_2} \\
&= \|\hat{\Sigma}_{yz} - \Sigma_{yz}^*\|_{\ell_2}.
\end{aligned}$$

By similar arguments, we also have $\|\hat{\Sigma} - \Sigma^*\|_{\ell_2} \geq \|\hat{\Sigma}_{zy} - \Sigma_{zy}^*\|_{\ell_2}$. Given that

$\varepsilon_0 \leq \lambda_{\min}(\Sigma^*) \leq \lambda_{\max}(\Sigma^*) \leq \frac{1}{\varepsilon_0}$, we then apply the sub-multiplicative property to obtain

$$\begin{aligned}
\|\hat{\Sigma}_{yz}(\Sigma_{zz}^*)^{-1}(\hat{\Sigma}_{zy} - \Sigma_{zy}^*)\|_{\ell_2} &\leq C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_2} \\
\|(\hat{\Sigma}_{yz} - \Sigma_{yz}^*)(\Sigma_{zz}^*)^{-1}\Sigma_{zy}^*\|_{\ell_2} &\leq C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_2} \\
\|\hat{\Sigma}_{yz}(\hat{\Sigma}_{zz}^{-1} - (\Sigma_{zz}^*)^{-1})\hat{\Sigma}_{zy}\|_{\ell_2} &\leq C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_2}
\end{aligned}$$

Thus we can combine the above bounds to bound J_2 as

$$J_2 = \|\hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \hat{\Sigma}_{zy} - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \Sigma_{zy}^*\|_{\ell_2} \leq C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_2}$$

In addition, due to the fact that $J_1 = \|\hat{\Sigma}_{yy} - \Sigma_{yy}^*\|_{\ell_2} \leq \|\hat{\Sigma} - \Sigma^*\|_{\ell_2}$, we have

$$J \leq J_1 + J_2 \leq C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_2} = C \cdot \|\hat{\Sigma} - \Sigma^*\|_{\ell_1} = O_P(\xi_{n,p})$$

This completes the proof of Theorem 1. $\square$

Proof of Theorem 2

Proof

To simplify notation, we define $\mu_{y|z} = \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \cdot g(z)$,

$\Sigma_{y|z} = \Sigma_{yy}^* - \Sigma_{yz}^* (\Sigma_{zz}^*)^{-1} \Sigma_{zy}^*$, $\hat{\mu}_{y|z} = \hat{\Sigma}_{yz} (\hat{\Sigma}_{zz})^{-1} \cdot \hat{g}(z)$ and

$\hat{\Sigma}_{y|z} = \hat{\Sigma}_{yy} - \hat{\Sigma}_{yz} \hat{\Sigma}_{zz}^{-1} \hat{\Sigma}_{zy}$. Then $L = \mu_{y|z}^* + \Phi^{-1}(\tau) \text{vdiag}(\Sigma_{y|z}^*)$. Define

$\hat{L} = \hat{\mu}_{y|z} (\hat{\Sigma}_{zz})^{-1} \cdot \hat{g}(z) + \Phi^{-1}(\tau) \text{vdiag}(\hat{\Sigma}_{yy} - \hat{\Sigma}_{yz} (\hat{\Sigma}_{zz})^{-1} \hat{\Sigma}_{zy})$, and then we have

$\hat{L} = \hat{\mu}_{y|z} + \Phi^{-1}(\tau) \text{vdiag}(\hat{\Sigma}_{y|z})$. In addition, we introduce $\tilde{Q}_\tau(y|z) = h^{-1}(\hat{L})$.

In order to bound $\hat{Q}_\tau(\mathbf{y}|\mathbf{z}) = Q_\tau(\mathbf{y}|\mathbf{z})$, we use the triangle inequality to obtain that

$$\begin{aligned}\|\hat{Q}_\tau(\mathbf{y}|\mathbf{z}) - Q_\tau(\mathbf{y}|\mathbf{z})\|_{\max} &\leq \|\hat{Q}_\tau(\mathbf{y}|\mathbf{z}) - \tilde{Q}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} + \|\tilde{Q}_\tau(\mathbf{y}|\mathbf{z}) - Q_\tau(\mathbf{y}|\mathbf{z})\|_{\max} \\ &\leq \|\hat{\mathbf{h}}^{-1}(\hat{\mathbf{L}}) - \mathbf{h}^{-1}(\hat{\mathbf{L}})\|_{\max} + \|\mathbf{h}^{-1}(\hat{\mathbf{L}}) - \mathbf{h}^{-1}(\mathbf{L})\|_{\max} \\ &\equiv I + J.\end{aligned}$$

In the sequel, we derive probability bounds for I and J respectively.

