

TRANSFER LEARNING FOR HIGH-DIMENSIONAL QUANTILE REGRESSION VIA CONVOLUTION SMOOTHING

Yijiao Zhang and Zhongyi Zhu*

Fudan University

Abstract: This paper studies the high-dimensional quantile regression problem under the transfer learning framework, where possibly related source datasets are available to make improvements on the estimation or prediction based solely on the target data. In the oracle case with known transferable sources, a smoothed two-step transfer learning algorithm based on convolution smoothing is proposed and the ℓ_1/ℓ_2 estimation error bounds of the corresponding estimator are also established. To avoid including non-informative sources, we propose to select the transferable sources adaptively and establish its selection consistency under regular conditions. Monte Carlo simulations as well as an empirical analysis of gene expression data demonstrate the effectiveness of the proposed procedure.

Key words and phrases: High-dimensional data, quantile regression, regularization, smoothing, transfer learning.

1. Introduction

The increasing availability of datasets from multiple sources has provided us with unprecedented opportunities to get a better understanding of the data-limited target problem. For example, for the task of drug sensitivity prediction, the drug response data for the target type of cancer may be limited, but source data for another cancer type may be sufficient (Turki, Wei and Wang, 2017). However, there is no free lunch. Along with the satisfactory sample size of source studies comes the heterogeneity between the source and the target. Intuitively, the more related the source to the target, the more improvement may be made in learning about the target. This motivates transfer learning (Torrey and Shavlik, 2010; Pan and Yang, 2009; Weiss, Khoshgoftaar and Wang, 2016; Niu et al., 2020), which attempts to improve a learner from one domain by transferring information from a related but different domain. A considerable amount of research has clearly shown the success of transfer learning in many real-world applications, including ride dispatching (Wang et al., 2018), medical images analyses (Yu et al., 2022), and human activity recognition (Hirooka et al., 2022), etc.

This paper aims to investigate the effect of transfer learning on quantile regression (QR) in a high-dimensional setting. Ever since the influential work of

*Corresponding author. E-mail: zhuzy@fudan.edu.cn

Koenker and Bassett (1978), numerous scholars have explored the theoretical characterization of QR in various areas. See Koenker (2017) for a detailed review. Compared to the traditional conditional mean regression, QR captures the heterogeneous impact of regressors on different parts of the distribution. It also exhibits robustness to heteroscedastic and heavy-tailed errors. Moreover, in view of the frequently-collected high-dimensional data in various application domains including genomics, tomography, and finance, we focus on the regime where the dimension of covariates is substantially larger than the sample size.

A limited number of studies have sought to examine the possibility of interaction between quantile regression and multiple sources of data. Among them, Fan, Xue and Zou (2016) considered the multi-task quantile regression problem under the transnormal model, which may be restricted in practice. Moreover, large-scale covariance matrix estimation is needed in their method. In the field of data integration, Dai, Müller and Carroll (2023) proposed to use multiple datasets together with multiple quantiles to select variables simultaneously. They established model selection consistency and asymptotic normality of their estimator, but theoretical results about the estimation error are left unknown. More importantly, it needs to be emphasized that, in contrast to multi-task learning and data integration, where the datasets are equally important, the roles of the source and target tasks are no longer symmetric in transfer learning since we care most about the performance on the target data.

Despite the popularity of transfer learning on the practical side, statistical views of it as well as the theoretical underpinnings of related algorithms have been less than satisfactory. In the context of nonparametric classification, Cai and Wei (2021) proposed an adaptive classifier and established its minimax optimality. For high-dimensional data analysis, existing research has focused on penalized mean regression. Bastani (2021) considered linear regression and derived the estimation error bound for transferring knowledge from a single source to the target. Li, Cai and Li (2022a) investigated the case of multiple sources and proved the minimax optimal rate of their estimator. Tian and Feng (2023) further extended their work to generalized linear models and provided analogous theoretical guarantees. The benefits for confidence interval construction are also studied. These studies used a two-step framework, which was further applied to Gaussian graphical models with false discovery rate control (Li, Cai and Li, 2022b; He et al., 2022) and functional learning models (Lin and Reimherr, 2024). By measuring the similarity between the source and target by the difference of corresponding regression coefficients, Li, Cai and Li (2022a) and Tian and Feng (2023) also allow some of the source studies to be non-informative. However, one restriction of these studies is that they assume homoscedastic random errors with sub-gaussian distributions. While in real life, it is common to observe heteroscedastic variance (Delaigle and Meister, 2007, 2008) and heavy-tailed noises (Fan, Li and Wang, 2017; Sun, Zhou and Fan, 2020). Furthermore, the similarity between the source and target may vary across

different tails of the outcome distribution, which can not be reflected by the mean regression model.

Motivated by the above concerns, given a target dataset $\{((\mathbf{x}_i^{(0)})^\top, y_i^{(0)})\}_{i=1}^{n_0}$ and K independent source datasets $\{[(\mathbf{x}_i^{(k)})^\top, y_i^{(k)}]_{i=1}^{n_k}\}_{k=1}^K$, we investigate the conditional quantile of the response $y_i^{(k)}$ conditional on the covariates $\mathbf{x}_i^{(k)} \in \mathbb{R}^p$ at a given quantile level $\tau \in (0, 1)$, denoted by $\mathcal{Q}_{y_i^{(k)}|\mathbf{x}_i^{(k)}}(\tau)$. We consider a linear QR model, that is, $\mathcal{Q}_{y_i^{(k)}|\mathbf{x}_i^{(k)}}(\tau) = (\mathbf{x}_i^{(k)})^\top \mathbf{w}^{(k)}(\tau)$ with $\mathbf{w}^{(k)}(\tau) \in \mathbb{R}^p$ being the coefficient vector. We omit the dependence of $\mathbf{w}^{(k)}(\tau)$ on τ hereafter. The preceding model can be rewritten as

$$\mathbf{y}^{(k)} = \mathbf{X}^{(k)} \mathbf{w}^{(k)} + \boldsymbol{\epsilon}^{(k)}, \quad k = 0, \dots, K, \quad (1.1)$$

where $\mathbf{y}^{(k)} = (y_1^{(k)}, \dots, y_{n_k}^{(k)})^\top \in \mathbb{R}^{n_k}$, $\mathbf{X}^{(k)} = (\mathbf{x}_1^{(k)}, \dots, \mathbf{x}_{n_k}^{(k)})^\top \in \mathbb{R}^{n_k \times p}$ is the design matrix, and $\boldsymbol{\epsilon}^{(k)} \in \mathbb{R}^{n_k}$ is the noise vector with the i -th element $\epsilon_i^{(k)}$ satisfying $\mathbb{P}(\epsilon_i^{(k)} \leq 0 | \mathbf{x}_i^{(k)}) = \tau$. The target parameter $\boldsymbol{\beta} := \mathbf{w}^{(0)}$, which satisfies $\|\boldsymbol{\beta}\|_0 = s \ll p$, is of our primary interest. Besides, we do not assume ℓ_0 -sparsity on $\mathbf{w}^{(k)}$ for $k = 1, \dots, K$. Following the spirit of Li, Cai and Li (2022a) and Tian and Feng (2023), we define the k -th contract vector $\boldsymbol{\delta}^{(k)} = \boldsymbol{\beta} - \mathbf{w}^{(k)}$ and the transferability of the k -th study as $\|\boldsymbol{\delta}^{(k)}\|_1$. In general, we prefer $\|\boldsymbol{\delta}^{(k)}\|_1$ to be sufficiently small to guarantee performance improvement of transfer learning. With this insight in mind, the transferable set of source studies is defined as $\mathcal{A}_\eta = \{1 \leq k \leq K : \|\boldsymbol{\delta}^{(k)}\|_1 \leq \eta\}$.

Compared with existing works for the (generalized) linear models under the transfer learning framework (Li, Cai and Li, 2022a; Tian and Feng, 2023), our proposed QR model has several advantages:

1. It offers a more complete picture of the target problem by varying τ ;
2. It could handle heterogeneity due to either heteroscedastic variance or other forms of non-location-scale covariate effects;
3. It relaxes the distributional conditions on the error terms, which is more robust to tail behavior;
4. It allows the difference $\boldsymbol{\delta}^{(k)}$ as well as the transferable set \mathcal{A}_η to change with τ , which is more flexible.

