




Transfer Learning Under High-Dimensional Generalized Linear Models

Ye Tian & Yang Feng


To cite this article: Ye Tian & Yang Feng (2023) Transfer Learning Under High-Dimensional Generalized Linear Models, Journal of the American Statistical Association, 118:544, 2684-2697, DOI: [10.1080/01621459.2022.2071278](https://doi.org/10.1080/01621459.2022.2071278)

To link to this article: <https://doi.org/10.1080/01621459.2022.2071278>

 View supplementary material [↗](#)

 Published online: 27 Jun 2022.

 Submit your article to this journal [↗](#)

 Article views: 5931

 View related articles [↗](#)

 View Crossmark data [↗](#)

 Citing articles: 28 View citing articles [↗](#)



Transfer Learning Under High-Dimensional Generalized Linear Models

Ye Tian^a and Yang Feng^b 

^aDepartment of Statistics, Columbia University, New York, NY; ^bDepartment of Biostatistics, School of Global Public Health, New York University, New York, NY

ABSTRACT

In this work, we study the transfer learning problem under high-dimensional generalized linear models (GLMs), which aim to improve the fit on *target* data by borrowing information from useful *source* data. Given which sources to transfer, we propose a transfer learning algorithm on GLM, and derive its ℓ_1/ℓ_2 -estimation error bounds as well as a bound for a prediction error measure. The theoretical analysis shows that when the target and sources are sufficiently close to each other, these bounds could be improved over those of the classical penalized estimator using only target data under mild conditions. When we don't know which sources to transfer, an *algorithm-free* transferable source detection approach is introduced to detect informative sources. The detection consistency is proved under the high-dimensional GLM transfer learning setting. We also propose an algorithm to construct confidence intervals of each coefficient component, and the corresponding theories are provided. Extensive simulations and a real-data experiment verify the effectiveness of our algorithms. We implement the proposed GLM transfer learning algorithms in a new R package `glmtrans`, which is available on CRAN. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received May 2021
Accepted April 2022

KEYWORDS

Generalized linear models;
High-dimensional inference;
Lasso; Negative transfer;
Sparsity; Transfer learning

1. Introduction

Nowadays, a great deal of machine learning algorithms has been successfully applied in our daily life. Many of these algorithms require sufficient training data to perform well, which sometimes can be limited. For example, from an online merchant's view, it could be difficult to collect enough personal purchase data for predicting the customers' purchase behavior and recommending corresponding items. However, in many cases, some related datasets may be available in addition to the limited data for the original task. In the merchant-customer example, we may also have the customers' clicking data in hand, which is not exactly the same as but shares similarities with the purchase data. How to use these additional data to help with the original target task motivates a well-known concept in computer science: *transfer learning* (Torrey and Shavlik 2010; Weiss, Khoshgoftaar, and Wang 2016). As its name indicates, in a transfer learning problem, we aim to transfer some useful information from similar tasks (*sources*) to the original task (*target*), in order to boost the performance on the target. To date, transfer learning has been widely applied in a number of machine learning applications, including the customer review classification (Pan and Yang 2009), medical diagnosis (Hajiramezani and Zamani 2018), and ride dispatching in ride-sharing platforms (Wang et al. 2018), etc. Compared with the rapidly growing applications, there has been little discussion about the theoretical guarantee of transfer learning. Besides, although transfer learning has been prevailing in computer

science community for decades, far less attention has been paid to it among statisticians. More specifically, transfer learning can be promising in the *high-dimensional* data analysis, where the sample size is much less than the dimension with some sparsity structure in the data (Tibshirani 1996). The impact of transfer learning in high-dimensional generalized linear models (GLMs) with sparsity structure is not quite clear up to now. In this article, we are trying to fill the gap by developing transfer learning tools in high-dimensional GLM inference problem, and providing corresponding theoretical guarantees.

Prior to our article, there are a few pioneering works exploring transfer learning under the high-dimensional setting. Bastani (2021) studied the single-source case when the target data comes from a high-dimensional GLM with limited sample size while the source data size is sufficiently large than the dimension. A two-step transfer learning algorithm was developed, and the estimation error bound was derived when the contrast between target and source coefficients is ℓ_0 -sparse. Li, Cai, and Li (2021) further explored the multi-source high-dimensional linear regression problem where both target and source samples are high-dimensional. The ℓ_2 -estimation error bound under ℓ_q -regularization ($q \in [0, 1]$) was derived and proved to be minimax optimal under some conditions. In Li, Cai, and Li (2020), the analysis was extended to the Gaussian graphical models with false discovery rate control. Other related research on transfer learning with theoretical guarantee includes the nonparametric classification model (Cai and Wei 2021; Reeve, Cannings, and Samworth 2021) and the analysis under general

functional classes via transfer exponents (Hanneke and Kpotufe 2020a, 2020b), etc. In addition, during the past few years, there have been some related works studying parameter sharing under the regression setting. For instance, Chen, Owen, and Shi (2015) and Zheng et al. (2019) developed the so-called “data enriched model” for linear and logistic regression under a single-source setting, where the properties of the oracle tuned estimator with a quadratic penalty were studied. Gross and Tibshirani (2016) and Ollier and Viallon (2017) explored the so-called “data shared Lasso” under the multi-task learning setting, where ℓ_1 penalties of all contrasts are considered.

In this work, we contribute to transfer learning under a high-dimensional context from three perspectives. First, we extend the results of Bastani (2021) and Li, Cai, and Li (2021), by proposing multi-source transfer learning algorithms on generalized linear models (GLMs) and we assume both target and source data to be high-dimensional. We assume the contrast between target and source coefficients to be ℓ_1 -sparse, which differs from the ℓ_0 -sparsity considered in Bastani (2021). The theoretical analysis shows that when the target and sources are sufficiently close to each other, the estimation error bound of target coefficients could be improved over that of the classical penalized estimator using only target data under mild conditions. Moreover, the error rate is shown to be minimax optimal under certain conditions. To the best of our knowledge, this is the initial study of the multi-source transfer learning framework under the high-dimensional GLM setting. Second, as we mentioned, transferring sources that are close to the target can bring benefits. However, some sources might be far away from the target, and transferring them can be harmful. This phenomenon is often called *negative transfer* in literature (Torrey and Shavlik 2010). We will show the impact of negative transfer in simulation studies in Section 4.1. To avoid this issue, we develop an *algorithm-free* transferable source detection algorithm, which can help identify informative sources. And with certain conditions satisfied, the algorithm is shown to be able to distinguish useful sources from useless ones. Third, all existing works of transfer learning on high-dimensional regression only focus on the point estimate of the coefficient, which is not sufficient for statistical inference. How transfer learning can benefit the confidence interval construction remains unclear. We propose an algorithm on the basis of our two-step transfer learning procedure and nodewise regression (Van de Geer et al. 2014), to construct the confidence interval for each coefficient component. The corresponding asymptotic theories are established.

The rest of this article is organized as follows. Section 2 first introduces GLM basics and transfer learning settings under high-dimensional GLM, then presents a general algorithm (where we know which sources are useful) and the transferable source detection algorithm (where useful sources are automatically detected). At the end of Section 2, we develop an algorithm to construct confidence intervals. Section 3 provides the theoretical analysis on the algorithms, including ℓ_1 and ℓ_2 -estimation error bounds of the general algorithm, detection consistency property of the transferable source detection algorithm, and asymptotic theories for the confidence interval construction. We conduct extensive simulations and a real-data study in Section 4, and the results demonstrate the effectiveness

of our GLM transfer learning algorithms. In Section 5, we review our contributions and shed light on some interesting future research directions. Additional simulation results and theoretical analysis, as well as all the proofs, are relegated to supplementary materials.

2. Methodology

We first introduce some notations to be used throughout the article. 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 $\|\mathbf{x}\| = (\sum_{i=1}^p |x_i|^q)^{1/q}$ ($q \in (0, 2]$), and ℓ_0 -“norm” as $\|\mathbf{x}\|_0 = \#\{j : x_j \neq 0\}$. For a matrix $\mathbf{A}_{p \times q} = [a_{ij}]_{p \times q}$, its 1-norm, 2-norm, ∞ -norm, and max-norm are defined as $\|\mathbf{A}\|_1 = \sup_j \sum_{i=1}^p |a_{ij}|$, $\|\mathbf{A}\|_2 = \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$, $\|\mathbf{A}\|_\infty = \sup_i \sum_{j=1}^q |a_{ij}|$, and $\|\mathbf{A}\|_{\max} = \sup_{i,j} |a_{ij}|$, respectively. For two nonzero real sequences $\{a_n\}_{n=1}^\infty$ and $\{b_n\}_{n=1}^\infty$, we use $a_n \ll b_n$, $b_n \gg a_n$ or $a_n = o(b_n)$ to represent $|a_n/b_n| \rightarrow 0$, where we defined the new command “ o ” at the beginning to the text file as follows: as $n \rightarrow \infty$. And $a_n \lesssim b_n$ or $a_n = O(b_n)$ means $\sup_n |a_n/b_n| < \infty$. Expression $a_n \asymp b_n$ means that a_n/b_n converges to some positive constant. For two random variable sequences $\{x_n\}_{n=1}^\infty$ and $\{y_n\}_{n=1}^\infty$, notation $x_n \lesssim_p y_n$ or $x_n = O_p(y_n)$ means that for any $\epsilon > 0$, there exists a positive constant M such that $\sup_n \mathbb{P}(|x_n/y_n| > M) \leq \epsilon$. And for two real numbers a and b , we use $a \vee b$ and $a \wedge b$ to represent $\max(a, b)$ and $\min(a, b)$, respectively. Without specific notes, the expectation \mathbb{E} , variance var , and covariance cov are calculated based on all randomness.

2.1. Generalized Linear Models (GLMs)

Given the predictors $\mathbf{x} \in \mathbb{R}^p$, if the response y follows the generalized linear models (GLMs), then its conditional distribution takes the form

$$y|\mathbf{x} \sim \mathbb{P}(y|\mathbf{x}) = \rho(y) \exp\{y\mathbf{x}^T \mathbf{w} - \psi(\mathbf{x}^T \mathbf{w})\},$$

where $\mathbf{w} \in \mathbb{R}^p$ is the coefficient, ρ and ψ are some known univariate functions. $\psi'(\mathbf{x}^T \mathbf{w}) = \mathbb{E}(y|\mathbf{x})$ is called the *inverse link function* (McCullagh and Nelder 1989). Another important property is that $\text{var}(y|\mathbf{x}) = \psi''(\mathbf{x}^T \mathbf{w})$, which follows from the fact that the distribution belongs to the exponential family. It is ψ that characterizes different GLMs. For example, in linear model with Gaussian noise, we have a continuous response y and $\psi(u) = \frac{1}{2}u^2$; in the logistic regression model, y is binary and $\psi(u) = \log(1 + e^u)$; and in Poisson regression model, y is a nonnegative integer and $\psi(u) = e^u$. For most GLMs, ψ is strictly convex and infinitely differentiable.