To bound I , we first use the triangle inequality again to bound

$$\begin{aligned}\|\hat{\mathbf{L}} - \mathbf{L}\|_{\max} &\leq \|\hat{\boldsymbol{\mu}}_{\mathbf{y}|\mathbf{z}} - \boldsymbol{\mu}_{\mathbf{y}|\mathbf{z}}^*\|_{\max} + \Phi^{-1}(\tau) \cdot \|\text{vdiag}(\hat{\Sigma}_{\mathbf{y}|\mathbf{z}}) - \text{vdiag}(\Sigma_{\mathbf{y}|\mathbf{z}}^*)\|_{\max} \\ &\leq \|\hat{\boldsymbol{\mu}}_{\mathbf{y}|\mathbf{z}} - \boldsymbol{\mu}_{\mathbf{y}|\mathbf{z}}^*\|_{\max} + \Phi^{-1}(\tau) \cdot \|\hat{\Sigma}_{\mathbf{y}|\mathbf{z}} - \Sigma_{\mathbf{y}|\mathbf{z}}^*\|_{\ell_2} \\ &\leq O_P\left(\sqrt{\frac{\log n \log p_0}{n^\kappa}} + \xi_{n,p} \sqrt{\log n}\right) + \Phi^{-1}(\tau) \cdot O_P(\xi_{n,p}) \\ &= O_P\left(\sqrt{\frac{\log n \log p_0}{n^\kappa}} + \xi_{n,p} \sqrt{\log n}\right)\end{aligned}$$

Notice that $\sqrt{n^{-\kappa} \log n \log p_0} + \xi_{n,p} \sqrt{\log n} = o(1) \leq |\mathbf{L}|$. Thus, with probability 1, we have

$$|\hat{\mathbf{L}}| \leq 2 \cdot |\mathbf{L}|, \text{ and then, } \Phi(-2 \cdot |\mathbf{L}|) \leq \Phi(\hat{\mathbf{L}}) \leq \Phi(2 \cdot |\mathbf{L}|). \quad (23)$$

In view of (23), we use the probability inequality from Page 75 of Serfling (1980) to obtain

$$Pr(|\hat{F}_{p_0+j}^{-1}(t) - F_{p_0+j}^{-1}(t)| > \varepsilon) \leq 2 \exp(-2n\delta_\varepsilon^2)$$

where $\delta_\varepsilon = \min_{t \in [-2L_j], [2L_j]} \{F_j(t+\varepsilon) - F_j(t), F_j(t) - F_j(t-\varepsilon)\}$, for any $j = 1, \dots, q_0$ and $t \in [\Phi(-2|L_j|), \Phi(2|L_j|)] \subset \mathcal{J}_j$. Recall that $\hat{\mathbf{f}} = (\Phi^{-1} \circ \hat{F}_1, \dots, \Phi^{-1} \circ \hat{F}_{p_0}) = (\hat{\mathbf{g}}, \hat{\mathbf{h}})$. Note that $\hat{h}_j^{-1}(\hat{L}_j) - h_j^{-1}(\hat{L}_j) = \hat{F}_j^{-1}(\Phi(\hat{L}_j)) - F_j^{-1}(\Phi(\hat{L}_j))$ is the j -th element of $\hat{\mathbf{h}}^{-1}(\hat{\mathbf{L}}) - \mathbf{h}^{-1}(\hat{\mathbf{L}})$ for $j = 1, \dots, q_0$. Now we use the union bound to derive