However, the adoption of QR plays the role of a double-edged sword. More specifically, under the transfer learning framework, three critical issues need to be addressed: (i) *what to transfer*, (ii) *how to transfer*, and (iii) *when to transfer* (Pan and Yang, 2009). The first one can be similarly solved by defining some common component as that in Li, Cai and Li (2022a) and Tian and Feng (2023). A natural solution to the second issue is to adopt the two-step framework. Unfortunately, due to the non-differentiability of the quantile loss function, this can be not

only technically challenging but also computationally expensive, especially in the transfer learning setting when the auxiliary source data may be extremely large. Besides, the (possibly) unrelated tasks pose additional challenges to the last issue, which involves correctly identifying the transferable set \mathcal{A}_η to circumvent negative transfer (Zhang et al., 2022; Niu et al., 2020).

To tackle the aforementioned challenges, we combine a smoothed two-step procedure with a source detection algorithm for transferring high-dimensional QR models. To deal with the non-smoothness, we employ a recently developed convolution smoothing technique (Fernandes, Guerre and Horta, 2021; He et al., 2023) to smooth the piecewise linear quantile loss. Convolution smoothing and convex relaxation enable us to use gradient-based algorithms which are much more scalable to large-scale datasets. At the same time, delicate analysis of the smoothing bandwidths is also needed to control the smoothing bias. A distributed QR transfer approach is also proposed for the ease of computation burden. To avoid negative transfer, we propose to evaluate the change of performance on the target with and without each source and exclude unrelated sources which lead to worse performance. This idea originates from the works by Eaton, Desjardins and Lane (2008) and Tian and Feng (2023).

In theory, our contributions are fourfold. Firstly, by choosing smoothing bandwidths properly to have the smoothing bias controlled, we derived the ℓ_1 and ℓ_2 estimation error bounds for our smoothed two-step estimator, which is complementary to the existing results in the mean regression world. The convergence rate is shown to be faster than that of the single-task smoothed high-dimensional quantile regression established in Tan, Wang and Zhou (2022). Secondly, we show that the transferable set can be consistently detected by the proposed clustering method when there is a sufficiently large gap between the positive and negative sources. Thirdly, the statistical property of our distributed QR transfer estimator is also established. As a byproduct, two lemmas related to the local restricted strong convexity (RSC) of the empirical smoothed quantile loss are established, which provide a core result for establishing error bounds for our smoothed two-step QR estimators and may be of independent interest.

The rest of the paper is organized as follows. In Section 2, we present our smoothed two-step estimator with the known transferable set, followed by a source detection algorithm to select the transferable set. Section 3 is dedicated to providing theoretical guarantees for our proposed method, including estimation error bounds and the selection consistency of the source detection algorithm. We demonstrate our proposed method on simulated data in Section 4 and an empirical analysis of Genotype-Tissue Expression (GTEx) data in Section 5. We conclude with some discussions on possible extensions in Section 6. The proofs of all theoretical results, additional simulation results, and an extension to the distributed QR transfer are relegated to the Supplementary Material.

We finish this section with notation. Throughout this paper, we use bold capitalized letters (e.g. \mathbf{X}, \mathbf{A}) to denote matrices and use bold little letters (e.g. \mathbf{x}, \mathbf{y}) to denote vectors. For a p -dimensional vector $\mathbf{x} = (x_1, \dots, x_p)^T$, we denote its ℓ_q -norm as $\|\mathbf{x}\|_q = (\sum_{i=1}^p |x_i|^q)^{1/q}$ ($q \in (0, 2]$), ℓ_0 -norm as $\|\mathbf{x}\|_0 = \#\{j : x_j \neq 0\}$ and ℓ_∞ -norm as $\|\mathbf{x}\|_\infty = \max_j |x_j|$. We use $|\mathcal{I}|$ to denote the cardinal number of a set \mathcal{I} . For any positive integer, n , we denote the index set $\{1, \dots, n\}$ as $[n]$. We use $\mathbf{I}\{\cdot\}$ to denote the indicator function. For a matrix $\mathbf{A}_{p \times q} = [a_{ij}]_{m \times n}$, let $\|\mathbf{A}\|_1 = \sup_j \sum_{i=1}^m |a_{ij}|$ and $\|\mathbf{A}\|_{\max} = \sup_{i,j} |a_{ij}|$ denote its ℓ_1 -norm and max-norm respectively. For any $k \times k$ symmetric, positive semidefinite matrix $\mathbf{A} \in \mathbb{R}^{m \times m}$, we use $\Lambda_{\min}(\mathbf{A})$ and $\Lambda_{\max}(\mathbf{A})$ to denote its minimum and maximum eigenvalue. For any two real numbers a and b , we write $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. For two sequences of non-negative numbers $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, $a_n \lesssim b_n$ or $b_n \gtrsim a_n$ indicate that $\sup_n |a_n/b_n| < \infty$; $a_n \asymp b_n$ is equivalent to $a_n \lesssim b_n$ and $b_n \lesssim a_n$. We use $a_n \ll b_n, b_n \gg a_n$ or $a_n = o(b_n)$ to represent $|a_n/b_n| \rightarrow 0$ as $n \rightarrow \infty$.

2. Methodology

2.1. QR basics

We start by providing a brief introduction to QR. Consider the QR model

$$y_i = \mathbf{x}_i^\top \mathbf{w} + \epsilon_i, \quad \mathbb{P}(\epsilon_i \leq 0 \mid \mathbf{x}_i) = \tau, \quad i = 1, \dots, n \tag{2.1}$$

where $\mathbf{w} \in \mathbb{R}^p$ is the coefficient vector with $\|\mathbf{w}\|_0 = s$. Following Belloni and Chernozhukov (2011), we can estimate \mathbf{w} by fitting the ℓ_1 -penalized quantile regression (ℓ_1 -QR):

$$\tilde{\mathbf{w}} \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n \rho_\tau(y_i - \mathbf{x}_i^\top \mathbf{w}) + \lambda_{\mathbf{w}} \|\mathbf{w}\|_1 \right\}, \tag{2.2}$$

where $\rho_\tau(x) = x\{\tau - \mathbf{I}(x \leq 0)\}$ is the check function. However, if we fit ℓ_1 -QR based solely on the target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ to obtain an estimator $\tilde{\boldsymbol{\beta}}_{\text{tar}}$, its performance can be limited by the sample size n_0 of the target data, which motivates us to incorporate other related sources to make some improvements.

2.2. Smoothed two-step QR transfer

We start with the case when the transferable set \mathcal{A}_η is known. As there is heterogeneity between the target and source data, it leads us to think about the first issue in transfer learning: *What to transfer*. The key point is to find a common component between the source and the target. In view of this, we pool the target and the sources in the transferable set \mathcal{A}_η together and consider the

parameter $\mathbf{w}^{\mathcal{A}_\eta}$ identified by the following moment equation

$$\sum_{k \in \mathcal{A}_\eta \cup \{0\}} \alpha_k \mathbb{E} \left(\mathbf{x}^{(k)} \left[F_{\epsilon^{(k)} | \mathbf{x}^{(k)}} \left\{ (\mathbf{x}^{(k)})^\top (\mathbf{w}^{\mathcal{A}_\eta} - \mathbf{w}^{(k)}) \mid \mathbf{x}^{(k)} \right\} - \tau \right] \right) = 0, \quad (2.3)$$

where $\alpha_k = n_k / (n_{\mathcal{A}_\eta} + n_0)$ with $n_{\mathcal{A}_\eta} = \sum_{k \in \mathcal{A}_\eta} n_k$ and $F_{\epsilon^{(k)} | \mathbf{x}^{(k)}}(\cdot \mid \mathbf{x}^{(k)})$ is the conditional distribution function of $\epsilon^{(k)}$ given $\mathbf{x}^{(k)}$. Further denote $f_{\epsilon^{(k)} | \mathbf{x}^{(k)}}(\cdot)$ as the conditional density of $\epsilon^{(k)}$ given $\mathbf{x}^{(k)}$, and $\mathbf{M}_k = \mathbb{E} \left\{ \int_0^1 \mathbf{x}^{(k)} (\mathbf{x}^{(k)})^\top f_{\epsilon^{(k)} | \mathbf{x}^{(k)}}(t(\mathbf{x}^{(k)}))^\top (\mathbf{w}^{\mathcal{A}_\eta} - \mathbf{w}^{(k)}) dt \right\}$. Then we can write