2.2. Target Data, Source Data, and Transferring Level

In this article, we consider the following multi-source transfer learning problem. Suppose we have the *target* dataset $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ and K *source* datasets with the k th source denoted as $(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})$, where $\mathbf{X}^{(k)} \in \mathbb{R}^{n_k \times p}$, $\mathbf{y}^{(k)} \in \mathbb{R}^{n_k}$ for $k = 0, \dots, K$. The i th row of $\mathbf{X}^{(k)}$ and the i th element of $\mathbf{y}^{(k)}$ are denoted as $\mathbf{x}_i^{(k)}$ and

$y_i^{(k)}$, respectively. The goal is to transfer useful information from source data to obtain a better model for the target data. We assume the responses in the target and source data all follow the generalized linear model:

$$y^{(k)}|\mathbf{x} \sim \mathbb{P}(y|\mathbf{x}) = \rho(y) \exp\{y\mathbf{x}^T \mathbf{w}^{(k)} - \psi(\mathbf{x}^T \mathbf{w}^{(k)})\}, \quad (1)$$

for $k = 0, \dots, K$, with possibly different coefficient $\mathbf{w}^{(k)} \in \mathbb{R}^p$, the predictor $\mathbf{x} \in \mathbb{R}^p$, and some known univariate functions ρ and ψ . Denote the target parameter as $\boldsymbol{\beta} = \mathbf{w}^{(0)}$. Suppose the target model is ℓ_0 -sparse, which satisfies $\|\boldsymbol{\beta}\|_0 = s \ll p$. This means that only s of the p variables contribute to the target response. Intuitively, if $\mathbf{w}^{(k)}$ is close to $\boldsymbol{\beta}$, the k th source could be useful for transfer learning.

Define the k th contrast $\boldsymbol{\delta}^{(k)} = \boldsymbol{\beta} - \mathbf{w}^{(k)}$. Given some $\mathcal{A} \subseteq \{1, \dots, K\}$, we denote $h = \max_{k \in \mathcal{A}} \|\boldsymbol{\delta}^{(k)}\|$ and we call this h as the *transferring level* of \mathcal{A} . We denote this index set \mathcal{A} as \mathcal{A}_h to emphasize its associated transferring level. Note that in general, h can be any positive value. However, in our regime of interest, h shall be reasonably small to guarantee that transferring sources in \mathcal{A}_h is beneficial. Denote $n_{\mathcal{A}_h} = \sum_{k \in \mathcal{A}_h} n_k$, $\alpha_k = \frac{n_k}{n_{\mathcal{A}_h} + n_0}$ for $k \in \{0\} \cup \mathcal{A}_h$ and $K_{\mathcal{A}_h} = |\mathcal{A}_h|$.

Note that in (1), we assume GLMs of the target and all sources share the same inverse link function ψ . After a careful examination of our proofs for theoretical properties in Section 3, we find that these theoretical results still hold even when the target and each source have their own function ψ , as long as these GLMs satisfy Assumptions 1 and 3 (to be presented in Section 3.1). It means that transferring information across different GLM families is possible. For simplicity, in the following discussion, we assume all these GLMs belong to the same family and hence have the same function ψ .

2.3. Two-Step GLM Transfer Learning

We first introduce a general transfer learning algorithm on GLMs, which can be applied to transfer all sources in a given index set \mathcal{A} . The algorithm is motivated by the ideas in Bastani (2021) and Li, Cai, and Li (2021), which we call a *two-step transfer learning algorithm*. The main strategy is to first transfer the information from those sources by pooling all the data to obtain a rough estimator, then correct the bias in the second step using the target data. More specifically, we fit a GLM with ℓ_1 -penalty by pooled samples first, then fit the contrast in the second step using only the target by another ℓ_1 -regularization. The detailed algorithm (\mathcal{A} -Trans-GLM) is presented in Algorithm 1. The transferring step could be understood as to solve the following equation w.r.t. $\mathbf{w} \in \mathbb{R}^p$:

$$\sum_{k \in \{0\} \cup \mathcal{A}} \left[(\mathbf{X}^{(k)})^T \mathbf{y}^{(k)} - \sum_{i=1}^{n_k} \psi'((\mathbf{w})^T \mathbf{x}_i^{(k)}) \mathbf{x}_i^{(k)} \right] = \mathbf{0}_p,$$

which converges to the solution of its population version under certain conditions

$$\sum_{k \in \{0\} \cup \mathcal{A}} \alpha_k \mathbb{E} \left\{ [\psi'((\mathbf{w}^A)^T \mathbf{x}^{(k)}) - \psi'((\mathbf{w}^{(k)})^T \mathbf{x}^{(k)})] \mathbf{x}^{(k)} \right\} = \mathbf{0}_p, \quad (2)$$

where $\alpha_k = \frac{n_k}{n_{\mathcal{A}} + n_0}$. Notice that in the linear case, \mathbf{w}^A can be explicitly expressed as a linear transformation of the true parameter $\mathbf{w}^{(k)}$, that is, $\mathbf{w}^A = \boldsymbol{\Sigma}^{-1} \sum_{k \in \{0\} \cup \mathcal{A}} \alpha_k \boldsymbol{\Sigma}^{(k)} \mathbf{w}^{(k)}$, where

$\boldsymbol{\Sigma}^{(k)} = \mathbb{E}[\mathbf{x}^{(k)}(\mathbf{x}^{(k)})^T]$ and $\boldsymbol{\Sigma} = \sum_{k \in \{0\} \cup \mathcal{A}} \alpha_k \mathbb{E}[\mathbf{x}^{(k)}(\mathbf{x}^{(k)})^T]$ (Li, Cai, and Li 2021).

To help readers better understand the algorithm, we draw a schematic in Section S.1.1 of supplementary materials. We refer interested readers who want to get more intuitions to that.

Algorithm 1: \mathcal{A} -Trans-GLM

Input: target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, source data $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k=1}^K$, penalty parameters $\lambda_{\mathbf{w}}$ and $\lambda_{\boldsymbol{\delta}}$, transferring set \mathcal{A}

Output: the estimated coefficient vector $\hat{\boldsymbol{\beta}}$

1 **Transferring step:** Compute $\hat{\mathbf{w}}^{\mathcal{A}} \leftarrow$

$$\arg \min_{\mathbf{w}} \left\{ \frac{1}{n_{\mathcal{A}} + n_0} \sum_{k \in \{0\} \cup \mathcal{A}} \left[-(\mathbf{y}^{(k)})^T \mathbf{X}^{(k)} \mathbf{w} + \sum_{i=1}^{n_k} \psi(\mathbf{w}^T \mathbf{x}_i^{(k)}) \right] + \lambda_{\mathbf{w}} \|\mathbf{w}\|_1 \right\}$$

2 **Debiasing step:** Compute

$$\hat{\boldsymbol{\delta}}^{\mathcal{A}} \leftarrow \arg \min_{\boldsymbol{\delta}} \left\{ -\frac{1}{n_0} (\mathbf{y}^{(0)})^T \mathbf{X}^{(0)} (\hat{\mathbf{w}}^{\mathcal{A}} + \boldsymbol{\delta}) + \frac{1}{n_0} \sum_{i=1}^{n_0} \psi((\hat{\mathbf{w}}^{\mathcal{A}} + \boldsymbol{\delta})^T \mathbf{x}_i^{(0)}) + \lambda_{\boldsymbol{\delta}} \|\boldsymbol{\delta}\|_1 \right\}$$

3 Let $\hat{\boldsymbol{\beta}} \leftarrow \hat{\mathbf{w}}^{\mathcal{A}} + \hat{\boldsymbol{\delta}}^{\mathcal{A}}$

4 Output $\hat{\boldsymbol{\beta}}$

2.4. Transferable Source Detection

As we described, Algorithm 1 can be applied only if we are certain about which sources to transfer, which in practice may not be known a priori. Transferring certain sources may not improve the performance of the fitted model based on only target, and can even lead to worse performance. In transfer learning, we say *negative transfer* happens when the source data leads to an inferior performance on the target task (Pan and Yang 2009; Torrey and Shavlik 2010; Weiss, Khoshgoftaar, and Wang 2016). How to avoid negative transfer has become an increasingly popular research topic.

Here we propose a simple, *algorithm-free*, and *data-driven* method to determine an informative transferring set $\hat{\mathcal{A}}$. We call this approach a transferable source *detection* algorithm and refer to it as Trans-GLM.

We sketch this detection algorithm as follows. First, divide the target data into three folds, that is, $\{(\mathbf{X}^{(0)[r]}, \mathbf{y}^{(0)[r]})\}_{r=1}^3$. Note that we choose three folds only for convenience. We also explored other fold number choices in the simulation. See Section S.1.3.3 in the supplementary materials. Second, run the transferring step on each source data and every two folds of target data. Then, for a given loss function, we calculate its value on the left-out fold of target data and compute the average cross-validation loss $\hat{L}_0^{(k)}$ for each source. As a benchmark, we also fit Lasso on every choice of two folds of target data and calculate the loss on the remaining fold. The average cross-validation loss $\hat{L}_0^{(0)}$ is viewed as the loss of target. Finally, the difference between $\hat{L}_0^{(k)}$ and $\hat{L}_0^{(0)}$ is calculated and compared with some threshold, and sources with a difference less than the threshold will be recruited into $\hat{\mathcal{A}}$. Under the GLM setting, a natural loss function

is the negative log-likelihood. For convenience, suppose n_0 is divisible by 3. According to (1), for any coefficient estimate \mathbf{w} , the average of negative log-likelihood on the r th fold of target data $(\mathbf{X}^{(0)[r]}, \mathbf{y}^{(0)[r]})$ is

$$\hat{L}_0^{[r]}(\mathbf{w}) = -\frac{1}{n_0/3} \sum_{i=1}^{n_0/3} \log \rho(y_i^{(0)[r]}) - \frac{1}{n_0/3} (\mathbf{y}^{(0)[r]})^T \mathbf{X}^{(0)} \mathbf{w} + \frac{1}{n_0/3} \sum_{i=1}^{n_0/3} \psi(\mathbf{w}^T \mathbf{x}_i^{(0)[r]}). \tag{3}$$

The detailed algorithm is presented as Algorithm 2.

Algorithm 2: Trans-GLM

Input: target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, all source data $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k=1}^K$, a constant $C_0 > 0$, penalty parameters $\{\{\lambda^{(k)[r]}\}_{k=0}^K\}_{r=1}^3$

Output: the estimated coefficient vector $\hat{\boldsymbol{\beta}}$, and the determined transferring set $\hat{\mathcal{A}}$

- 1 **Transferable source detection:** Randomly divide $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ into three sets of equal size as $\{(\mathbf{X}^{(0)[i]}, \mathbf{y}^{(0)[i]})\}_{i=1}^3$
- 2 **for** $r = 1$ **to** 3 **do**
- 3 $\hat{\boldsymbol{\beta}}^{(0)[r]} \leftarrow$ fit the Lasso on $\{(\mathbf{X}^{(0)[i]}, \mathbf{y}^{(0)[i]})\}_{i=1}^3 \setminus (\mathbf{X}^{(0)[r]}, \mathbf{y}^{(0)[r]})$ with penalty parameter $\lambda^{(0)[r]}$
- 4 $\hat{\boldsymbol{\beta}}^{(k)[r]} \leftarrow$ run Step 1 in Algorithm 1 with $(\{(\mathbf{X}^{(0)[i]}, \mathbf{y}^{(0)[i]})\}_{i=1}^3 \setminus (\mathbf{X}^{(0)[r]}, \mathbf{y}^{(0)[r]}) \cup (\mathbf{X}^{(k)}, \mathbf{y}^{(k)})$ and penalty parameter $\lambda^{(k)[r]}$ for all $k \neq 0$
- 5 Calculate the loss function $\hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(k)[r]})$ on $(\mathbf{X}^{(0)[r]}, \mathbf{y}^{(0)[r]})$ for $k = 1, \dots, K$
- 6 **end**
- 7 $\hat{L}_0^{(k)} \leftarrow \sum_{r=1}^3 \hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(k)[r]})/3, \hat{L}_0^{(0)} \leftarrow \sum_{r=1}^3 \hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(0)[r]})/3, \hat{\sigma} = \sqrt{\sum_{r=1}^3 (\hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(0)[r]}) - \hat{L}_0^{(0)})^2/2}$
- 8 $\hat{\mathcal{A}} \leftarrow \{k \neq 0 : \hat{L}_0^{(k)} - \hat{L}_0^{(0)} \leq C_0(\hat{\sigma} \vee 0.01)\}$
- 9 **$\hat{\mathcal{A}}$ -Trans-GLM:** $\hat{\boldsymbol{\beta}} \leftarrow$ run Algorithm 1 using $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k \in \{0\} \cup \hat{\mathcal{A}}}$
- 10 Output $\hat{\boldsymbol{\beta}}$ and $\hat{\mathcal{A}}$

It’s important to point out that Algorithm 2 does not require the input of h . We will show that $\hat{\mathcal{A}} = \mathcal{A}_h$ for some specific h if certain conditions hold, in Section 3.2. Furthermore, under these conditions, transferring with $\hat{\mathcal{A}}$ will lead to a faster convergence rate compared to Lasso fitted on only the target data, when target sample size n_0 falls into some regime. This is the reason that this algorithm is called the *transferable* source detection algorithm.