$$\begin{aligned}Pr(\|\hat{Q}_\tau(\mathbf{y}|\mathbf{z}) - \tilde{Q}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} > \varepsilon) &\leq \sum_{j=1}^{q_0} Pr(|\hat{h}_j^{-1}(\hat{L}_j) - h_j^{-1}(\hat{L}_j)| > \varepsilon) \\ &\leq \sum_{j=1}^{q_0} Pr(|\hat{F}_{p_0+j}^{-1}(\Phi(\hat{L}_j)) - F_{p_0+j}^{-1}(\Phi(\hat{L}_j))| > \varepsilon) \\ &\leq \sum_{j=1}^{q_0} \max_{t \in [\Phi(-2|L_j|), \Phi(2|L_j|)]} Pr(|\hat{F}_{p_0+j}^{-1}(t) - F_{p_0+j}^{-1}(t)| > \varepsilon) \\ &\leq q_0 \cdot \exp(-2nM^2\varepsilon^2)\end{aligned}$$

where we use the fact that $\delta_\varepsilon \geq \min_{t \in \mathcal{J}_j} \psi_j(t) \varepsilon \geq M\varepsilon$ in the last inequality. Thus, we have

$$\|\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z}) - \tilde{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} = O_P\left(\frac{1}{M} \sqrt{\frac{\log q_0}{n}}\right)$$

To bound J , we apply the mean-value theorem to the j -th element of $\mathbf{h}^{-1}(\hat{\mathbf{L}}) - \mathbf{h}^{-1}(\mathbf{L})$, that is, $h_j^{-1}(\hat{L}_j) - h_j^{-1}(L_j) = F_j^{-1}(\Phi(\hat{L}_j)) - F_j^{-1}(\Phi(L_j))$. Namely,

$$h_j^{-1}(\hat{L}_j) - h_j^{-1}(L_j) = (h_j^{-1})'(L_j) \cdot (\hat{L}_j - L_j) = \frac{\phi(\tilde{L}_j)}{\psi_{p_0+j}(F_{p_0+j}^{-1}(\Phi(\tilde{L}_j)))} \cdot (\hat{L}_j - L_j)$$

where $\tilde{L}_j$ is no the line segment between $\hat{L}_j$ and L_j , and in the second equality we use the fact that $(h_j^{-1})'(t) = (F_{p_0+j}^{-1} \circ \Phi)'(t) = \phi(t) / \psi_{p_0+j}(F_{p_0+j}^{-1}(\Phi(t)))$. With probability 1, we have

$$\|\tilde{\mathbf{L}}\| \leq \|\tilde{\mathbf{L}} - \mathbf{L}\| + \|\mathbf{L}\| \leq \|\hat{\mathbf{L}} - \mathbf{L}\| + \|\mathbf{L}\| \leq 2\|\mathbf{L}\|.$$

Then, $\phi(\tilde{L}_j) / \psi_{p_0+j}(F_{p_0+j}^{-1}(\Phi(\tilde{L}_j))) \leq \phi(\tilde{L}_j) / (\min_{x \in \mathcal{J}_j} \psi_{p_0+j}(x)) \leq (M \sqrt{2\pi})^{-1}$. Thus we have

$$\begin{aligned} \|\tilde{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z}) - \mathbf{Q}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} &= \|\mathbf{h}^{-1}(\hat{\mathbf{L}}) - \mathbf{h}^{-1}(\mathbf{L})\|_{\max} \\ &\leq \max_{j=1, \dots, q_0} \frac{\phi(\tilde{L}_j)}{\psi_{p_0+j}(F_{p_0+j}^{-1}(\Phi(\tilde{L}_j)))} \cdot \|\hat{\mathbf{L}} - \mathbf{L}\|_{\max} \\ &= O_P\left(\frac{1}{M} \sqrt{\frac{\log n \log p_0}{n^\kappa}} + \frac{\xi_{n,p}}{M} \sqrt{\log n}\right) \end{aligned}$$

Therefore, we obtain the desired bound for $\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z}) - \mathbf{Q}_\tau(\mathbf{y}|\mathbf{z})$ as follows,

$$\|\hat{\mathbf{Q}}_\tau(\mathbf{y}|\mathbf{z}) - \mathbf{Q}_\tau(\mathbf{y}|\mathbf{z})\|_{\max} \leq I + J = O_P\left(\frac{1}{M} \sqrt{\frac{\log q_0}{n}} + \frac{1}{M} \sqrt{\frac{\log n \log p_0}{n^\kappa}} + \frac{\xi_{n,p}}{M} \sqrt{\log n}\right).$$