$$\mathbf{w}^{\mathcal{A}_\eta} = \boldsymbol{\beta} - \boldsymbol{\delta}^{\mathcal{A}_\eta},$$

where $\boldsymbol{\delta}^{\mathcal{A}_\eta} = (\sum_{k \in \mathcal{A}_\eta \cup \{0\}} \alpha_k \mathbf{M}_k)^{-1} \sum_{k \in \mathcal{A}_\eta \cup \{0\}} \alpha_k \mathbf{M}_k \boldsymbol{\delta}^{(k)}$. The bias $\boldsymbol{\delta}^{\mathcal{A}_\eta}$ is a weighted average of $\boldsymbol{\delta}^{(k)}$, which is supposed to be ℓ_1 -sparse under certain conditions. Since the moment equation (2.3) incorporates information from the sources and the target, $\mathbf{w}^{\mathcal{A}_\eta}$ can be seen as a shared knowledge which can be transferred.

Now we move to consider the second issue: *how to transfer*. A commonly used heuristic in transfer learning is to adopt the two-step framework (Bastani, 2021; Li, Cai and Li, 2022a; Tian and Feng, 2023). Specifically, in the first transferring step, we can pool the sources and target together to obtain an estimator $\hat{\mathbf{w}}^{\mathcal{A}_\eta}$ and then further correct the bias of $\hat{\mathbf{w}}^{\mathcal{A}_\eta}$ by calibrating it using the target data. Both steps could be accomplished by fitting ℓ_1 -QR similar to that in (2.2). However, the nondifferentiable quantile loss function brings challenges to both computation and theory establishment.

On the computation side, the minimization problem in (2.2) can be reformulated as a linear program (Wang, Wu and Li, 2012; Peng and Wang, 2015) of which the computation complexity grows with both p and n . In the setting of multiple sources as well as high-dimensional data, this approach may suffer from heavy computational costs. On the theoretical side, due to the nonsmooth loss and (possibly) non-sparsity of $\hat{\mathbf{w}}^{\mathcal{A}_\eta}$, it is difficult to establish theoretical guarantees.

To overcome these challenges, we propose a smoothed version of the two-step procedure by smoothing the piecewise linear quantile loss via convolution. Recall the QR model in (2.1). The ℓ_1 -penalized smoothed quantile regression (ℓ_1 -SQR) estimator (Tan, Wang and Zhou, 2022) is defined as

$$\hat{\mathbf{w}} \in \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^p} \left\{ \frac{1}{nh} \sum_{i=1}^n \int_{-\infty}^{\infty} \rho_\tau(u) K \left(\frac{u + \mathbf{x}_i^\top \mathbf{w} - y_i}{h} \right) du + \lambda \|\mathbf{w}\|_1 \right\}, \quad (2.4)$$

where h is the smoothing bandwidth, λ is a tuning parameter that controls the model complexity, and $K : \mathbb{R} \rightarrow [0, \infty)$ is a symmetric, non-negative kernel that integrates to one. For ease of notation, we use ℓ_1 -SQR($\{(\mathbf{x}_i, y_i)\}_{i \in \mathcal{I}}; \lambda, h$) to denote the estimator in (2.4) given a dataset $\{(\mathbf{x}_i, y_i)\}_{i \in \mathcal{I}}$ with index set \mathcal{I} , a

tuning parameter λ and a bandwidth h .

The convolution-type smoothing first introduced by Fernandes, Guerre and Horta (2021) yields a convex and twice differentiable loss, which enables the use of gradient-based algorithms and hence eases the computational burden. Additionally, it is also convenient for statistical analysis. Consider the single task quantile regression based solely on the target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$. In the low-dimensional regime where $p \ll n_0$, Fernandes, Guerre and Horta (2021) provided a comprehensive asymptotic analysis for the unpenalized smoothed QR estimator, followed by an in-depth finite sample theory in He et al. (2023). More related to our work, Tan, Wang and Zhou (2022) investigated convolution smoothing for high-dimensional QR on both the theoretical and computational sides. They showed that with a proper yet flexible choice of the bandwidth, the ℓ_1 -SQR estimator shares the same ℓ_1 and ℓ_2 error upper bounds as the ℓ_1 -QR estimator (Belloni and Chernozhukov, 2011), which are $s\sqrt{\log p/n_0}$ and $\sqrt{s \log p/n_0}$ respectively. In addition, a coordinate descent algorithm and an alternating direction method of multiplier algorithm were also proposed for solving the ℓ_1 -SQR problem in (2.4) with the uniform kernel and Gaussian kernel respectively.

Combining this smoothing procedure with our two-step framework leads to Algorithm 1 (Oracle-Trans-SQR). Here $\mathcal{I}_k = [n_k]$ for $k = 0, \dots, K$. The smoothed two-step estimator $\hat{\beta}$ is expected to converge fast to our target β provided a substantial sample size of source studies in the transferring step along with a sufficiently small bias in the debiasing step.

Algorithm 1 Oracle-Trans-SQR Algorithm.

Input: Target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, source data $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k \in \mathcal{A}_\eta}$, penalty parameters $(\lambda_w, \lambda_\delta)$ and bandwidths (h_w, h_δ) .

- 1: **Transferring:** $\hat{w}^{\mathcal{A}_\eta} \leftarrow \ell_1\text{-SQR}(\{\{(\mathbf{x}_i^{(k)}, y_i^{(k)})\}_{i \in \mathcal{I}_k}\}_{k \in \mathcal{A}_\eta \cup \{0\}}; \lambda_w, h_w)$
- 2: **Debiasing:** $\hat{\delta}^{\mathcal{A}_\eta} \leftarrow \ell_1\text{-SQR}(\{(\mathbf{x}_i^{(0)}, y_i^{(0)}) - (\mathbf{x}_i^{(0)})^\top \hat{w}^{\mathcal{A}_\eta}\}_{i \in \mathcal{I}_0}; \lambda_\delta, h_\delta)$

Output: $\hat{\beta} = \hat{w}^{\mathcal{A}_\eta} + \hat{\delta}^{\mathcal{A}_\eta}$

Due to the memory constraints of a single machine, we also provide a distributed version of Algorithm 1 with corresponding statistical properties, which are relegated to the Supplementary Material.

2.3. Transferable set detection

In practice, the transferable set \mathcal{A}_η may be unknown, which makes the problem more intractable. Indeed, the performance of learning on the target may be reduced if irrelevant sources are included in transfer learning, which is defined as *negative transfer* (Pan and Yang, 2009; Ge et al., 2014). It forces us to take the third issue into consideration: *when to transfer*.

To avoid negative transfer, Li, Cai and Li (2022a) and Lin and Reimherr (2024) proposed to aggregate a collection of candidate estimators, which leads to a relatively robust estimator. However, previous research on QR model aggregation has been restricted to the out-of-sample prediction error (Lu and Su, 2015; Wang et al., 2023). Relatively little is understood about the optimality of the estimation error after model aggregation.

Intuited by (Eaton, Desjardins and Lane, 2008), we propose to select the sources with positive transfer by thresholding the change in performance on the target between learning with and without each source. Specifically, we randomly split the target index set \mathcal{I}_0 into the training part $\mathcal{I}_0^{\text{tr}}$ and the validation part $\mathcal{I}_0^{\text{va}}$ with equal size. We first obtain a benchmark estimate $\hat{\beta}^{(0)}$ by fitting ℓ_1 -SQR on the target training data. Next, we carry out ℓ_1 -SQR based on the k -th source data as well as the target training data to obtain an estimate $\hat{\beta}^{(k)}$, which is exactly the transferring step in Algorithm 1 with $\mathcal{A}_\eta = \{k\}$ and $\mathcal{I}_0 = \mathcal{I}_0^{\text{tr}}$. Then we examine the difference of quantile losses on the target validation data based on $\hat{\beta}^{(k)}$ and $\hat{\beta}^{(0)}$, denoted by $\hat{T}^{(k)} = \widehat{Q}^{(0)}(\hat{\beta}^{(k)}; \mathcal{I}_0^{\text{va}}) - \widehat{Q}^{(0)}(\hat{\beta}^{(0)}; \mathcal{I}_0^{\text{va}})$, where $\widehat{Q}^{(0)}(\mathbf{w}; \mathcal{I}) = 1/|\mathcal{I}| \sum_{i \in \mathcal{I}} \rho_\tau\{y_i^{(0)} - (\mathbf{x}_i^{(0)})^\top \mathbf{w}\}$. We call $\hat{T}^{(k)}$ the transferability index of the k -th study, which is preferable when it is smaller. The transferable set \mathcal{A}_η is then estimated by the source datasets whose transferability indices are lower than a prespecified threshold. The detailed source detection procedure is presented in Algorithm 2.