2.5. Confidence Intervals

In previous sections, we’ve discussed how to obtain a point estimator of the target coefficient vector $\boldsymbol{\beta}$ from the two-step transfer learning approach. In this section, we would like to

construct the asymptotic confidence interval (CI) for each component of $\boldsymbol{\beta}$ based on that point estimate.

As described in the introduction, there have been quite a few works on high-dimensional GLM inference in the literature. In the following, we will propose a transfer learning procedure to construct CI based on the desparsified Lasso (Van de Geer et al. 2014). Recall that desparsified Lasso contains two main steps. The first step is to learn the inverse Fisher information matrix of GLM by nodewise regression (Meinshausen and Bühlmann 2006). The second step is to “debias” the initial point estimator and then construct the asymptotic CI. Here, the estimator $\hat{\boldsymbol{\beta}}$ from Algorithm 1 can be used as an initial point estimator. Intuitively, if the predictors from target and source data are similar and satisfy some sparsity conditions, it might be possible to use Algorithm 1 for learning the inverse Fisher information matrix of target data, which effectively combines the information from target and source data.

Before formalizing the procedure to construct the CI, let’s first define several additional notations. For any $\mathbf{W} \in \times^p$, denote $\mathbf{W}_w^{(k)} = \text{diag}(\sqrt{\psi''((\mathbf{x}_1^{(k)})^T \mathbf{w})}, \dots, \sqrt{\psi''((\mathbf{x}_{n_k}^{(k)})^T \mathbf{w})})$, $\mathbf{X}_w^{(k)} = \mathbf{W}_w^{(k)} \mathbf{X}^{(k)}$, $\boldsymbol{\Sigma}_w^{(k)} = \mathbb{E}[\mathbf{x}^{(k)} (\mathbf{x}^{(k)})^T \psi''((\mathbf{x}^{(k)})^T \mathbf{w})]$ and $\hat{\boldsymbol{\Sigma}}_w^{(k)} = n_k^{-1} (\mathbf{X}_w^{(k)})^T \mathbf{X}^{(k)} \cdot \mathbf{X}_{w,j}^{(k)}$ represents the j th column of $\mathbf{X}_w^{(k)}$ and $\mathbf{X}_{w,-j}^{(k)}$ represents the matrix $\mathbf{X}_w^{(k)}$ without the j th column. $\hat{\boldsymbol{\Sigma}}_{w,j,-j}^{(k)}$ represents the j th row of $\hat{\boldsymbol{\Sigma}}_w^{(k)}$ without the diagonal (j, j) element, and $\hat{\boldsymbol{\Sigma}}_{w,j,j}^{(k)}$ is the diagonal (j, j) element of $\hat{\boldsymbol{\Sigma}}_w^{(k)}$.

Next, we explain the details of the CI construction procedure in Algorithm 3. In Step 1, we obtain a point estimator $\hat{\boldsymbol{\beta}}$ from \mathcal{A} -Trans-GLM (Algorithm 1), given a specific transferring set \mathcal{A} . Then in Steps 2–4, we estimate the target inverse Fisher information matrix $(\boldsymbol{\Sigma}_\beta^{(0)})^{-1}$ as

$$\hat{\Theta} = \text{diag}(\hat{\tau}_1^{-2}, \dots, \hat{\tau}_p^{-2}) \begin{pmatrix} 1 & -\hat{\gamma}_{1,2}^{(0)} & \dots & -\hat{\gamma}_{1,p}^{(0)} \\ -\hat{\gamma}_{2,1}^{(0)} & 1 & \dots & -\hat{\gamma}_{2,p}^{(0)} \\ \vdots & \vdots & \ddots & \vdots \\ -\hat{\gamma}_{p,1}^{(0)} & -\hat{\gamma}_{p,2}^{(0)} & \dots & 1 \end{pmatrix}. \tag{4}$$

Finally in Step 5, we “debias” $\hat{\boldsymbol{\beta}}$ using the target data to get a new point estimator $\hat{\mathbf{b}}$ which is asymptotically unbiased as

$$\hat{\mathbf{b}} = \hat{\boldsymbol{\beta}} + \frac{1}{n_0} \hat{\Theta} (\mathbf{X}^{(0)})^T [\mathbf{Y}^{(0)} - \boldsymbol{\psi}'(\mathbf{X}^{(0)} \hat{\boldsymbol{\beta}})], \tag{5}$$

where $\boldsymbol{\psi}'(\mathbf{X}^{(0)} \hat{\boldsymbol{\beta}}) := (\boldsymbol{\psi}'((\mathbf{x}_1^{(0)})^T \hat{\boldsymbol{\beta}}), \dots, \boldsymbol{\psi}'((\mathbf{x}_{n_0}^{(0)})^T \hat{\boldsymbol{\beta}}))^T \in \mathbb{R}^{n_0}$.

It’s necessary to emphasize that the confidence level $(1 - \alpha)$ is for every single CI rather than for all p CIs simultaneously. As discussed in Sections 2.2 and 2.3 of Van de Geer et al. (2014), it is possible to get simultaneous CIs for different coefficient components and do multiple hypothesis tests when the design is fixed. In other cases, for example, random design in different replications (which we focus on in this article), multiple hypothesis testing might be more challenging.

3. Theory

In this section, we will establish theoretical guarantees on the three proposed algorithms. Section 3.1 provides a detailed anal-

Algorithm 3: Confidence interval construction via node-wise regression

Input: target data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, source data $\{(\mathbf{X}^{(k)}, \mathbf{y}^{(k)})\}_{k=1}^K$, penalty parameters $\{\lambda_j\}_{j=1}^p$ and $\{\tilde{\lambda}_j\}_{j=1}^p$, transferring set \mathcal{A} , confidence level $(1 - \alpha)$

Output: Level- $(1 - \alpha)$ confidence interval \mathcal{I}_j for β_j with $j = 1, \dots, p$

1 Compute $\hat{\boldsymbol{\beta}}$ via Algorithm 1

2 Compute $\hat{\mathbf{y}}_j^{\mathcal{A}} \leftarrow$

$$\arg \min_{\mathbf{y}} \left\{ -\frac{1}{2(n_{\mathcal{A}} + n_0)} \sum_{k \in \{0\} \cup \mathcal{A}} \|\mathbf{X}_{\hat{\boldsymbol{\beta}}, j}^{(k)} - \mathbf{X}_{\hat{\boldsymbol{\beta}}, -j}^{(k)} \mathbf{y}\|_2^2 + \lambda_j \|\mathbf{y}\|_1 \right\}$$

for $j = 1, \dots, p$

3 Compute $\hat{\boldsymbol{\theta}}_j \leftarrow$

$$\arg \min_{\boldsymbol{\theta}} \left\{ -\frac{1}{2n_0} \|\mathbf{X}_{\hat{\boldsymbol{\beta}}, j}^{(0)} - \mathbf{X}_{\hat{\boldsymbol{\beta}}, -j}^{(0)} (\hat{\mathbf{y}}_j^{\mathcal{A}} + \boldsymbol{\theta})\|_2 + \tilde{\lambda}_j \|\boldsymbol{\theta}\|_1 \right\}$$

4 Compute $\hat{\mathbf{y}}_j^{(0)} \leftarrow \hat{\mathbf{y}}_j^{\mathcal{A}} + \hat{\boldsymbol{\theta}}_j$, $\hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}} \leftarrow \sum_{k \in \{0\} \cup \mathcal{A}} \frac{n_k}{n_{\mathcal{A}} + n_0} \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}}^{(k)}$,

$$\hat{\tau}_j^2 = \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}, j, j} - \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}, j, -j} \hat{\mathbf{y}}_j$$

and calculate $\hat{\boldsymbol{\Theta}}$ via (4), where

$$\hat{\mathbf{y}}_j^{(0)} = (\hat{y}_{j,1}^{(0)}, \dots, \hat{y}_{j,j-1}^{(0)}, \hat{y}_{j,j+1}^{(0)}, \dots, \hat{y}_{j,p}^{(0)})^T.$$

5 Compute

$$\mathcal{I}_j \leftarrow [\hat{b}_j - \hat{\boldsymbol{\Theta}}_j^T \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}} \hat{\boldsymbol{\theta}}_j q_{\alpha/2} / \sqrt{n_0}, \hat{b}_j + \hat{\boldsymbol{\Theta}}_j^T \hat{\boldsymbol{\Sigma}}_{\hat{\boldsymbol{\beta}}} \hat{\boldsymbol{\theta}}_j q_{\alpha/2} / \sqrt{n_0}]$$

for $j = 1, \dots, p$, where \hat{b}_j is the j th component of $\hat{\mathbf{b}}$ in (5), and $q_{\alpha/2}$ is the $\alpha/2$ -left tail quantile of $\mathcal{N}(0, 1)$

6 Output $\{\mathcal{I}_j\}_{j=1}^p$

ysis of Algorithm 1 with transferring set \mathcal{A}_h , which we denote as \mathcal{A}_h -Trans-GLM. Section 3.2 introduces certain conditions, under which we show that the transferring set $\hat{\mathcal{A}}$ detected by Algorithm 2 (Trans-GLM) is equal to \mathcal{A}_h for some h with high probability. Section 3.3 presents the analysis of Algorithm 3 with transferring set \mathcal{A}_h , where we prove a central limit theorem. For the proofs and some additional theoretical results, refer to supplementary materials.

3.1. Theory on \mathcal{A}_h -Trans-GLM

We first impose some common assumptions about GLM.

Assumption 1. ψ is infinitely differentiable and strictly convex. We call a second-order differentiable function ψ strictly convex if $\psi''(x) > 0$.

Assumption 2. For any $\mathbf{a} \in \mathbb{R}^p$, $\mathbf{a}^T \mathbf{x}_i^{(k)}$'s are iid $\kappa_u \|\mathbf{a}\|_2^2$ -sub-Gaussian variables with zero mean for all $k = 0, \dots, K$, where κ_u is a positive constant. Denote the covariance matrix of $\mathbf{x}^{(k)}$ as $\boldsymbol{\Sigma}^{(k)}$, with $\inf_k \lambda_{\min}(\boldsymbol{\Sigma}^{(k)}) \geq \kappa_l > 0$, where κ_l is a constant.

Assumption 3. At least one of the following assumptions hold: (M_ψ , U and \bar{U} are some positive constants)

- (i) $\|\psi''\|_\infty \leq M_\psi < \infty$ a.s.;
- (ii) $\sup_k \|\mathbf{x}^{(k)}\|_\infty \leq U < \infty$ a.s., $\sup_k \sup_{|z| \leq \bar{U}} \psi''((\mathbf{x}^{(k)})^T \mathbf{w}^{(k)} + z) \leq M_\psi < \infty$ a.s.