This completes the proof of Theorem 2. $\square$

Proof of Theorem 3

Proof

Define the rank-based soft-thresholding estimator

$$\hat{\mathbf{R}}_{\text{soft}}^s = \left(s\lambda(\hat{r}_{jl}^s) \right)_{p \times p}$$

where $s\lambda(\cdot)$ applies the soft-thresholding rule for the off-diagonal elements of $\hat{\mathbf{R}}^s$. Under the event $\{\|\tilde{\mathbf{R}}^s - \sum^*_{\max} \leq c \cdot (n^{-1} \log p)^{1/2}\}$, we can follow the same line of proof as in Bickel & Levina (2008b) to show that

$$\sup_{\sum^* \in \mathcal{G}_q} \|\hat{\sum}_{\text{soft}}^s - \sum^*\|_{\ell_1} \leq C \cdot s_0 \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}}.$$

Given that $\lambda_{\min}(\sum^*) \geq \varepsilon_0 \gg s_0(\log p/n)^{\frac{1-q}{2}} + \varepsilon$, it is easy to see that $\lambda_{\min}(\hat{\sum}_{\text{soft}}^s) \geq \varepsilon$.

Thus, $\hat{\sum}_{\text{soft}}^s$ is the unique solution to the convex rank-based positive-definite ℓ_1 penalization under the same event, which immediately implies that

$$\sup_{\sum^* \in \mathcal{G}_q} \|\hat{\sum}_{\ell_1}^s - \sum^*\|_{\ell_1} \leq C \cdot s_0 \left(\frac{\log p}{n} \right)^{\frac{1-q}{2}}.$$

We cite the entry-wise estimation bound of $\tilde{\mathbf{R}}^s$ from Xue & Zou (2012) that

$$Pr(\|\tilde{\mathbf{R}}^s - \sum^*\|_{\max} \leq c \cdot (n^{-1} \log p)^{1/2} p) \rightarrow 1,$$

as n tends to ∞ . This completes the proof of Theorem 2. $\square$

Proof of Theorem 5

Proof

As demonstrated in Bickel & Levina (2008a), it is sufficient to bound $\|\hat{\Theta}_{chol}^s - \Theta^*\|_{\ell_1}$. Denote by $B_k(\mathbf{U}) = (u_{jl} I_{\{|j|-k\}})_{p \times p}$ the k -banded estimator of $\mathbf{U} = (u_{jl})_{p \times p}$. Let $\mathbf{A}_k = B_k(\mathbf{A})$ be the k -banded estimator of $\mathbf{A}$, and $\mathbf{D}_k$ be the corresponding diagonal matrix. Define $\Theta_k = (\mathbf{I} - \mathbf{A}_k)' \mathbf{D}_k^{-1} (\mathbf{I} - \mathbf{A}_k)$. Denote by $B_k(\Theta^*)$ the k -banded estimator of Θ^* .

We first apply the triangle inequality to obtain

$$\|\hat{\Theta}_{chol}^s - \Theta^*\|_{\ell_1} \leq \|\hat{\Theta}_{chol}^s - \Theta_k\|_{\ell_1} + \|\Theta_k - B_k(\Theta^*)\|_{\ell_1} + \|B_k(\Theta^*) - \Theta^*\|_{\ell_1}. \quad (24)$$

Notice that $\|\Theta_k - B_k(\Theta^*)\|_{\ell_1}$ in (24) can be bounded as in Bickel & Levina (2008a), i.e.

$$\|\Theta_k - B_k(\Theta^*)\|_{\ell_1} \leq Ck^{-\alpha}.$$

Also by the definition of $\mathcal{H}_\alpha$, it is easy to see that $\|B_k(\Theta^*) - \Theta^*\|_{\ell_1} \leq Ck^{-\alpha}$ holds. Then it is sufficient for us to bound $\|\hat{\Theta}_{chol}^s - \Theta_k\|_{\ell_2}$ only.