Algorithm 2 Trans-SQR Algorithm.

Input: Target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, source data $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k=1}^K$, penalty parameters $\{\lambda^{(k)}\}_{k=0}^K$, bandwidths $\{h^{(k)}\}_{k=0}^K$ and a threshold t .

1: **Estimating transferability index:**

- (1) Randomly split $\mathcal{I}_0 = [n_0]$ into $\mathcal{I}_0 = \mathcal{I}_0^{\text{tr}} \cup \mathcal{I}_0^{\text{va}}$ with $|\mathcal{I}_0^{\text{va}}| = \lfloor n_0/2 \rfloor$.
- (2) Compute $\hat{\beta}^{(0)} \leftarrow \ell_1\text{-SQR}(\{\mathbf{x}_i^{(0)}, y_i^{(0)}\}_{i \in \mathcal{I}_0^{\text{tr}}}; \lambda^{(0)}, h^{(0)})$.
- (3) For $k = 1, \dots, K$, compute

$$\hat{\beta}^{(k)} \leftarrow \ell_1\text{-SQR}(\{\mathbf{x}_i^{(k)}, y_i^{(k)}\}_{i \in \mathcal{I}_k} \cup \{\mathbf{x}_i^{(0)}, y_i^{(0)}\}_{i \in \mathcal{I}_0^{\text{tr}}}; \lambda^{(k)}, h^{(k)}).$$

- (4) For $k = 1, \dots, K$, compute $\hat{T}^{(k)} = \widehat{Q}^{(0)}(\hat{\beta}^{(k)}; \mathcal{I}_0^{\text{va}}) - \widehat{Q}^{(0)}(\hat{\beta}^{(0)}; \mathcal{I}_0^{\text{va}})$.

2: **Source detection:** Let $\hat{\mathcal{A}} = \{k : \hat{T}^{(k)} < t(\widehat{Q}^{(0)}(\hat{\beta}^{(0)}; \mathcal{I}_0^{\text{va}}) \vee 0.01)\}$.

3: **Trans-SQR:** $\hat{\beta} \leftarrow$ run Algorithm 1 with $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k \in \hat{\mathcal{A}} \cup \{0\}}$

Output: $\hat{\beta}$

Under regular conditions, the detected transferable set $\hat{\mathcal{A}}$ will be shown to exactly identify the unknown \mathcal{A}_η for some η with high probability. Furthermore, in comparison to the model aggregation method proposed in Li, Cai and Li (2022a), our method greatly eases the computation burden since we do not need to execute the oracle transferring procedure repeatedly based on a set of candidate estimates of \mathcal{A}_η .

3. Theoretical Guarantees

In this section, we will investigate the statistical properties of the proposed algorithms in Section 2. We first establish high probability bounds of our QR transfer estimator in Algorithm 1 in Section 3.1. In Section 3.2, we establish the consistency of the source detection procedure in Algorithm 2.

3.1. Estimation error of Oracle-Trans-SQR

We begin by imposing some regularity conditions in smoothed quantile regression. Denote the covariance matrix of $\mathbf{x}^{(k)}$ as $\Sigma^{(k)} = (\sigma_{ij}^{(k)})_{1 \leq i, j \leq p}$.

Condition 1. For $k = 0, \dots, K$ and some constant $b_0 > 0$, $\inf_k f_{\epsilon^{(k)}|\mathbf{x}^{(k)}}(t) \geq f_l > 0$ for all $|t| \leq b_0$ and $\sup_k f_{\epsilon^{(k)}|\mathbf{x}^{(k)}}(0) \leq f_u < \infty$ almost surely. Besides, there exists a constant $l_0 > 0$ such that $|f_{\epsilon^{(k)}|\mathbf{x}^{(k)}}(u) - f_{\epsilon^{(k)}|\mathbf{x}^{(k)}}(v)| \leq l_0|u - v|$ for all $u, v \in \mathbb{R}$ almost surely.

Condition 2. The kernel function $K : \mathbb{R} \rightarrow [0, \infty)$ is symmetric around zero, and satisfies

$$(a) \int_{-\infty}^{\infty} K(u)du = 1 \text{ and } \int_{-\infty}^{\infty} u^2 K(u)du < \infty;$$

$$(b) \kappa_\ell = \int_{-\infty}^{\infty} |u|^\ell K(u)du < \infty \text{ for } \ell = 1, 2, \dots \text{ and } \kappa_1 = \min_{|u| \leq 1} K(u) > 0.$$

Condition 3. For $k = 0, \dots, K$, the covariate vector $\mathbf{x}^{(k)} \in \mathbb{R}^p$ is sub-Gaussian: there exists some $v_1 \geq 1$ such that $\mathbb{P}(|\mathbf{u}^\top \mathbf{x}^{(k)}| \geq v_1 \|\mathbf{u}\|_2 \cdot t) \leq 2e^{-t^2/2}$ for all $\mathbf{u} \in \mathbb{R}^p$ and $t \geq 0$. In addition, $\mu_k = \sup_{\mathbf{u} \in \mathbb{S}^{p-1}} \mathbb{E}|\langle \mathbf{u}, \mathbf{x} \rangle|^k < \infty$ for $k = 1, \dots, 4$ and $0 < \gamma_p \leq \inf_k \Lambda_{\min}(\Sigma^{(k)}) \leq \sup_k \Lambda_{\max}(\Sigma^{(k)}) \leq \gamma_1$.

Remark 1. Condition 1 imposes regular conditions on the conditional density of the random errors. We allow for heteroscedastic errors. Condition 2 is commonly used for kernel functions (Fernandes, Guerre and Horta, 2021; Tan, Wang and Zhou, 2022). Some detailed comments can be found in Remark 4.1 of Tan, Wang and Zhou (2022). Condition 3 assumes sub-Gaussian covariates with well-behaved covariance structures. We comment that our theoretical analysis will still carry through if we impose a weaker restricted eigenvalue condition similar to that in Bickel, Ritov and Tsybakov (2009).

The next condition characterizes the differences between $\mathbf{M}_{\mathcal{A}_\eta}$ and \mathbf{M}_k defined in Section 2.2, which is used to ensure that the common component $\mathbf{w}^{\mathcal{A}_\eta}$ to be transferred is close to our target β .

Condition 4. Denote $\mathbf{M}_{\mathcal{A}_\eta} = \sum_{k \in \mathcal{A}_\eta \cup \{0\}} \alpha_k \mathbf{M}_k$ with α_k defined in (2.3). For some constant c_M , it holds that $\sup_{k \in \mathcal{A}_\eta} \|\mathbf{M}_{\mathcal{A}_\eta}^{-1} \mathbf{M}_k\|_1 = c_M < \infty$.

Remark 2. Condition 4 is related in spirit to Assumption 4 in Tian and Feng (2023) as well as Condition 4 in Li, Cai and Li (2022a). We regard Condition 4 as reasonable as it does not require exact sparsity of $\mathbf{M}_{\mathcal{A}_\eta}^{-1} \mathbf{M}_k$. For example, the

ℓ_1 -norm of the inverses are bounded for symmetric diagonally dominant positive matrices (Li, 2008) as well as matrices with its (i, j) -th component being $\rho^{|i-j|}$ for some $\rho \in (-1, 1)$. Besides, it also holds true for banded matrices with a fixed bandwidth (Demko, 1977).