Assumption 1 imposes the *strict convexity* and differentiability of ψ , which is satisfied by many popular distribution families, such as Gaussian, binomial, and Poisson distributions. Note that we do not require ψ to be *strongly convex* (that is, $\exists C > 0$, such that $\psi''(x) > C$), which relaxes Assumption 4 in Bastani (2021). It is easy to show that ψ in logistic regression is in general not strongly convex with unbounded predictors. **Assumption 2** requires the predictors in each source to be sub-Gaussian with a well-behaved correlation structure. **Assumption 3** is motivated by Assumption (GLM 2) in the full-length version of Negahban et al. (2009), which is imposed to restrict ψ'' in a bounded region in some sense. Note that linear regression and logistic regression satisfy condition (i), while Poisson regression with coordinate-wise bounded predictors and ℓ_1 -bounded coefficients satisfies condition (ii).

Besides these common conditions on GLM, as discussed in Section 2.3, to guarantee the success of \mathcal{A}_h -Trans-GLM, we have to make sure that the estimator from the transferring step is close enough to $\boldsymbol{\beta}$. Therefore, we introduce the following assumption, which guarantees $\mathbf{w}^{\mathcal{A}_h}$ defined in (2) with $\mathcal{A} = \mathcal{A}_h$ is close to $\boldsymbol{\beta}$.

Assumption 4. Denote $\tilde{\boldsymbol{\Sigma}}_h = \sum_{k \in \{0\} \cup \mathcal{A}_h} \alpha_k \mathbb{E} \left[\int_0^1 \psi''((\mathbf{x}^{(k)})^T \boldsymbol{\beta} + t(\mathbf{x}^{(k)})^T (\mathbf{w}^{\mathcal{A}_h} - \boldsymbol{\beta})) dt \cdot \mathbf{x}^{(k)} (\mathbf{x}^{(k)})^T \right]$ and $\tilde{\boldsymbol{\Sigma}}_h^{(k)} = \mathbb{E} \left[\int_0^1 \psi''((\mathbf{x}^{(k)})^T \boldsymbol{\beta} + t(\mathbf{x}^{(k)})^T (\mathbf{w}^{(k)} - \boldsymbol{\beta})) dt \cdot \mathbf{x}^{(k)} (\mathbf{x}^{(k)})^T \right]$. It holds that $\sup_{k \in \{0\} \cup \mathcal{A}_h} \|\tilde{\boldsymbol{\Sigma}}_h^{-1} \tilde{\boldsymbol{\Sigma}}_h^{(k)}\|_1 < \infty$.

Remark 1. A sufficient condition for Assumption 4 to hold is $(\tilde{\boldsymbol{\Sigma}}_{\mathbf{w}^{\mathcal{A}_h}, \boldsymbol{\beta}}^{(k)})^{-1} \tilde{\boldsymbol{\Sigma}}_{\mathbf{w}^{(k')}, \boldsymbol{\beta}}^{(k')}$ has positive diagonal elements and is diagonally dominant, for any $k \neq k'$ in \mathcal{A}_h , where $\tilde{\boldsymbol{\Sigma}}_{\mathbf{w}, \boldsymbol{\beta}}^{(k)} := \mathbb{E} \left[\int_0^1 \psi''((\mathbf{x}^{(k)})^T \boldsymbol{\beta} + t(\mathbf{x}^{(k)})^T (\mathbf{w} - \boldsymbol{\beta})) dt \cdot \mathbf{x}^{(k)} (\mathbf{x}^{(k)})^T \right]$ for any $\mathbf{w} \in \mathbb{R}^p$.

In the linear case, this assumption can be further simplified as a restriction on heterogeneity between target predictors and source predictors. More discussions can be found in Condition 4 of Li, Cai, and Li (2021). Now, we are ready to present the following main result for the \mathcal{A}_h -Trans-GLM algorithm. Given some $\mathcal{A} \subseteq \{1, \dots, K\}$, define the parameter space as

$$\Xi(s, h) = \left\{ \boldsymbol{\beta}, \{\mathbf{W}^{(k)}\}_{k=1}^K : \|\boldsymbol{\beta}\|_0 \leq s, \sup_{k \in \mathcal{A}} \|\mathbf{W}^{(k)} - \boldsymbol{\beta}\|_1 \leq h \right\}.$$

Again, we denote the index set \mathcal{A} as \mathcal{A}_h to emphasize its associated transferring level.

Theorem 1 (ℓ_1/ℓ_2 -estimation error bound of \mathcal{A}_h -Trans-GLM with Assumption 4). Assume Assumptions 1, 2, and 4 hold. Suppose $h \ll \sqrt{\frac{n_0}{\log p}}$, $h \leq C$, $n_0 \geq C \log p$ and $n_{\mathcal{A}_h} \geq Cs \log p$, where $C > 0$ is a constant. Also assume Assumption 3.(i) holds or Assumption 3.(ii) with $h \leq C' U^{-1} \bar{U}$ for some $C' > 0$ holds. We take $\lambda_{\mathbf{w}} = C_{\mathbf{w}} \sqrt{\frac{\log p}{n_{\mathcal{A}_h} + n_0}}$ and $\lambda_{\delta} = C_{\delta} \sqrt{\frac{\log p}{n_0}}$, where $C_{\mathbf{w}}$ and C_{δ} are sufficiently large positive constants. Then

$$\sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \lesssim \left(\frac{s \log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/2} + \left[\left(\frac{\log p}{n_0} \right)^{1/4} h^{1/2} \right] \wedge h \right) \geq 1 - n_0^{-1}, \quad (6)$$

$$\sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 \lesssim s \left(\frac{\log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/2} + h \right) \geq 1 - n_0^{-1}. \quad (7)$$

Remark 2. When $h \ll s\sqrt{\frac{\log p}{n_0}}$, $n_{\mathcal{A}_h} \gg n_0$, the upper bounds in (6) and (7) are better than the classical Lasso ℓ_2 -bound $\mathcal{O}_p\left(\sqrt{\frac{s \log p}{n_0}}\right)$ and ℓ_1 -bound $\mathcal{O}_p\left(s\sqrt{\frac{\log p}{n_0}}\right)$ using only target data.

Similar to [Theorem 2](#) in Li, Cai, and Li (2021), we can show the following lower bound of ℓ_1/ℓ_2 -estimation error in regime $\Xi(s, h)$ in the minimax sense.

Theorem 2 (ℓ_1/ℓ_2 -minimax estimation error bound). Assume [Assumptions 1, 2, and 4](#) hold. Also assume [Assumption 3.\(i\)](#) holds or [Assumption 3.\(ii\)](#) with $n_0 \gtrsim s^2 \log p$ holds. Then

$$\begin{aligned} \inf_{\hat{\boldsymbol{\beta}}} \sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \gtrsim \left(\frac{s \log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/2} + \left(\frac{s \log p}{n_0} \right)^{1/2} \wedge \left[\left(\frac{\log p}{n_0} \right)^{1/4} h^{1/2} \right] \wedge h \right) &\geq \frac{1}{2}, \\ \inf_{\hat{\boldsymbol{\beta}}} \sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 \gtrsim s \left(\frac{\log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/2} + \left[s \left(\frac{\log p}{n_0} \right)^{1/2} \right] \wedge h \right) &\geq \frac{1}{2}. \end{aligned}$$

Remark 3. [Theorem 2](#) indicates that under conditions on h required by [Theorem 1](#) ($h \lesssim s\sqrt{\log p/n_0}$), \mathcal{A}_h -Trans-GLM achieves the minimax optimal rate of ℓ_1/ℓ_2 -estimation error bound.

Next, we present a similar upper bound, which is weaker than the bound above but holds without requiring [Assumption 4](#).

Theorem 3 (ℓ_1/ℓ_2 -estimation error bound of \mathcal{A}_h -Trans-GLM without [Assumption 4](#)). Assume [Assumptions 1 and 2](#) hold. Suppose $h \ll \sqrt{\frac{n_0}{\log p}}$, $h \leq Cs^{-1/2}$, $n_0 \geq C \log p$ and $n_{\mathcal{A}_h} \geq Cs \log p$, where $C > 0$ is a constant. Also assume [Assumption 3.\(i\)](#) holds or [Assumption 3.\(ii\)](#) with $h \leq C'U^{-1}\bar{U}$ for some $C' > 0$ holds. We take $\lambda_{\mathbf{w}} = C_{\mathbf{w}}\left(\sqrt{\frac{\log p}{n_{\mathcal{A}_h} + n_0}} + h\right)$ and $\lambda_{\delta} = C_{\delta}\sqrt{\frac{\log p}{n_0}}$, where $C_{\mathbf{w}}$ and C_{δ} are sufficiently large positive constants. Then

$$\sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_2 \lesssim \left(\frac{s \log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/2} + \sqrt{sh} + \left[\left(\frac{\log p}{n_0} \right)^{1/4} h^{1/2} \right] \wedge h \right) \geq 1 - n_0^{-1},$$

$$\sup_{\xi \in \Xi(s,h)} \mathbb{P} \left(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_1 \lesssim s\sqrt{\frac{\log p}{n_{\mathcal{A}_h} + n_0}} + sh \right) \geq 1 - n_0^{-1}.$$

Remark 4. When $h \ll \sqrt{\frac{\log p}{n_0}}$ and $n_{\mathcal{A}_h} \gg n_0$, the upper bounds in (i) and (ii) are better than the classical Lasso bound $\mathcal{O}_p\left(\sqrt{\frac{\log p}{n_0}}\right)$ with target data.

Comparing the results in [Theorems 1 and 3](#), we know that with [Assumption 4](#), we could get sharper ℓ_1/ℓ_2 -estimation error bounds.

3.2. Theory on the Transferable Source Detection Algorithm

In this section, we will show that under certain conditions, our transferable set detection algorithm (Trans-GLM) can recover the level- h transferring set \mathcal{A}_h for some specific h , that is, $\hat{\mathcal{A}} = \mathcal{A}_h$ with high probability. Under these conditions, transferring with $\hat{\mathcal{A}}$ will lead to a faster convergence rate compared to Lasso fitted on the target data, when the target sample size n_0 falls into certain regime. But as we described in [Section 2.4, Algorithm 2](#) does not require any explicit input of h .

The corresponding population version of $\hat{L}_0^{[r]}(\mathbf{w})$ defined in (3) is

$$\begin{aligned} L_0(\mathbf{w}) &= -\mathbb{E}[\log \rho(y^{(0)})] - \mathbb{E}[y^{(0)} \mathbf{w}^T \mathbf{x}^{(0)}] + \mathbb{E}[\psi(\mathbf{w}^T \mathbf{x}^{(0)})] \\ &= -\mathbb{E}[\log \rho(y^{(0)})] - \mathbb{E}[\psi'(\mathbf{w}^T \mathbf{x}^{(0)}) \mathbf{w}^T \mathbf{x}^{(0)}] \\ &\quad + \mathbb{E}[\psi(\mathbf{w}^T \mathbf{x}^{(0)})]. \end{aligned}$$

Based on [Assumption 6](#), similar to (2), for $\{k\}$ -Trans-GLM ([Algorithm 1](#) with $\mathcal{A} = \{k\}$) used in [Algorithm 2](#), consider the following population version of estimators from the transferring step with respect to target data and the k th source data, which is the solution $\boldsymbol{\beta}^{(k)}$ of equation $\sum_{j \in \{0,k\}} \alpha_j^{(k)} \mathbb{E} \left\{ [\psi'((\boldsymbol{\beta}^{(k)})^T \mathbf{x}^{(k)}) - \psi'((\mathbf{w}^{(k)})^T \mathbf{x}^{(k)})] \mathbf{x}^{(k)} \right\} = 0$, where $\alpha_0^{(k)} = \frac{2n_0/3}{2n_0/3+n_k}$ and $\alpha_k^{(k)} = \frac{n_k}{2n_0/3+n_k}$. Define $\boldsymbol{\beta}^{(0)} = \boldsymbol{\beta}$. Next, let's impose a general assumption to ensure the identifiability of some \mathcal{A}_h by Trans-GLM.