Under the event $\{\|\hat{\mathbf{R}}^s - \sum^*\|_{\max} \leq \zeta\}$ with $k\zeta = \alpha(1)$, we follow (A14) – (A15) in Bickel & Levina (2008a) to show that $\|\hat{\mathbf{A}}^s - \mathbf{A}_k\|_{\max} \leq C\zeta$, and meanwhile, follow (A16) – (A17)

in Bickel & Levina (2008a) to show that $\|\hat{D}^s - D_k\|_{\max} \leq Ck^2\zeta^2$. Then by using triangular inequality, we obtain the upper bound for $\|\hat{\Theta}_{chol}^s - \Theta_k\|_{\ell_1}$ under the same event as follows,

$$\begin{aligned} \|\hat{\Theta}_{chol}^s - \Theta_k\|_{\ell_1} &= \|(I - \hat{A}^s)'(\hat{D}^s)^{-1}(I - \hat{A}^s) - (I - A_k)'D_k^{-1}(I - A_k)\|_{\ell_1} \\ &\leq \|(A_k - \hat{A}^s)'D_k^{-1}(I - \hat{A}^s)\|_{\ell_1} + \|(I - A_k)'D_k^{-1}(A_k - \hat{A}^s)\|_{\ell_1} + \\ &\quad \|(I - \hat{A}^s)'((\hat{D}^s)^{-1} - D_k^{-1})(I - \hat{A}^s)\|_{\ell_1} \\ &\leq C\|A_k - \hat{A}^s\|_{\ell_1} + C\|\hat{D}^s - D_k\|_{\ell_1}, \end{aligned}$$

where the last inequality used the fact that $\|D^{-1}\| \leq C_1$ and $\|I - A\| \leq C_2$ by Lemma A.2 of Bickel & Levina (2008a) and also the fact that $k\zeta = \alpha(1)$. Note that

$$\|\hat{D}^s - D_k\|_{\ell_1} \leq \|\hat{D}^s - D_k\|_{\max} \leq Ck^2\zeta^2$$

and

$$\|A_k - \hat{A}^s\|_{\ell_1} \leq k \cdot \|A_k - \hat{A}^s\|_{\max} \leq Ck\zeta.$$

By using the fact that $k\zeta = \alpha(1)$ again, we have

$$\|\hat{\Theta}_{chol}^s - \Theta_k\|_{\ell_1} \leq Ck\lambda + Ck^2\lambda^2 \leq Ck\zeta.$$

Next we cite an entry-wise estimation bound from Xue & Zou (2012) that for any positive quantity ζ satisfying $n\zeta = \alpha(1)$, there exists some fixed constant c such that

$$Pr(\|\tilde{R}^s - \sum^*\|_{\max} > \zeta) \leq p^2 \exp(-cn\zeta^2).$$

Since $\hat{R}^s$ is the nearest correlation matrix projection of $\tilde{R}^s$, then we have

$$\|\hat{R}^s - \sum^*\|_{\max} \leq 2\|\tilde{R}^s - \sum^*\|_{\max} \leq 2\lambda.$$

Thus, taking $\zeta = c \cdot n^{-1/2} \log^{1/2} p$ yields the upper bound for $\|\hat{\Theta}_{chol}^s - \Theta^*\|_{\ell_1}$. This completes the proof of Theorem 3. $\square$

References

- Adam B, Qu Y, Davis J, Ward M, Clements M, Cazares L, Semmes O, Schellhammer P, Yasui Y, Feng Z, Wright G. Serum protein fingerprinting coupled with a pattern-matching algorithm distinguishes prostate cancer from benign prostate hyperplasia and healthy men. *Cancer Research*. 2002; 62:3609–3614. [PubMed: 12097261]
- Bickel P, Levina E. Regularized estimation of large covariance matrices. *The Annals of Statistics*. 2008a; 36:199–227.