Condition 5. Assume that $\|\mathbf{x}^{(k)}\|_\infty \leq B_k$ almost surely for some $B_k > 0$ and $B_k \|\boldsymbol{\delta}_{\mathcal{A}} - \boldsymbol{\delta}^{(k)}\|_1 \leq b_0$, for all $k \in \mathcal{A}_\eta \cup \{0\}$.

Remark 3. The upper bound B_k in Condition 5 can be $O(1)$ when $\mathbf{x}^{(k)}$ has bounded components and $O(\log p)$ for general sub-Gaussian designs. Therefore, a sufficient condition for Condition 5 to hold is $\eta \lesssim 1$ in the former case and $\eta \lesssim 1/\log p$ in the latter.

With the above preparations, we are ready to present the main result for our Oracle-Trans-SQR algorithm. Consider the parameter space

$$\Theta(s, \eta) = \left\{ \mathbf{B} = \left(\boldsymbol{\beta}, \boldsymbol{\delta}^{(1)}, \dots, \boldsymbol{\delta}^{(K)} \right) : \|\boldsymbol{\beta}\|_0 \leq s, \sup_{k \in \mathcal{A}_\eta} \|\boldsymbol{\delta}^{(k)}\|_1 \leq \eta \right\}.$$

Theorem 1. Assume Conditions 1–5 hold. Suppose that $n_0 \gtrsim \log p$, $n_{\mathcal{A}_\eta} \gtrsim s^2 \log p$, and $\eta = O(1)$. Let $\lambda_{\mathbf{w}} \asymp \{\log p / (n_{\mathcal{A}_\eta} + n_0)\}^{1/2}$ and $\lambda_{\boldsymbol{\delta}} \asymp (\log p / n_0)^{1/2}$. Choose the bandwidths as $h_{\mathbf{w}} \asymp \{s \log p / (n_{\mathcal{A}_\eta} + n_0)\}^{1/4}$ and $h_{\boldsymbol{\delta}} \asymp (\log p / n_0)^{1/4}$. Then for the estimator $\hat{\boldsymbol{\beta}}$ obtained from the Oracle-Trans-SQR algorithm, there exist positive constants c_1 and c_2 such that

$$\begin{aligned} \inf_{\mathbf{B} \in \Theta(s, \eta)} \mathbb{P} \left\{ \|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 \lesssim s \left(\frac{\log p}{n_{\mathcal{A}_\eta} + n_0} \right)^{1/2} + \eta \right\} &\geq 1 - c_1 \exp(-c_2 \log p), \\ \inf_{\mathbf{B} \in \Theta(s, \eta)} \mathbb{P} \left[\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \lesssim \left(\frac{\log p}{n_0} \right)^{1/4} \left\{ \sqrt{s} \left(\frac{\log p}{n_{\mathcal{A}_\eta} + n_0} \right)^{1/4} + \sqrt{\eta} \right\} \right] \\ &\geq 1 - c_1 \exp(-c_2 \log p). \end{aligned}$$

Remark 4. The above theorem established the convergence rate of $\hat{\boldsymbol{\beta}}$ from the Oracle-Trans-SQR algorithm under a proper choice of bandwidths. For single-task smoothed quantile regression based solely on the target data, Tan, Wang and Zhou (2022) obtained the rate for ℓ_2 bound as $\sqrt{s \log p / n_0}$. Hence, our proposed estimator will enjoy a sharper convergence rate provided that $\eta \ll s \sqrt{\log p / n_0}$ and $n_{\mathcal{A}} \gg n_0$. These results parallel analogous results in Li, Cai and Li (2022a) for linear regression and in Tian and Feng (2023) for generalized linear models.

3.2. Source detection consistency

Next, we investigate the theoretical property of the Trans-SQR algorithm. The key ingredient of this part is to establish the detection consistency of the transferable set \mathcal{A}_η . Let $\mathcal{A}_\eta^c = \{1, \dots, K\} \setminus \mathcal{A}_\eta$ and $n_{\min} = \min_{1 \leq k \leq K} n_k$. Similar

to (2.3), let $\boldsymbol{\beta}^{(k)}$ be the parameter identified by

$$\sum_{k' \in \{0, k\}} \alpha_{k'} \mathbb{E} \left(\mathbf{x}^{(k')} \left[F_{\epsilon^{(k')} | \mathbf{x}^{(k')}} \left\{ (\mathbf{x}^{(k')})^\top (\boldsymbol{\beta}^{(k)} - \mathbf{w}^{(k')}) \mid \mathbf{x}^{(k')} \right\} - \tau \right] \right) = 0.$$

Condition 6. Suppose that there exist some positive s' and η' independent of k such that $\boldsymbol{\beta}^{(k)}$ can be decomposed into $\boldsymbol{\beta}^{(k)} = \boldsymbol{\iota}^{(k)} + \boldsymbol{\varpi}^{(k)}$ with $\|\boldsymbol{\iota}^{(k)}\|_0 \leq s'$ and $\|\boldsymbol{\varpi}^{(k)}\|_1 \leq \eta'$ for $k = 1, \dots, K$.

Remark 5. Condition 6 imposes a “weak” sparse structure on $\boldsymbol{\beta}^{(k)}$, which is similar to Assumption 6 in Tian and Feng (2023). We allow the support of $\boldsymbol{\iota}^{(k)}$, denoted by \mathcal{S}_k , to be significantly different for $k \in \mathcal{A}_\eta$ and $k \in \mathcal{A}_\eta^c$.

Define $\mathbf{R}^{(k)} = \mathbb{E} \left\{ \int_0^1 \mathbf{x}^{(0)} (\mathbf{x}^{(0)})^\top f_{\epsilon^{(0)} | \mathbf{x}^{(0)}}(t(\mathbf{x}^{(0)})^\top (\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta})) dt \right\}$. Let $\Omega_{\max} = \sqrt{s' \log p / (n_{\min} + n_0)} + \{\log p / (n_{\min} + n_0)\}^{1/4} \sqrt{\eta'} + \{\log p / (n_{\min} + n_0)\}^{3/8} (s')^{-1/8} \eta'$, $\Omega_0 = \sqrt{s \log p / n_0}$, and $c_n = (\mu_1 + 2v_1 \sqrt{\log p / n_0})(\Omega_{\max} + \Omega_0)$. Here μ_1 and v_1 are defined in Condition 3.

Condition 7. Suppose that there exist some η and a constant C such that $\inf_{k \in \mathcal{A}_\eta^c} \Lambda_{\min}(\mathbf{R}^{(k)}) := \underline{\lambda} > 0$, $\sup_{k \in \mathcal{A}_\eta} \Lambda_{\max}(\mathbf{R}^{(k)}) := \bar{\lambda} < \infty$ and

$$\begin{aligned} & \inf_{k \in \mathcal{A}_\eta^c} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2^2 \\ & \geq \underline{\lambda}^{-1} \left(\bar{\lambda} \sup_{k \in \mathcal{A}_\eta} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2^2 + 12v_1 \sqrt{\frac{\log p}{n_0}} \sup_{k \in [K]} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 + 2Cc_n \right). \end{aligned}$$

Remark 6. Condition 7 ensures the identifiability of some \mathcal{A}_η by Trans-SQR. It assumes that for sources not in the transferable set \mathcal{A}_η , there is a significant gap between the population-level coefficient from the transferring step and the true coefficient of the target data. It is parallel to Assumption 5 in Tian and Feng (2023).

Theorem 2. Assume that Conditions 1–6 hold and Condition 7 is satisfied for some η . Suppose that $n_{\mathcal{A}_\eta} \gtrsim s^2 \log p$. Let $\lambda_w \asymp \{\log p / (n_{\mathcal{A}_\eta} + n_0)\}^{1/2}$, and $\lambda_\delta \asymp (\log p / n_0)^{1/2}$. For $k = 0, \dots, K$, let $\lambda^{(k)} \asymp \{\log p / (n_k + n_0)\}^{1/2}$ and $h^{(k)} \asymp \{s' \log p / (n_k + n_0)\}^{1/4}$. Choose the threshold t as $\bar{\lambda} \sup_{k \in \mathcal{A}_\eta} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2^2 + \sqrt{\log p / n_0} \sup_{k \in [K]} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 + c_n \lesssim t \lesssim \underline{\lambda} / 2 \inf_{k \in \mathcal{A}_\eta^c} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2^2$. Then for the detected set $\hat{\mathcal{A}}$ in Algorithm 2, we have

$$\mathbb{P}(\hat{\mathcal{A}} = \mathcal{A}_\eta) \rightarrow 1.$$

Consequently, the estimator $\hat{\boldsymbol{\beta}}$ obtained from the Trans-SQR algorithm enjoys the same ℓ_1/ℓ_2 -estimation error upper bounds in Theorem 1 with high probability.