Assumption 5 (Identifiability of \mathcal{A}_h). Denote $\mathcal{A}_h^c = \{1, \dots, K\} \setminus \mathcal{A}_h$. Suppose for some h and for all k , we have

$$\begin{aligned} \mathbb{P} \left(\sup_r |\hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(k)[r]}) - \hat{L}_0^{[r]}(\boldsymbol{\beta}^{(k)})| > \gamma_1^{(k)} + \zeta \Gamma_1^{(k)} \right) &\lesssim g_1^{(k)}(\zeta), \\ \mathbb{P} \left(\sup_r |\hat{L}_0^{[r]}(\boldsymbol{\beta}^{(k)}) - L_0(\boldsymbol{\beta}^{(k)})| > \zeta \Gamma_2^{(k)} \right) &\lesssim g_2^{(k)}(\zeta), \end{aligned}$$

where $g_1^{(k)}(\zeta), g_2^{(k)}(\zeta) \rightarrow 0$ as $\zeta \rightarrow \infty$. Assume $\inf_{k \in \mathcal{A}_h^c} \lambda_{\min}(\mathbb{E}[\int_0^1 \psi''((1-t)(\mathbf{x}^{(0)})^T \boldsymbol{\beta} + t(\mathbf{x}^{(0)})^T \boldsymbol{\beta}^{(k)}) dt \cdot \mathbf{x}^{(0)}(\mathbf{x}^{(0)})^T]) := \underline{\lambda} > 0$, and

$$\|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 \geq \underline{\lambda}^{-1/2} \left[C_1 \left(\sqrt{\Gamma_1^{(0)}} \vee \sqrt{\Gamma_2^{(0)}} \vee 1 \right) + \sqrt{2\Upsilon_1^{(k)}} \right], \quad \forall k \in \mathcal{A}_h^c \quad (8)$$

$$\begin{aligned} \Upsilon_1^{(k)} + \Gamma_1^{(k)} + \Gamma_2^{(k)} + h^2 &= \mathcal{O}(1), \forall k \in \mathcal{A}_h; \\ \Gamma_1^{(k)} = \mathcal{O}(1), \Gamma_2^{(k)} &= \mathcal{O}(1), \forall k \in \mathcal{A}_h^c, \end{aligned} \quad (9)$$

where $C_1 > 0$ is sufficiently large.

Remark 5. Here we use generic notations $\Upsilon_1^{(k)}, \Gamma_1^{(k)}, \Gamma_2^{(k)}, g_1^{(k)}(\zeta)$, and $g_2^{(k)}(\zeta)$. We show their explicit forms under linear, logistic, and Poisson regression models in Proposition 1 in Section S.1.2.1 of supplementary materials.

Remark 6. Condition (8) guarantees that for the sources not in \mathcal{A}_h , there is a sufficiently large gap between the population-level coefficient from the transferring step and the true coefficient of target data. Condition (9) guarantees the variations of $\sup_r |\hat{L}_0^{[r]}(\hat{\boldsymbol{\beta}}^{(k)[r]}) - \hat{L}_0^{[r]}(\boldsymbol{\beta}^{(k)})|$ and $\sup_r |\hat{L}_0^{[r]}(\boldsymbol{\beta}^{(k)}) - L_0(\boldsymbol{\beta}^{(k)})|$ are shrinking as the sample sizes go to infinity.

Based on Assumption 5, we have the following detection consistency property.

Theorem 4 (Detection consistency of \mathcal{A}_h). For Algorithm 2 (Trans-GLM), with Assumption 5 satisfied for some h , for any $\delta > 0$, there exist constants $C'(\delta)$ and $N = N(\delta) > 0$ such that when $C_0 = C'(\delta)$ and $\min_{k \in \{0\} \cup \mathcal{A}_h} n_k > N(\delta)$,

$$\mathbb{P}(\hat{\mathcal{A}} = \mathcal{A}_h) \geq 1 - \delta.$$

Then Algorithm 2 has the same high-probability upper bounds of ℓ_1/ℓ_2 -estimation error as those in Theorems 1 and 3 under the same conditions, respectively.

Remark 7. We would like to emphasize again that Algorithm 2 does not require the explicit input of h . Theorem 4 tells us that the transferring set $\hat{\mathcal{A}}$ suggested by Trans-GLM will be \mathcal{A}_h for some h , under certain conditions.

Next, we attempt to provide a sufficient and more explicit condition (Corollary 1) to ensure that Assumption 5 hold. Recalling the procedure of Algorithm 2, we note that it relies on using the negative log-likelihood as the similarity metric between target and source data, where the accurate estimation of coefficients or log-likelihood for GLM under the high-dimensional setting depends on the sparsity structure. Therefore, in order to provide an explicit and sufficient condition for Assumption 6 to hold, we now impose a “weak” sparsity assumption on both $\mathbf{w}^{(k)}$ and $\boldsymbol{\beta}^{(k)}$ with $k \in \mathcal{A}_h^c$ for some h . Note that the source data in \mathcal{A}_h automatically satisfy the sparsity condition due to the definition of \mathcal{A}_h .

Assumption 6. For some h and any $k \in \mathcal{A}_h^c$, we assume $\mathbf{w}^{(k)}$ and $\boldsymbol{\beta}^{(k)}$ can be decomposed as follows with some s' and $h' > 0$:

- (i) $\mathbf{w}^{(k)} = \boldsymbol{\zeta}^{(k)} + \boldsymbol{\vartheta}^{(k)}$, where $\|\boldsymbol{\zeta}^{(k)}\|_0 \leq s'$ and $\|\boldsymbol{\vartheta}^{(k)}\|_1 \leq h'$;
- (ii) $\boldsymbol{\beta}^{(k)} = \boldsymbol{\iota}^{(k)} + \boldsymbol{\omega}^{(k)}$, where $\|\boldsymbol{\iota}^{(k)}\|_0 \leq s'$ and $\|\boldsymbol{\omega}^{(k)}\|_1 \leq h'$.

Corollary 1. Assume Assumptions 1, 2, 6, and $\inf_{k \in \mathcal{A}_h^c} \lambda_{\min}(\mathbb{E}[\int_0^1 \psi''((1-t)\mathbf{x}^{(0)})^T \boldsymbol{\beta} + t\mathbf{x}^{(0)})^T \boldsymbol{\beta}^{(k)} dt \cdot \mathbf{x}^{(0)}(\mathbf{x}^{(0)})^T]) := \underline{\lambda} > 0$ hold. Also assume $\sup_{k \in \mathcal{A}_h^c} \|\boldsymbol{\beta}^{(k)}\|_\infty < \infty, \sup_k \|\mathbf{w}^{(k)}\|_\infty < \infty$. Let $\lambda^{(k)[r]} = C \left(\sqrt{\frac{\log p}{n_k + n_0}} + h \right)$ when $k \in \mathcal{A}_h$, $\lambda^{(k)[r]} = C \sqrt{\frac{\log p}{n_k + n_0}} \cdot (1 \vee \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 \vee \|\mathbf{w}^{(k)} - \boldsymbol{\beta}\|_2)$ when $k \in \mathcal{A}_h^c$ and $\lambda^{(0)[r]} = C \sqrt{\frac{\log p}{n_0}}$ for some sufficiently large constant $C > 0$. Then we have the following sufficient conditions to make Assumption 5 hold for logistic, linear and Poisson regression models. Denote $\Omega = \sqrt{h'} \left(\frac{\log p}{\min_{k \in \mathcal{A}_h} n_k + n_0} \right)^{1/4} + \left(\frac{s' \log p}{\min_{k \in \mathcal{A}_h} n_k + n_0} \right)^{1/4} [(s \vee s')^{1/4} + \sqrt{h'}] + \left(\frac{\log p}{\min_{k \in \mathcal{A}_h} n_k + n_0} \right)^{1/8} (h')^{1/4} [(s \vee s')^{1/8} + (h')^{1/4}]$.

(i) For logistic regression models, we require

$$\begin{aligned} \inf_{k \in \mathcal{A}_h} n_k &\gg s \log p, \quad n_0 \gg \{[s \vee s' + (h')^2] \vee \Omega^2\} \cdot \log K, \\ \inf_{k \in \mathcal{A}_h^c} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 &\gtrsim \left(\frac{s \log p}{n_0} \right)^{1/4} \vee 1 + \Omega, \quad h \ll s^{-1/2}. \end{aligned}$$

(ii) For linear models, we require

$$\begin{aligned} \inf_{k \in \mathcal{A}_h} n_k &\gg s^2 \log p, \\ n_0 &\gg \{[(s \vee s')^2 + (h')^4] \vee [(s \vee s' + (h')^2)\Omega^2]\} \cdot \log K, \\ \inf_{k \in \mathcal{A}_h^c} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 &\gtrsim \left(\frac{s^2 \log p}{n_0} \right)^{1/4} \vee 1 + [(s')^{1/4} + \sqrt{h'}] \Omega, \\ &h \ll s^{-1}. \end{aligned}$$

(iii) For Poisson regression models, we require

$$\begin{aligned} \inf_{k \in \mathcal{A}_h} n_k &\gg s^2 \log p, \\ n_0 &\gg [(s \vee s' + h') \vee \Omega^2] \cdot \log K, \quad U(s \vee s' + h \vee h') \lesssim 1, \\ \inf_{k \in \mathcal{A}_h^c} \|\boldsymbol{\beta}^{(k)} - \boldsymbol{\beta}\|_2 &\gtrsim \left(\frac{s \log p}{n_0} \right)^{1/4} \vee 1 + [(s')^{1/4} + \sqrt{h'}] \Omega, \\ &h \ll s^{-1}. \end{aligned}$$

Under Assumptions 1, 2, and the sufficient conditions derived in Corollary 1, by Theorem 4, we can conclude that $\hat{\mathcal{A}} = \mathcal{A}_h$ for some h . Note that we don't impose Assumption 4 here. Remark 4 indicates that, for \mathcal{A}_h -Trans-GLM to have a faster convergence rate than Lasso on target data, we need $h \ll \sqrt{\frac{\log p}{n_0}}$ and $n_{\mathcal{A}_h} \gg n_0$. Suppose $s' \asymp s$ and $h' \lesssim s^{1/2}$. Then for logistic regression models, when $n_0 \gg s \log(p \vee K)$, we can claim that as long as the source is “helpful”, i.e. there exists such an $h \ll \sqrt{\frac{\log p}{n_0}}$, Trans-GLM is able to detect \mathcal{A}_h . The same conclusion holds for linear models when $n_0 \gg s^2 \log(p \vee K)$ and for Poisson models when $n_0 \gg (s \log K) \vee (s^2 \log p)$. This implies that when target sample size n_0 is large enough and there exists an h such that \mathcal{A}_h is helpful, Trans-GLM can successfully

detect it with high probability and subsequently lead to a better ℓ_1/ℓ_2 -estimation error bound than that of the classical Lasso on target data.