- Bickel P, Levina E. Covariance regularization by thresholding. *The Annals of Statistics*. 2008b; 36:2577–2604.
- Bien J, Tibshirani R. Sparse estimation of a covariance matrix. *Biometrika*. 2011; 98:807–820. [PubMed: 23049130]
- Box GEP, Cox DR. An analysis of transformations (with discussion). *Journal of the Royal Statistical Society Series B*. 1964; 26:211–252.
- Bradic J, Fan J, Wang W. Penalized composite quasi-likelihood for ultrahigh dimensional variable selection. *Journal of the Royal Statistical Society: Series B*. 2011; 73:325–349.
- Breiman L, Friedman J. Estimating optimal transformations for multiple regression and correlation (with discussion). *Journal of the American Statistical Association*. 1985; 80:580–598.
- Cai T, Zhou H. Minimax estimation of large covariance matrices under ℓ_1 norm (with discussion). *Statistica Sinica*. 2012; 22:1319–1378.
- Danaher P, Wang P, Witten D. The joint graphical lasso for inverse covariance estimation across multiple classes. *Journal of the Royal Statistical Society: Series B*. 2013 to appear.
- Devlin S, Gnanadesikan R, Kettenring J. Robust estimation and outlier detection with correlation coefficients. *Biometrika*. 1975; 62:531–545.
- Duchi J, Shalev-Shwartz S, Singer Y, Chandra T. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. *Proceedings of the 25th International Conference on Machine Learning*. 2008:272–279.
- Fan J, Liao Y, Mincheva M. Large covariance estimation by thresholding principal orthogonal complements (with discussion). *Journal of the Royal Statistical Society: Series B*. 2013; 75:603–680.
- Fan J, Ke T, Liu H, Xia L. QUADRO: A supervised dimension reduction method via rayleigh quotient optimization. *The Annals of Statistics*. 2013 to appear.
- Fan Y, Härdle W, Wang W, Zhu L. Composite quantile regression for the single-index model. *SFB 649 Discussion Papers*. 2013 No. 2013-010.
- Friedman JH. Greedy function approximation: a gradient boosting machine. *The Annals of Statistics*. 2001; 29:1189–1232.
- Friedman JH. Stochastic gradient boosting. *Computational Statistics & Data Analysis*. 2002; 38:367–378.
- Friedman JH, Stuetzle W. Projection pursuit regression. *Journal of the American statistical Association*. 1981; 76:817–823.
- Han F, Liu H. Semiparametric principal component analysis. *Advances in Neural Information Processing Systems*. 2012:171–179.
- Huang J, Liu N, Pourahmadi M, Liu L. Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika*. 2006; 93:85–98.
- Kendall, M. *Rank Correlation Methods*. Charles Griffin and Co. Ltd.; London: 1948.
- Koenker R, Bassett GJ. Regression quantiles. *Econometrica*. 1978; 46:33–50.
- Koenker R, Geling O. Reappraising medfly longevity: a quantile regression survival analysis. *Journal of the American Statistical Association*. 2001; 96:458–468.
- Koenker R, Xiao Z. Quantile autoregression (with discussion). *Journal of the American Statistical Association*. 2006; 101:980–990.
- Levina E, Rothman A, Zhu J. Sparse estimation of large covariance matrices via a nested lasso penalty. *The Annals of Applied Statistics*. 2008; 2:245–263.
- Lin Y, Jeon Y. Discriminant analysis through a semiparametric model. *Biometrika*. 2003; 90:379–392.
- Liu H, Han F, Yuan M, Lafferty J, Wasserman L. High dimensional semiparametric gaussian copula graphical models. *The Annals of Statistics*. 2012; 40:2293–2326.
- Liu H, Lafferty J, Wasserman L. The nonparanormal: semiparametric estimation of high dimensional undirected graphs. *Journal of Machine Learning Research*. 2009; 10:1–37.
- Ma S, Xue L, Zou H. Alternating direction methods for latent variable gaussian graphical model selection. *Neural Computation*. 2013; 25:2172–2198. [PubMed: 23607561]
- Mai Q, Zou H. Semiparametric sparse discriminant analysis in ultra-high dimensions. *Manuscript*. 2012