Remark 7. Theorem 2 indicates that the source detection procedure in Algorithm 2 can consistently detect the underlying transferable set \mathcal{A}_η .

4. Simulation

In this section, we conduct numerical experiments to evaluate the empirical performance of the proposed transfer learning methods in comparison with some other alternatives. Specifically, we consider the following methods, including L1-SQR (ℓ_1 -SQR with only target data), Oracle-TSQR (Alg. 1), Oracle-TQR (Alg. 1 with both the transferring and debiasing steps solved by ℓ_1 -QR), Naive-TSQR (Alg. 1 with all sources) and TSQR (Alg. 2).

4.1. Data generation

We consider a target study with $n_0 = 150$ and $K = 20$ source studies, each with $n_1, \dots, n_K = 100$. The dimension $p = 500$ for both the target and the source data. We consider the following generative model for $k = 0, \dots, K$:

$$y^{(k)} = 0.5 + (\mathbf{x}^{(k)})^\top \mathbf{w}^{(k)} + \epsilon^{(k)},$$

where $\mathbf{x}^{(0)} \sim N_p(\mathbf{0}, \Sigma)$ with $\Sigma = (0.7^{|i-j|})_{1 \leq i, j \leq p}$ and $\mathbf{x}^{(k)} \sim \mathcal{N}(\mathbf{0}_p, \Sigma + \epsilon \epsilon^\top)$ with $\epsilon \sim \mathcal{N}(\mathbf{0}_p, \delta^2 \mathbf{I}_p)$ for $k = 1, \dots, K$. We consider $\delta = 0.3$ here. See Section S5.1 of the Supplementary Material for additional simulations under more heterogeneous designs with larger δ 's. We generate $\epsilon^{(k)}$ from two different distributions: (i) standard normal distribution, $\epsilon^{(k)} \sim \mathcal{N}(0, 1)$; and (ii) t distribution with degrees of freedom 3, $\epsilon^{(k)} \sim t(3)$.

Now we specify the coefficients in each study. For the target, we set $s = 16$, $w_j^{(0)} = 0.5 \mathbf{I}\{j \in [s]\}$. Denote $r_i^{(k)}$ as an independent Rademacher random variable (taking values in $\{1, -1\}$ with equal probability). For the source studies, we set

$$w_j^{(k)} = \begin{cases} w_j^{(0)} + \eta/200 * r_j^{(k)} \mathbf{I}\{j \in H^{(k)}\}, & k \in \mathcal{A}_\eta, \\ \eta/100 * r_j^{(k)} \mathbf{I}\{j \in H^{(k)} \cup [s]\}, & k \in \mathcal{A}_\eta^c, \end{cases}$$

where $H^{(k)}$ is a random subset of $\{s+1, \dots, p\}$ with $|H^{(k)}| = 200$.

We consider $\eta \in \{5, 10, 15\}$ for Gaussian errors and $\eta \in \{10, 20, 30\}$ for t errors. We vary $A = |\mathcal{A}_\eta| \in \{0, 4, 8, \dots, 20\}$ and repeat simulations 100 times. For smoother QR, we use the Gaussian kernel for smoothing with regularization parameter selected by five-fold cross-validation (CV) and choose the bandwidths as $\max[0.05, \sqrt{\tau(1-\tau)}\{\log(p)/n\}^{1/4}]$ as recommended in Tan, Wang and Zhou (2022) for a specific n in the corresponding problem. It needs to be emphasized that this particular choice of bandwidths is by no means optimal numerically. One can also select the optimal bandwidths by CV. We have also noted that the performance of our proposed estimators is not sensitive to the choice of the bandwidths, see Section S2.2 of the Supplementary Material for additional simulation results about the numerical sensitivity to the bandwidths.

We set the threshold $t = 0.2$ for the TQR estimator for simplicity although one can still choose t by CV. We found that the performance of TQR is not

sensitive to the choice of t as long as t falls into a proper interval. We also tried $t = 0.05, 0.1, 0.15$ and the results are similar.

For non-smoothed QR, the regularization parameter is selected by five-fold CV when $A = 0$ and BIC when $A \geq 4$ respectively, since CV in ℓ_1 -QR greatly increases the computational burden as the sample size increases. BIC leads to similar performance as that of CV when the sample size is relatively large while spending much less time than CV. We evaluate the performance of each method in terms of the estimation error $\|\hat{\beta} - \beta\|_2^2$ as well as the computation time.

4.2. Simulation results

Figure 1 and Figure 2 provide a comparison of ℓ_2 -estimation errors of the five methods mentioned above under Gaussian errors and t errors respectively. The average computation time of one replication for the five methods under Gaussian errors is reported in Figure 3, where we take an average of over 100 replications and various (τ, h) for each method and each A . We highlight four conclusions in the following.

Firstly, as we expect, Oracle-TSQR enjoys the best performance across all scenarios as it transfers knowledge exactly from the ideal transferable set \mathcal{A}_η . Moreover, smaller η and bigger A produce smaller estimation errors. In contrast, they do not affect the performance of the single-task ℓ_1 -SQR. This corroborates our main message that one can benefit from transfer learning, which is consistent with results in Theorem 1.

Secondly, the worse performance of Naive-TSQR compared to the ℓ_1 -SQR when the size of the transferable set A is small confirms our concern about the negative transfer. A larger η leads to more severe negative transfer when A is small and therefore a larger A is needed for Naive-TSQR to perform better than ℓ_1 -SQR. It also demonstrates the necessity of the source detection procedure in our Trans-SQR algorithm.

Thirdly, our proposed TSQR almost nearly matches the oracle estimation errors obtained by Oracle-TSQR. This supports our theory on source detection consistency in Theorem 2.

To illustrate the benefit of smoothing, we note in Figure 3 that oracle-TSQR and even TSQR have an overwhelming advantage in computational time compared to that of the non-smooth oracle-TQR especially when the sample size of transferable sources is large, even when we have already used BIC to select the regularization parameter for oracle-TQR in these cases. In addition, we also note that Oracle-TSQR performs slightly better than Oracle-TQR in terms of the estimation error when the size of the transferable set \mathcal{A}_η is relatively small. This phenomenon is also noted in Tan, Wang and Zhou (2022) thanks to the strong convexity brought by convolution smoothing.

Generally speaking, the above conclusions on the performance of different methods apply when the errors follow a t distribution.