3.3. Theory on Confidence Interval Construction

In this section, we will establish the theory for our confidence interval construction procedure described in Algorithm 3. First, we would like to review and introduce some notations. In Section 2.5, we defined $\Sigma_{\beta}^{(k)} = \mathbb{E}[\mathbf{x}^{(k)}(\mathbf{x}^{(k)})^T \psi''((\mathbf{x}^{(k)})^T \beta)]$. Let $\Theta = (\Sigma_{\beta}^{(0)})^{-1}$ and $K_{\mathcal{A}_h} = |\mathcal{A}_h|$. Define

$$\begin{aligned} \boldsymbol{\gamma}_j^{(k)} &= \arg \min_{\boldsymbol{\gamma} \in \mathbb{R}^{p-1}} \mathbb{E} \left\{ \psi''(\beta^T \mathbf{x}^{(k)}) \cdot [\mathbf{x}_j^{(k)} - (\mathbf{x}_{-j}^{(k)})^T \boldsymbol{\gamma}]^2 \right\} \\ &= (\Sigma_{\beta, -j, -j}^{(k)})^{-1} \Sigma_{\beta, -j, j}^{(k)}, \end{aligned}$$

which is closely related to $(\Sigma_{\beta}^{(k)})^{-1}$ and $\boldsymbol{\gamma}_j^{(0)}$ can be viewed as the population version of $\hat{\boldsymbol{\gamma}}_j^{(0)}$. And $\Sigma_{\beta, j, -j}^{(k)}$ represents the j th row without the (j, j) diagonal element of $\Sigma_{\beta}^{(k)}$. $\Sigma_{\beta, -j, -j}^{(k)}$ denotes the submatrix of $\Sigma_{\beta}^{(k)}$ without the j th row and j th column. Suppose

$$\sup_{k \in \mathcal{A}_h, j=1:p} \|(\Sigma_{\beta, -j, -j}^{(0)})^{-1} \Sigma_{\beta, -j, j}^{(0)} - (\Sigma_{\beta, -j, -j}^{(k)})^{-1} \Sigma_{\beta, -j, j}^{(k)}\|_1 \leq h_1,$$

$$\sup_{k \in \mathcal{A}_h, j=1:p} \left[\left| \Sigma_{\beta, j, j}^{(k)} - \Sigma_{\beta, j, j}^{(0)} \right| \vee \left| (\Sigma_{\beta, j, -j}^{(k)} - \Sigma_{\beta, j, -j}^{(0)}) \boldsymbol{\gamma}_j^{(0)} \right| \right] \leq h_{\max}.$$

Then by the definition of $\boldsymbol{\gamma}_j^{(k)}$,

$$\sup_{k \in \mathcal{A}_h, j=1:p} \|\boldsymbol{\gamma}_j^{(k)} - \boldsymbol{\gamma}_j^{(0)}\|_1 \lesssim h_1,$$

which is similar to our previous setting $\sup_{k \in \mathcal{A}_h} \|\mathbf{w}^{(k)} - \beta\|_1 \leq h$. This motivates us to apply a similar two-step transfer learning procedure (Steps 2–4 in Algorithm 3) to learn $\boldsymbol{\gamma}_j^{(0)}$ for $j = 1, \dots, p$. We impose the following set of conditions.

Assumption 7. Suppose the following conditions hold:

- (i) $\sup_{k \in \{0\} \cup \mathcal{A}_h} \|\mathbf{x}^{(k)}\|_{\infty} \leq U < \infty, \sup_{k \in \{0\} \cup \mathcal{A}_h} |(\mathbf{x}^{(k)})^T \mathbf{w}^{(k)}| \leq U' < \infty$ a.s.;
- (ii) $\sup_j \|\boldsymbol{\gamma}_j^{(0)}\|_0 / s < \infty, \sup_{j \in 1:p, k \in \{0\} \cup \mathcal{A}_h} |(\mathbf{x}^{(k)})^T \boldsymbol{\gamma}_j^{(0)}| \leq U'' < \infty$ a.s.;
- (iii) $\inf_{k \in \{0\} \cup \mathcal{A}_h} \lambda_{\min}(\Sigma_{\mathbf{w}^{(k)}}) \geq \underline{U} > 0$;
- (iv) $\sup_{k \in \{0\} \cup \mathcal{A}_h} \sup_{|z| \leq \bar{U}} \psi'''((\mathbf{x}^{(k)})^T \mathbf{w}^{(k)} + z) \leq M_{\psi} < \infty$ a.s.;
- (v) $\sup_{k \in \{0\} \cup \mathcal{A}_h} \|(\Sigma_{\beta, -j, -j}^{\mathcal{A}_h})^{-1} \Sigma_{\beta, -j, j}^{(k)}\|_1 < \infty$, where $\Sigma_{\beta}^{\mathcal{A}_h} = \sum_{k \in \{0\} \cup \mathcal{A}_h} \alpha_k \Sigma_{\beta}^{(k)}$;
- (vi) $\min_{k \in \mathcal{A}_h} n_k \gtrsim n_0, n_0 \gg \frac{s^3 (\log p)^2}{K_{\mathcal{A}_h}^2} \vee K_{\mathcal{A}_h}, n_{\mathcal{A}_h} + n_0 \gg s^2 \log p$;
- (vii) $h_1 \lesssim s^{-1/2} \wedge \left[\sqrt{\frac{n_0}{\log p}} \left(\frac{\sqrt{K_{\mathcal{A}_h}}}{s} \wedge 1 \right) \right], h_1 \vee h \ll \frac{K_{\mathcal{A}_h} n_0^{1/2}}{s^2 (\log p)^{3/2}} \wedge \frac{n_0^{1/4}}{s^{1/2} (\log p)^{1/4}}, hh_1^{1/2} \ll n_0^{-1/4} (\log p)^{-1/4} \left(\frac{K_{\mathcal{A}_h}}{s} \wedge 1 \right), h^{5/2} h_1 \ll n_0^{-3/4} (s \log p)^{-1/4}, h_1 \ll \frac{K_{\mathcal{A}_h}^{1/2} n_0^{1/2}}{s^{3/2} (\log p)^{1/2}} \wedge$

$$\begin{aligned} \frac{K_{\mathcal{A}_h}^{3/2} n_0^{1/2}}{s^{5/2} (\log p)^{3/2}}, h_1 h^{1/2} &\ll \frac{n_0^{1/4}}{s (\log p)^{1/4}} \wedge \frac{K_{\mathcal{A}_h} n_0^{1/4}}{s^2 (\log p)^{5/4}}, h \ll \frac{K_{\mathcal{A}_h}^{1/2}}{(s \log p)^{1/2}} \wedge \frac{1}{n_0^{1/4} (\log p)^{1/2}}, h_{\max} \ll s^{-1/2} \wedge \left(\frac{1}{s} \sqrt{\frac{K_{\mathcal{A}_h}}{\log p}} \right), \\ hh_{\max} &\ll n_0^{-1/2}. \end{aligned}$$

Remark 8. Conditions (i)–(iii) are motivated from conditions of Theorem 3.3 in Van de Geer et al. (2014). Note that in Van de Geer et al. (2014), they define $s_j = \|\boldsymbol{\gamma}_j^{(0)}\|_0$ and treat s_j and s as two different parameters. Here we require $\sup_j s_j \lesssim s$ just for simplicity (otherwise condition (vii) would be more complicated). Condition (iv) requires the inverse link function to behave well, which is similar to Assumption 3. Condition (v) is similar to Assumption 4 to guarantee the success of the two-step transfer learning procedure to learn $\boldsymbol{\gamma}^{(0)}$ in Algorithm 3 with a fast rate. Without condition (v), the conclusions in the following Theorem 5 may still hold but under a stronger condition on h, h_1 and h_{\max} , and the rate (11) may be worse. We do not explore the details in this article and leave them to interested readers. Conditions (vi) and (vii) require that the sample size is sufficiently large and the distance between target and source is not too large. In condition (vi), $\min_{k \in \mathcal{A}_h} n_k \gtrsim n_0$ is not necessary and the only reason we add it here is to simplify condition (vii).

Remark 9. When $\mathbf{x}^{(k)}$'s are from the same distribution, $h_1 = h_{\max} = 0$. In this case, we can drop the debiasing step to estimate $\hat{\boldsymbol{\gamma}}_j^{(0)}$ in Algorithm 3 as well as condition (v). Furthermore, condition (vii) can be significantly simplified and only $h \ll \frac{K_{\mathcal{A}_h} n_0^{1/2}}{s^2 (\log p)^{3/2}} \wedge \frac{n_0^{1/4}}{s^{1/2} (\log p)^{1/4}} \wedge \frac{K_{\mathcal{A}_h}^{1/2}}{(s \log p)^{1/2}} \wedge \frac{1}{n_0^{1/4} (\log p)^{1/2}}$ is needed.

Remark 10. From conditions (vi) and (vii), we can see that as long as $K_{\mathcal{A}_h} \lesssim s (\log p)^{2/3}$, the conditions become milder as $K_{\mathcal{A}_h}$ increases.

Now, we are ready to present our main result for Algorithm 3.

Theorem 5. Under Assumptions 1–4 and Assumption 7,

$$\frac{\sqrt{n_0}(\hat{b}_j - \beta_j)}{\sqrt{\widehat{\Theta}_j^T \widehat{\Sigma}_{\beta} \widehat{\Theta}_j}} \xrightarrow{d} \mathcal{N}(0, 1), \tag{10}$$

and

$$\begin{aligned} &|\widehat{\Theta}_j^T \widehat{\Sigma}_{\beta} \widehat{\Theta}_j - \Theta_{jj}| \\ &\lesssim s \sqrt{\frac{\log p}{n_{\mathcal{A}_h} + n_0}} + \sqrt{s} \left[h^{1/2} \left(\frac{\log p}{n_0} \right)^{1/4} \wedge h \right] \\ &\quad + (sh_1)^{1/2} \left(\frac{\log p}{n_0} \right)^{1/4} \\ &\quad + (sh_1)^{1/2} \left[\left(\frac{s \log p}{n_{\mathcal{A}_h} + n_0} \right)^{1/4} + \left(h^{1/4} \left(\frac{\log p}{n_0} \right)^{1/8} \right) \right] \wedge h^{1/2} \\ &\quad + \sqrt{sh_{\max}}, \end{aligned} \tag{11}$$

for $j = 1, \dots, p$, with probability at least $1 - K_{\mathcal{A}_h} n_0^{-1}$.

Theorem 5 guarantees that under certain conditions, the $(1 - \alpha)$ -confidence interval for each coefficient component obtained in **Algorithm 3** has approximately level $(1 - \alpha)$ when the sample size is large. Also, if we compare the rate of (11) with the rate $\mathcal{O}_p(s\sqrt{\log p/n_0})$ in Van de Geer et al. (2014) (see the proof of Theorem 3.1), we can see that when $h \ll s\sqrt{\frac{\log p}{n_0}}$, $h_1 \ll \sqrt{\frac{s \log p}{n_0}}$. $\left[s^{1/2} \wedge \left(\frac{n_{\mathcal{A}_h} + n_0}{n_0} \right)^{1/4} \right]$, $h_1^{1/2} h^{1/4} \ll s^{1/2} \left(\frac{\log p}{n_0} \right)^{3/8}$ and $h_{\max} \ll \sqrt{\frac{s \log p}{n_0}}$, the rate is better than that of desparsified Lasso using only target data.

4. Numerical Experiments

In this section, we demonstrate the power of our GLM transfer learning algorithms via extensive simulation studies and a real-data application. In the simulation part, we study the performance of different methods under various settings of h . The methods include *Trans-GLM* (**Algorithm 2**), *naïve-Lasso* (Lasso on target data), *\mathcal{A}_h -Trans-GLM* (**Algorithm 1** with $\mathcal{A} = \mathcal{A}_h$) and *Pooled-Trans-GLM* (**Algorithm 1** with all sources). In the real-data study, besides naïve-Lasso, Pooled-Trans-GLM, and Trans-GLM, additional methods are explored for comparison, including support vector machines (SVM), decision trees (Tree), random forests (RF) and Adaboost algorithm with trees (Boosting). We run these benchmark methods twice. First, we fit the models on only the target data, then at the second time, we fit them a combined data of target and all sources, which is called a pooled version. We use the original method name to denote the corresponding method implemented on target data, and add a prefix “Pooled” to denote the corresponding method implemented on target and all source data. For example, Pooled-SVM represents SVM fitted on all data from target and sources.

All experiments are conducted in R. We implement our GLM transfer learning algorithms in a new R package `glmtrans`, which is available on CRAN. More implementation details can be found in Section S.1.3.1 in the supplementary materials.