- Qi H, Sun D. A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM Journal on Matrix Analysis and Applications*. 2006; 28:360–385.
- Serfling, R. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons; 1980.
- Stone CJ. Additive regression and other nonparametric models. *The Annals of Statistics*. 1985; 13:689–705.
- Tibshirani R. Estimating transformations for regression via additivity and variance stabilization. *Journal of the American Statistical Association*. 1988; 83:394–405.
- Wang H, He X. Detecting differential expressions in GeneChip microarray studies: a quantile approach. *Journal of the American Statistical Association*. 2007; 102:104–112.
- Wegkamp M, Zhao Y. Adaptive estimation of the copula correlation matrix for semiparametric elliptical copulas. *Bernoulli*. 2015 to appear.
- Wu W, Pourahmadi M. Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika*. 2003; 90:831–844.
- Wu T, Yu K, Yu Y. Single-index quantile regression. *Journal of Multivariate Analysis*. 2010; 101:1607–1621.
- Xue L, Ma S, Zou H. Positive definite ℓ_1 penalized estimation of large covariance matrices. *Journal of the American Statistical Association*. 2012; 107:1480–1491.
- Xue L, Zou H. Regularized rank estimation of high-dimensional nonparanormal graphical models. *The Annals of Statistics*. 2012; 40:2541–2571.
- Xue L, Zou H. Rank-based tapering estimation of bandable correlation matrices. *Statistica Sinica*. 2014; 24:83–100.
- Zhao T, Roeder K, Liu H. Smooth-projected neighborhood pursuit for high-dimensional nonparanormal graph estimation. *Advances in Neural Information Processing Systems*. 2012:162–170.
- Zhu L, Huang M, Li R. Semiparametric quantile regression with high-dimensional covariates. *Statistica Sinica*. 2012; 22:1379–1401. [PubMed: 24501536]
- Zou H, Yuan M. Composite quantile regression and the oracle model selection theory. *The Annals of Statistics*. 2008; 36:1108–1126.

Figure 1.

Illustration of the non-normality issue of the protein mass spectroscopy data: (A1) and (A2) for the control data; (B1) and (B2) for the cancer data.

(A) Relative Prediction Errors for the control data

(B) Relative Prediction Errors for the cancer data

Figure 2.

The differences of prediction errors between the existing Cholesky-based method (Bickel & Levina 2008a) and the proposed rank-based estimate with $\tau = 0.5$. The solid dots below the horizontal line represent the mass/charge ratio indices where our proposed method outperforms the BL method.

(A) 95% Prediction Interval for the control data

(B) 95% Prediction Interval for the cancer data

Figure 3.

Prediction intervals for predicting protein mass spectroscopy intensities using our proposed method. The predict target, proposed point estimate ($\tau = 0.5$) and 95% prediction intervals are averaged over the testing set. The solid dots are the mass/charge ratio indices in which the proposed method outperforms the BL method.

Table 1

List of three regularized estimators in the simulation study.

Methods	Details
$(\hat{\sum}_{chol}^o, \hat{\Theta}_{chol}^o, \hat{t}_1^o(z))$	regularizing the “oracle” sample correlation matrix
$(\hat{\sum}_{chol}^s, \hat{\Theta}_{chol}^s, \hat{t}_1^s(z))$	regularizing the adjusted Spearman’s rank correlation matrix
$(\hat{\sum}_{chol}^n, \hat{\Theta}_{chol}^n, \hat{t}_1^n(z))$	regularizing the usual sample correlation matrix

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Performance of estimating the correlation matrix with $\sum_{chol}^o$ and $\sum_{chol}^n$. Each metric is averaged over 100 independent replications with standard errors in the bracket.

Method	Model 1			Model 2			Model 3		
	p=100	p=200	p=500	p=100	p=200	p=500	p=100	p=200	p=500
matrix ℓ_2 norm									
$\hat{\sum}_{chol}^o$	1.34 (0.05)	1.60 (0.13)	1.95 (0.27)	0.89 (0.03)	0.91 (0.04)	0.99 (0.08)	1.00 (0.02)	1.04 (0.05)	1.16 (0.12)
$\hat{\sum}_{chol}^s$	1.49 (0.07)	1.68 (0.11)	2.08 (0.21)	0.92 (0.03)	0.99 (0.05)	1.10 (0.14)	1.03 (0.02)	1.06 (0.04)	1.21 (0.12)
$\hat{\sum}_{chol}^n$	3.21 (0.06)	3.49 (0.09)	3.78 (0.24)	1.45 (0.02)	1.48 (0.04)	1.63 (0.11)	1.11 (0.03)	1.22 (0.05)	1.28 (0.11)
matrix ℓ_2 norm									
$\hat{\sum}_{chol}^o$	0.78 (0.03)	0.95 (0.06)	1.10 (0.13)	0.48 (0.01)	0.49 (0.02)	0.54 (0.04)	0.52 (0.01)	0.54 (0.02)	0.59 (0.05)
$\hat{\sum}_{chol}^s$	0.84 (0.03)	1.00 (0.06)	1.26 (0.11)	0.49 (0.01)	0.53 (0.02)	0.58 (0.05)	0.53 (0.01)	0.55 (0.02)	0.63 (0.06)
$\hat{\sum}_{chol}^n$	1.67 (0.05)	1.73 (0.08)	1.89 (0.14)	0.77 (0.01)	0.80 (0.02)	0.84 (0.04)	0.55 (0.01)	0.60 (0.02)	0.68 (0.05)