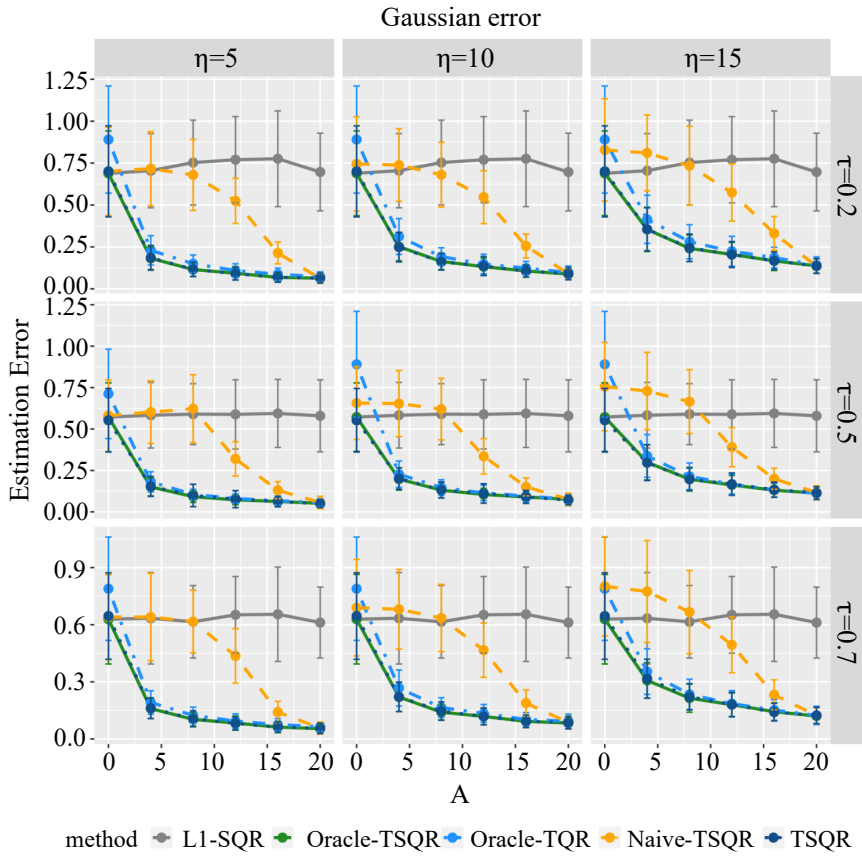


Figure 1. ℓ_2 estimation errors of various methods with Gaussian errors under quantile levels $\tau = 0.2, 0.5, 0.7$, averaged over 100 replications.

5. Application

Primary familial brain calcification (PFBC) is a rare, genetically dominant, inherited neurological disorder characterized by bilateral calcifications in the basal ganglia and other brain regions, and commonly presents motor, psychiatric, and cognitive symptoms. Recent studies have identified JAM2 as a novel causative gene of autosomal recessive PFBC (Cen et al., 2019; Schottlaender et al., 2020). JAM2 encodes junctional adhesion molecule 2, which is highly expressed in neurovascular unit-related cell types (endothelial cells and astrocytes) and is predominantly localized on the plasma membrane. It may be important in cell-cell adhesion and maintaining homeostasis in the central nervous system (CNS). Schottlaender et al. (2020) show that JAM2 variants lead to reduction of JAM2 mRNA expression and absence of JAM2 protein in patient’s fibroblasts, consistent with a loss-of-function mechanism. Therefore, it is of great interest in predicting the expression level of gene JAM2 in target brain tissues, especially at the lower quantile levels.

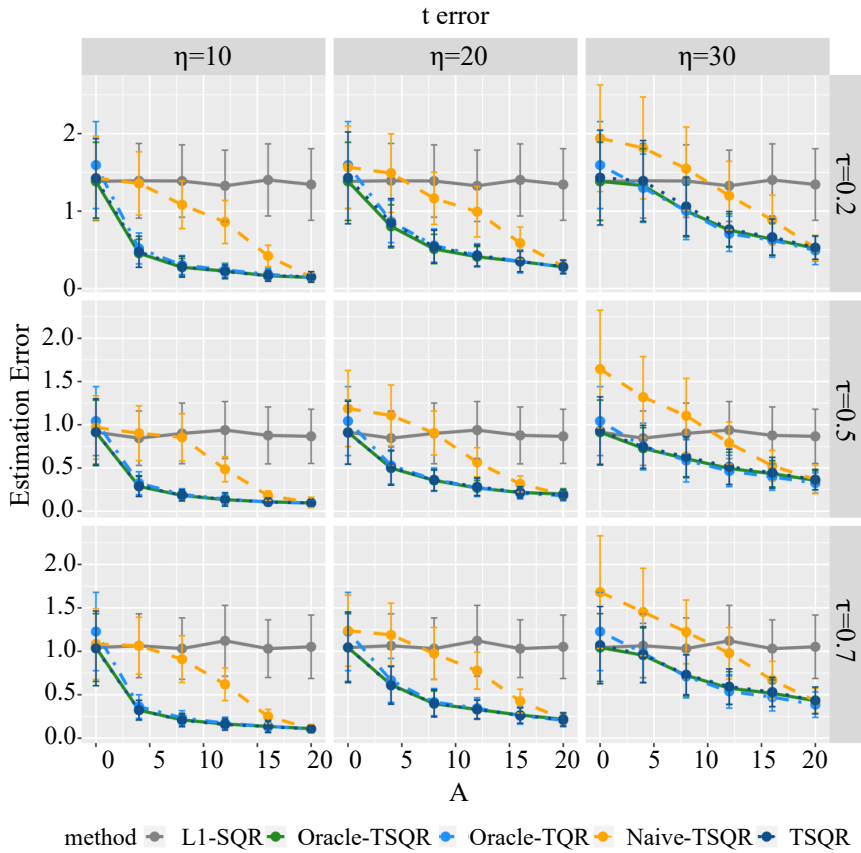


Figure 2. ℓ_2 estimation errors of various methods with t errors under quantile levels $\tau = 0.2, 0.5, 0.7$, averaged over 100 replications.

We apply the proposed transfer learning algorithm to the Genotype-Tissue Expression (GTEx) data, available at <https://gtexportal.org/>. This dataset contains gene expression levels from 49 tissues of 838 individuals, a total of 1,207,976 observations of 38,187 genes. We study the CNS gene regulations in different tissues. The CNS-related genes were assembled as MODULE_137, which includes 545 genes in total as well as 1,632 additional genes that are significantly enriched in the same experiments as the genes of the module (see http://robotics.stanford.edu/erans/cancer/modules/module_137 for a detailed description of this module).

Specifically, we are interested in predicting the expression level of gene JAM2 in a target tissue using other CNS genes. We consider 13 brain tissues as our target tissues. The corresponding models are estimated one by one. The average sample sizes of the target dataset and source datasets are 177 and 14,837 respectively. For each model, we split the target dataset into five folds with each fold being predicted using the remaining four folds as the training data. The

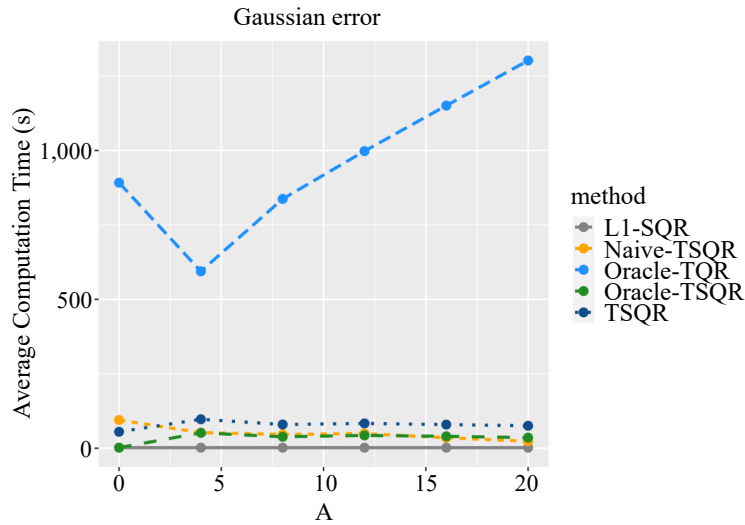


Figure 3. Computation time (in seconds) for each method and each A under Gaussian errors, averaged over 100 replications and various (τ, h) .

regularization parameters are selected by five-fold cross-validation.

Figure 4 reports the relative prediction errors of the Naive-TSQR and our proposed TSQR to the ℓ_1 -SQR estimator based solely on the target data. As we expected, transferring-based methods reduce prediction errors in most cases. In addition, the Naive-TSQR without the source detection procedure performs worse than L1-SQR on the Hemisphere and Cerebellum tissues, while our proposed TSQR reflects its robustness to avoid including unrelated sources across all scenarios. In the best scenario, TSQR reduces the prediction error by about 40%, see the prediction of the JAM2 expression at the 0.7-th quantile on the S.C.cervical tissue for example.

Another interesting finding is that in comparison to the analogous empirical analysis results in Li, Cai and Li (2022a), where transfer learning improves the prediction error by about 30% in the mean regression world for the Hippocampus tissue, here in the quantile regression world, we find that the inclusion of sources seems to be less useful at quantile levels 0.5 and 0.7 but becomes quietly helpful at the lower quantile level 0.2. This indicates that the similarity between the target and the sources may vary with the distribution of our responses. Since the expression level of JAM2 at the lower quantile is of more interest to us, our QR model for transfer learning is more flexible than that of the mean regression.