4.1. Simulations

4.1.1. Transfer Learning on \mathcal{A}_h

In this section, we study the performance of \mathcal{A}_h -Trans-GLM and compare it with that of naïve-Lasso. The purpose of the simulation is to verify that \mathcal{A}_h -Trans-GLM can outperform naïve-Lasso in terms of the target coefficient estimation error, when h is not too large.

Consider the simulation setting as follows. We set the target sample size $n_0 = 200$ and source sample size $n_k = 100$ for each $k \neq 0$. The dimension $p = 500$ for both target and source data. For the target, the coefficient is set to be $\beta = (0.5 \cdot \mathbf{1}_s, \mathbf{0}_{p-s})^T$, where $\mathbf{1}_s$ has all s elements 1 and $\mathbf{0}_{p-s}$ has all $(p - s)$ elements 0, where s is set to be 5. Denote $\mathcal{R}_p^{(k)}$ as p independent Rademacher variables (being -1 or 1 with equal probability) for any k . $\mathcal{R}_p^{(k)}$ is independent with $\mathcal{R}_p^{(k')}$ for any $k \neq k'$. For any source data k in \mathcal{A}_h , we set $\mathbf{w}^{(k)} = \beta + (h/p)\mathcal{R}_p^{(k)}$. For linear and logistic regression models, predictors from target $\mathbf{x}_i^{(0)} \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}_p, \Sigma)$ with $\Sigma = [\Sigma_{ij}]_{p \times p}$ where $\Sigma_{ij} = 0.5^{|j-i|}$, for all $i = 1, \dots, p$. And for $k \in \mathcal{A}_h$, we generate p -dimensional

predictors from $\mathcal{N}(\mathbf{0}_p, \Sigma + \epsilon \epsilon^T)$, where $\epsilon \sim \mathcal{N}(\mathbf{0}_p, 0.3^2 \mathbf{I}_p)$ and is independently generated. For Poisson regression model, predictors are from the same Gaussian distributions as those in linear and binomial cases with coordinate-wise truncation at ± 0.5 .

Note that naïve-Lasso is fitted on only target data, and \mathcal{A}_h -Trans-GLM denotes **Algorithm 1** on source data in \mathcal{A}_h as well as target data. We train naïve-Lasso and \mathcal{A}_h -Trans-GLM models under different settings of h and $K_{\mathcal{A}_h}$, then calculate the ℓ_2 -estimation error of β . All the experiments are replicated 200 times and the average ℓ_2 -estimation errors of \mathcal{A}_h -Trans-GLM and naïve-Lasso under linear, logistic, and Poisson regression models are shown in **Figure 1**.

From **Figure 1**, it can be seen that \mathcal{A}_h -Trans-GLM outperforms naïve-Lasso for most combinations of h and K . As more and more source data become available, the performance of \mathcal{A}_h -Trans-GLM improves significantly. This matches our theoretical analysis because the ℓ_2 -estimation error bounds in **Theorems 1** and **3** become sharper as $n_{\mathcal{A}_h}$ grows. When h increases, the performance of \mathcal{A}_h -Trans-GLM becomes worse.

We also apply the inference **Algorithm 3** with \mathcal{A}_h and compare it with desparsified Lasso (Van de Geer et al. 2014) on only target data. Recall the notations we used in **Section 3.3**. Here we consider 95% confidence intervals (CIs) for each component of coefficient β , and report three evaluation metrics in **Figure 2** when $h = 20$ under different $K_{\mathcal{A}_h}$: (i) the average of estimation error of Θ_{jj} over variables in the signal set S and noise set S^c (including the intercept), respectively (which we call “average estimation error”); (ii) the average CI coverage probability over variables in the signal set S and noise set S^c ; (iii) the average CI length over $j \in$ signal set S and noise set S^c . Note that there is no explicit formula of Θ for logistic and Poisson regression models. Here we approximated it through 5×10^6 Monte Carlo simulations. Notice that the average estimation error of \mathcal{A}_h -Trans-GLM declines as K increases, which agrees with our theoretical analysis in **Section 3.3**. As for the coverage probability, although CIs obtained by desparsified Lasso can achieve 95% coverage probability on S^c in linear and binomial cases, it fails to meet the 95% requirement of coverage probability on S in all three cases. In contrast, CIs provided by \mathcal{A}_h -Trans-GLM can achieve approximately 95% level when K is large on both S and S^c . Finally, the results of average CI length reveal that the CIs obtained by \mathcal{A}_h -Trans-GLM tend to be wider as K increases. Considering this together with the average estimation error and coverage probability, a possible explanation could be that desparsified Lasso might down-estimate Θ_{jj} which leads to too narrow CIs to cover the true coefficients. And \mathcal{A}_h -Trans-GLM offers a more accurate estimate of Θ_{jj} which results in wider CIs.

We also consider different $(\{n_k\}_{k=0}^K, p, s)$ settings with the results in the supplementary materials.

4.1.2. Transfer Learning When \mathcal{A}_h is Unknown

Different from the previous section, now we fix the total number of sources as $K = 10$. There are two types of sources, which belong to either \mathcal{A}_h or \mathcal{A}_h^c . Sources from \mathcal{A}_h have similar coefficients to the target one, while the coefficients of sources outside \mathcal{A}_h can be quite different. Intuitively, using more sources from \mathcal{A}_h benefits the estimation of the target coefficient. But in

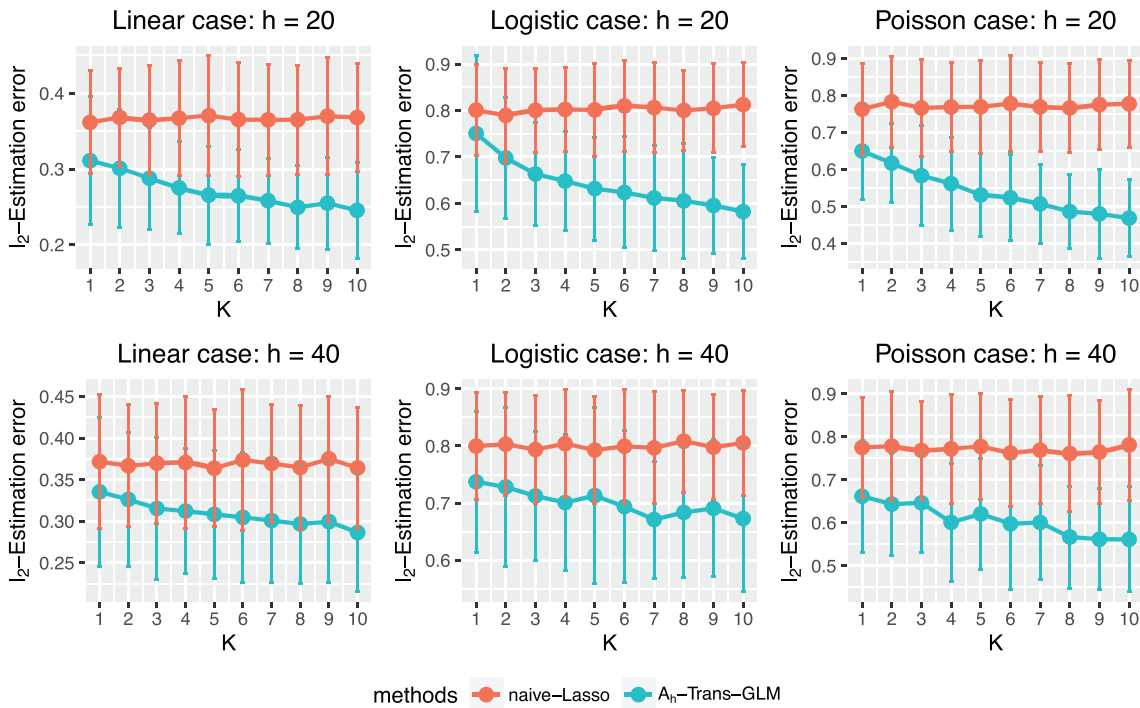


Figure 1. The average l_2 -estimation error of \mathcal{A}_h -Trans-GLM and naive-Lasso under linear, logistic and Poisson regression models with different settings of h and K . $n_0 = 200$ and $n_k = 100$ for all $k = 1, \dots, p, p = 500, s = 5$. Error bars denote the standard deviations.

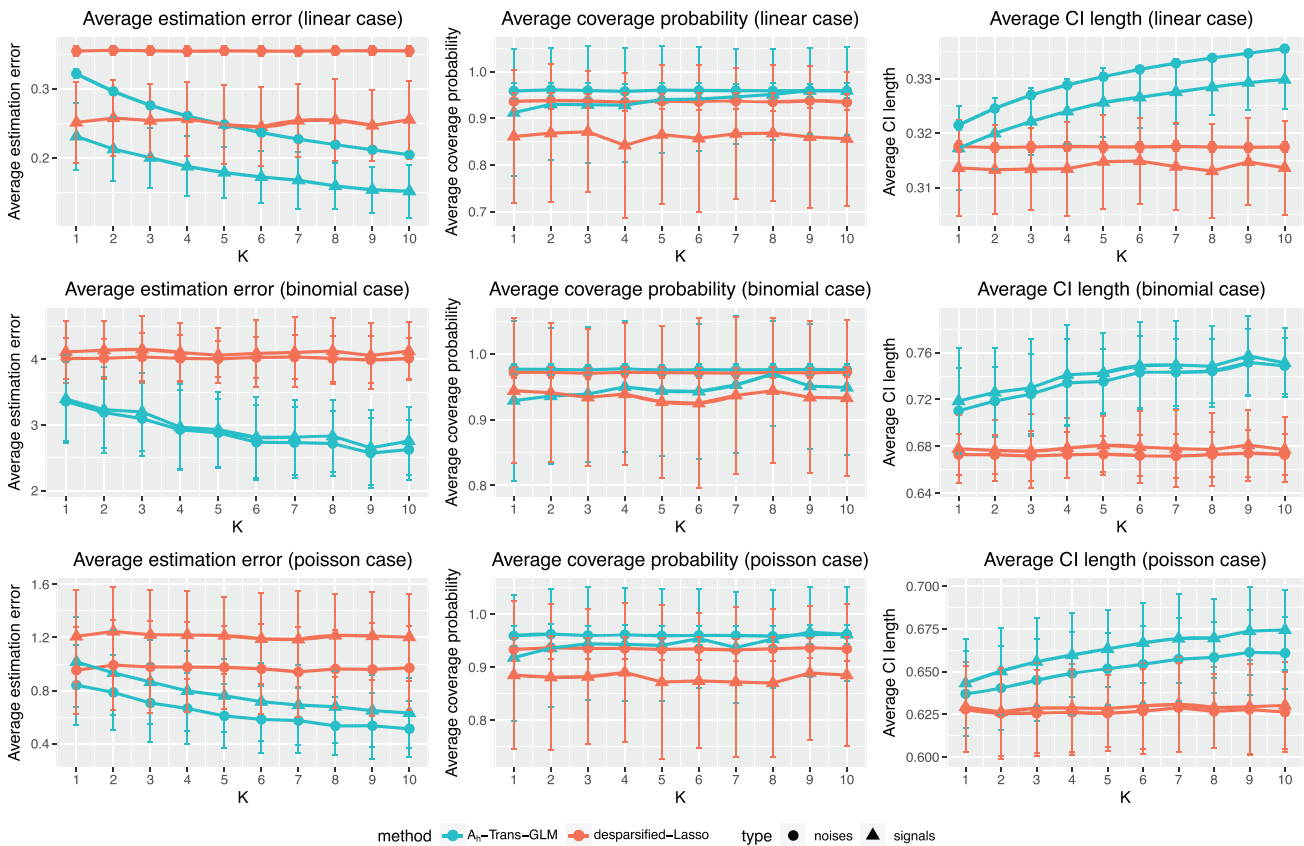


Figure 2. Three evaluation metrics of Algorithm 3 with \mathcal{A}_h (we denote it as \mathcal{A}_h -Trans-GLM) and desparsified Lasso on target data, under linear, logistic and Poisson regression models, with different settings of K . $h = 20$. $n_0 = 200$ and $n_k = 100$ for all $k = 1, \dots, p, p = 500, s = 5$. Error bars denote the standard deviations.