Table 3

Performance of estimating the inverse correlation matrix with $\hat{\Theta}_{chol}^{\circ}$, $\hat{\Theta}_{chol}^s$ and $\hat{\Theta}_{chol}^n$. Each metric is averaged over 100 independent replications with standard errors in the bracket.

Method	Model 1			Model 2			Model 3		
	p=100	p=200	p=500	p=100	p=200	p=500	p=100	p=200	p=500
	matrix ℓ_2 norm								
$\hat{\Theta}_{chol}^{\circ}$	4.07 (0.21)	4.72 (0.37)	5.14 (0.57)	1.83 (0.08)	1.95 (0.12)	1.99 (0.25)	2.21 (0.04)	2.29 (0.09)	2.37 (0.18)
$\hat{\Theta}_{chol}^s$	4.21 (0.24)	4.80 (0.32)	5.28 (0.55)	1.91 (0.07)	1.97 (0.14)	2.02 (0.25)	2.27 (0.04)	2.33 (0.07)	2.41 (0.16)
$\hat{\Theta}_{chol}^n$	7.18 (0.22)	7.55 (0.35)	8.02 (0.54)	3.29 (0.06)	3.48 (0.14)	3.70 (0.29)	2.81 (0.03)	2.93 (0.10)	3.08 (0.19)
	matrix ℓ_2 norm								
$\hat{\Theta}_{chol}^{\circ}$	3.18 (0.10)	3.41 (0.27)	4.07 (0.35)	1.26 (0.04)	1.36 (0.07)	1.41 (0.13)	1.68 (0.04)	1.74 (0.08)	1.75 (0.10)
$\hat{\Theta}_{chol}^s$	3.29 (0.16)	3.59 (0.25)	4.10 (0.30)	1.35 (0.04)	1.39 (0.07)	1.45 (0.08)	1.69 (0.04)	1.75 (0.08)	1.77 (0.11)
$\hat{\Theta}_{chol}^n$	4.98 (0.18)	5.15 (0.25)	5.33 (0.32)	2.64 (0.04)	2.81 (0.11)	2.89 (0.18)	2.29 (0.03)	2.32 (0.07)	2.39 (0.18)

Table 4

Performance of three optimal prediction methods $\hat{t}_1^s(z)$, $\hat{t}_1^o(z)$ and $\hat{t}_1^n(z)$. Prediction accuracy are measured by the relative prediction error and averaged over 100 independent replications, with the standard errors shown in the bracket.

Method	Model 1			Model 2			Model 3		
	p=100	p=200	p=500	p=100	p=200	p=500	p=100	p=200	p=500
$\hat{t}_1^s(z)$ vs. $\hat{t}_1^o(z)$	0.02 (0.00)	0.03 (0.00)	0.08 (0.01)	0.16 (0.02)	0.26 (0.03)	0.66 (0.04)	0.12 (0.01)	0.22 (0.01)	0.56 (0.04)
$\hat{t}_1^n(z)$ vs. $\hat{t}_1^o(z)$	0.11 (0.00)	0.19 (0.01)	0.21 (0.03)	13.50 (0.22)	26.07 (0.58)	63.80 (2.16)	7.38 (0.14)	14.70 (0.37)	35.68 (1.28)

Table 5

Testing for normality. This table shows the number of rejecting the normality hypothesis at different significance levels among 218 mass over charge ratio indices.

	significance level	Anderson-Darling	Cramer-von Mises	Kolmogorov-Smirnov
control data	0.05	158	138	119
	0.05/218	71	52	30
cancer data	0.05	196	177	156
	0.05/218	136	134	116