6. Discussion

In this article, we study the high-dimensional quantile regression problem under the transfer learning setting. We propose a two-step framework for QR




Figure 4. Prediction errors of the Naive-TSQR and TSQR relative to the L1-SQR for the expression level of gene *JAM2* in 13 brain tissues under quantile levels 0, 2, 0.5 and 0.7, evaluated via fivefold cross-validation. The black horizontal line here represents ratio 1.

transfer via convolution smoothing as well as a clustering-based source detection procedure to avoid negative transfer. Numerical experiments as well as an empirical study about GTEx data validate the effectiveness of our proposed method.

There are several avenues for future research. Firstly, there may exist subgroup structures in the target dataset, which may be exploited for precise transfer learning. Specifically, we may consider the target parameter β_g with $g \in \{1, \dots, G\}$ indicating the subgroup and borrow information from other sources with a similar group structure. Secondly, transfer learning of datasets with more complex structures, like matrices or tensors, may be of interest as well.

Supplementary Material

The Supplementary Material contains all technical proofs, additional simulation results about the sensitivity of our proposed estimators to more heterogeneous designs and the smoothing parameters, and a distributed version of the Oracle-Trans-SQR algorithm.

Acknowledgments

We sincerely thank the editor, associate editor, and two anonymous reviewers for their insightful comments, which have greatly enhanced the quality of our paper. This work is partially supported by the National Natural Science Foundation of China grants 12331009, 12071087, and 12226308.

References

- Bastani, H. (2021). Predicting with proxies: Transfer learning in high dimension. *Management Science* **67**, 2964–2984.
- Belloni, A. and Chernozhukov, V. (2011). ℓ_1 -penalized quantile regression in high-dimensional sparse models. *The Annals of Statistics* **39**, 82–130.
- Bickel, P. J., Ritov, Y. and Tsybakov, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics* **37**, 1705–1732.
- Cai, T. T. and Wei, H. (2021). Transfer learning for nonparametric classification: Minimax rate and adaptive classifier. *The Annals of Statistics* **49**, 100–128.
- Cen, Z., Chen, Y., Chen, S., Wang, H., Yang, D., Zhang, H. et al. (2019). Biallelic loss-of-function mutations in JAM2 cause primary familial brain calcification. *Brain* **143**, 491–502.
- Dai, G., Müller, U. U. and Carroll, R. J. (2023). Data integration in high dimension with multiple quantiles. *Statistica Sinica* **33**, 169–191.
- Delaigle, A. and Meister, A. (2007). Nonparametric regression estimation in the heteroscedastic errors-in-variables problem. *Journal of the American Statistical Association* **102**, 1416–1426.
- Delaigle, A. and Meister, A. (2008). Density estimation with heteroscedastic error. *Bernoulli* **14**, 562–579.
- Demko, S. (1977). Inverses of band matrices and local convergence of spline projections. *SIAM Journal on Numerical Analysis* **14**, 616–619.
- Eaton, E., Desjardins, M. and Lane, T. (2008). Modeling transfer relationships between learning tasks for improved inductive transfer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 317–332. Springer.
- Fan, J., Li, Q. and Wang, Y. (2017). Estimation of high dimensional mean regression in the absence of symmetry and light tail assumptions. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **79**, 247–265.
- Fan, J., Xue, L. and Zou, H. (2016). Multitask quantile regression under the transnormal model. *Journal of the American Statistical Association* **111**, 1726–1735.
- Fernandes, M., Guerre, E. and Horta, E. (2021). Smoothing quantile regressions. *Journal of Business & Economic Statistics* **39**, 338–357.
- Ge, L., Gao, J., Ngo, H., Li, K. and Zhang, A. (2014). On handling negative transfer and imbalanced distributions in multiple source transfer learning. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **7**, 254–271.
- He, X., Pan, X., Tan, K. M. and Zhou, W. (2023). Smoothed quantile regression with large-scale inference. *Journal of Econometrics* **232**, 367–388.
- He, Y., Li, Q., Hu, Q. and Liu, L. (2022). Transfer learning in high-dimensional semiparametric graphical models with application to brain connectivity analysis. *Statistics in Medicine* **41**, 4112–4129.

- Hirooka, K., Hasan, M. A. M., Shin, J. and Srizon, A. Y. (2022). Ensembled transfer learning based multichannel attention networks for human activity recognition in still images. *IEEE Access* **10**, 47051–47062.
- Koenker, R. (2017). Quantile regression: 40 years on. *Annual Review of Economics* **9**, 155–176.
- Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica: Journal of the Econometric Society* **46**, 33–50.
- Li, S., Cai, T. T. and Li, H. (2022a). Transfer learning for high-dimensional linear regression: Prediction, estimation and minimax optimality. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **84**, 149–173.
- Li, S., Cai, T. T. and Li, H. (2022b). Transfer learning in large-scale Gaussian graphical models with false discovery rate control. *Journal of the American Statistical Association* **118**, 2171–2183.
- Li, W. (2008). The infinity norm bound for the inverse of nonsingular diagonal dominant matrices. *Applied Mathematics Letters* **21**, 258–263.
- Lin, H. and Reimherr, M. (2024). On hypothesis transfer learning of functional linear models. In *Proceedings of the 41st International Conference on Machine Learning PMLR* **235**, 30252–30285.
- Lu, X. and Su, L. (2015). Jackknife model averaging for quantile regressions. *Journal of Econometrics* **188**, 40–58.
- Niu, S., Liu, Y., Wang, J. and Song, H. (2020). A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence* **1**, 151–166.
- Pan, S. J. and Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* **22**, 1345–1359.
- Peng, B. and Wang, L. (2015). An iterative coordinate descent algorithm for high-dimensional nonconvex penalized quantile regression. *Journal of Computational and Graphical Statistics* **24**, 676–694.
- Schottlaender, L. V., Abeti, R., Jaunmuktane, Z., Macmillan, C., Chelban, V., O’callaghan, B. et al. (2020). Bi-allelic JAM2 variants lead to early-onset recessive primary familial brain calcification. *The American Journal of Human Genetics* **106**, 412–421.
- Sun, Q., Zhou, W. and Fan, J. (2020). Adaptive huber regression. *Journal of the American Statistical Association* **115**, 254–265.
- Tan, K. M., Wang, L. and Zhou, W. (2022). High-dimensional quantile regression: Convolution smoothing and concave regularization. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **84**, 205–233.
- Tian, Y. and Feng, Y. (2023). Transfer learning under high-dimensional generalized linear models. *Journal of the American Statistical Association* **118**, 2684–2697.
- Torrey, L. and Shavlik, J. (2010). Transfer learning. In *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, 242–264. IGI Global.
- Turki, T., Wei, Z. and Wang, J. T. (2017). Transfer learning approaches to improve drug sensitivity prediction in multiple myeloma patients. *IEEE Access* **5**, 7381–7393.
- Wang, L., Wu, Y. and Li, R. (2012). Quantile regression for analyzing heterogeneity in ultra-high dimension. *Journal of the American Statistical Association* **107**, 214–222.
- Wang, M., Zhang, X., Wan, A. T., You, K. and Zou, G. (2023). Jackknife model averaging for high-dimensional quantile regression. *Biometrics* **79**, 178–189.
- Wang, Z., Qin, Z., Tang, X., Ye, J. and Zhu, H. (2018). Deep reinforcement learning with knowledge transfer for online rides order dispatching. In *2018 IEEE International Conference on Data Mining (ICDM)*, 617–626. IEEE.

- Weiss, K., Khoshgoftaar, T. M. and Wang, D. (2016). A survey of transfer learning. *Journal of Big Data* **3**, 1–40.
- Yu, X., Wang, J., Hong, Q., Teku, R., Wang, S. and Zhang, Y. (2022). Transfer learning for medical images analyses: A survey. *Neurocomputing* **489**, 230–254.
- Zhang, W., Deng, L., Zhang, L. and Wu, D. (2022). A survey on negative transfer. *IEEE/CAA Journal of Automatica Sinica* **10**, 305–329.

(Received November 2022; accepted June 2023)