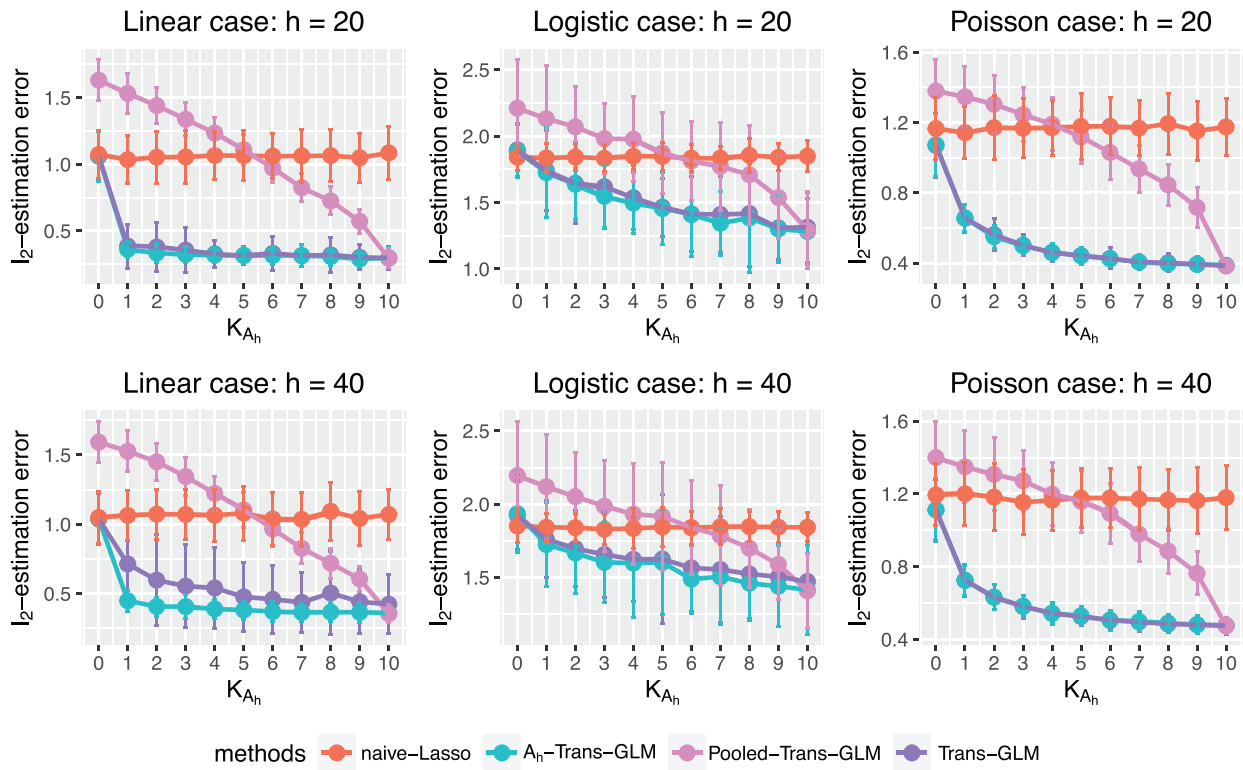


Figure 3. The average ℓ_2 -estimation error of various models with different settings of h and $K_{\mathcal{A}_h}$ when $K = 10$. $n_k = 200$ for all $k = 0, \dots, K$, $p = 2000$, $s = 20$. Error bars denote the standard deviations.

practice, \mathcal{A}_h may not be known a priori. As we argued before, Trans-GLM can detect useful sources automatically, therefore, it is expected to be helpful in such a scenario. Simulations in this section aim to justify the effectiveness of Trans-GLM.

Here is the detailed setting. We set the target sample size $n_0 = 200$ and source sample size $n_k = 200$ for all $k \neq 0$. The dimension $p = 2000$. Target coefficient is the same as the one used in Section 4.1.1 and we fix the signal number $s = 20$. Recall $\mathcal{R}_p^{(k)}$ denotes p independent Rademacher variables and $\mathcal{R}_p^{(k')}$ are independent for any $k \neq k'$. Consider $h = 20$ and 40 . For any source data k in \mathcal{A}_h , we set $\mathbf{w}^{(k)} = \boldsymbol{\beta} + (h/p)\mathcal{R}_p^{(k)}$. For linear and logistic regression models, predictors from target $\mathbf{x}_i^{(0)} \stackrel{\text{iid}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = [\Sigma_{ij}]_{p \times p}$ where $\Sigma_{ij} = 0.9^{|i-j|}$, for all $i = 1, \dots, n_0$. For the source, we generate p -dimensional predictors from independent t -distribution with degrees of freedom 4. For the target and sources of Poisson regression model, we generate predictors from the same Gaussian distribution and t -distribution, respectively, and truncate each predictor at ± 0.5 .

To generate the coefficient $\mathbf{w}^{(k)}$ for $k \notin \mathcal{A}_h$, we randomly generate $S^{(k)}$ of size s from $\{2s + 1, \dots, p\}$. Then, the j th component of coefficient $\mathbf{w}^{(k)}$ is set to be

$$w_j^{(k)} = \begin{cases} 0.5 + 2hr_j^{(k)}/p, & j \in \{s + 1, \dots, 2s\} \cup S^{(k)}, \\ 2hr_j^{(k)}/p, & \text{otherwise,} \end{cases}$$

where $r_j^{(k)}$ is a Rademacher variable. We also add an intercept 0.5. The generating process of each source data is independent. Compared to the setting in Section 4.1.1, the current setting is more challenging because source predictors come from t -distribution with heavier tails than sub-Gaussian tails. However,

although Assumption 2 is violated, in the following analysis, we will see that Trans-GLM can still succeed in detecting informative sources.

As before, we fit naive-Lasso on only target data. \mathcal{A}_h -Trans-GLM and Pooled-Trans-GLM represent Algorithm 1 on source data in \mathcal{A}_h and target data or all sources and target data, respectively. Trans-GLM runs Algorithm 2 by first identifying the informative source set $\hat{\mathcal{A}}$, then applying Algorithm 1 to fit the model on sources in $\hat{\mathcal{A}}$. We vary the values of $K_{\mathcal{A}_h}$ and h , and repeat simulations in each setting 200 times. The average ℓ_2 -estimation errors are summarized in Figure 3.

From Figure 3, it can be observed that in all three models, \mathcal{A}_h -Trans-GLM always achieves the best performance as expected since it transfers information from sources in \mathcal{A}_h . Trans-GLM mimics the behavior of \mathcal{A}_h -Trans-GLM very well, implying that the transferable source detection algorithm can successfully recover \mathcal{A}_h . When $K_{\mathcal{A}_h}$ is small, Pooled-Trans-GLM performs worse than naive-Lasso because of the negative transfer. As $K_{\mathcal{A}_h}$ increases, the performance of Pooled-Trans-GLM improves and finally matches those of \mathcal{A}_h -Trans-GLM and Trans-GLM when $K_{\mathcal{A}_h} = K = 10$.

4.2. A Real-Data Study

In this section, we study the 2020 U.S. presidential election results of each county. We only consider win or lose between two main parties, Democrats and Republicans, in each county. The 2020 county-level election result is available at https://github.com/tonmcg/US_County_Level_Election_Results_08-20. The response is the election result of each county. If Democrats win, we denote this county as class 1, otherwise, we denote it

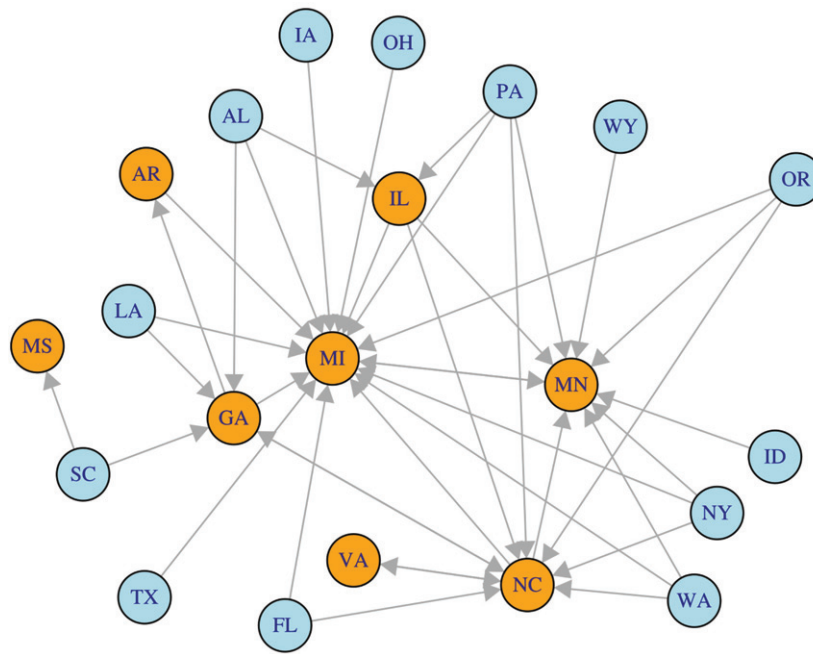


Figure 4. The transferability between different states for Trans-GLM.

as class 0. And we also collect the county-level information as the predictors, including the population and race proportions, from <https://www.kaggle.com/benhamner/2016-us-election>.

The goal is to explore the relationship between states in the election using transfer learning. We are interested in swing states with a large number of counties. Among 49 states (Alaska and Washington, D.C. excluded), we select the states where the proportion of counties voting Democrats falls in $[10\%, 90\%]$, and have at least 75 counties as target states. They include Arkansas (AR), Georgia (GA), Illinois (IL), Michigan (MI), Minnesota (MN), Mississippi (MS), North Carolina (NC), and Virginia (VA).

The original data includes 3111 counties and 52 county-level predictors. We also consider the pairwise interaction terms between predictors. After preprocessing, there are 3111 counties and 1081 predictors in the final data, belonging to 49 U.S. states.

We would like to investigate which states have a closer relationship with these target states by our transferable source detection algorithm. For each target state, we use it as the target data and the remaining 48 states as source datasets. Each time we randomly sample 80% of target data as training data and the remaining 20% is used for testing. Then we run Trans-GLM (Algorithm 2) and see which states are in the estimated transferring set $\hat{\mathcal{A}}$. We repeat the simulation 500 times and count the transferring frequency between every state pair. The 25 (directed) state pairs with the highest transferring frequencies are visualized in Figure 4. Each orange node represents a target state we mentioned above and blue nodes are source states. States with the top 25 transferring frequencies are connected with a directed edge.

From Figure 4, we observe that Michigan has a strong relationship with other states, since there is a lot of information transferable when predicting the county-level results in Michigan, Minnesota, and North Carolina. Another interesting finding is that states which are geographically close to each other

may share more similarities. For instance, see the connection between Indiana and Michigan, Ohio and Michigan, North Carolina and Virginia, South Carolina and Georgia, Alabama and Georgia, etc. In addition, one can observe that states in the Rust Belt also share more similarities. As examples, see the edges among Pennsylvania, Minnesota, Illinois, Michigan, New York, and Ohio, etc.

To further verify the effectiveness of our GLM transfer learning framework on this dataset and make our findings more convincing, we calculate the average test misclassification error rates for each of the eight target states. For comparison, we compare the performances of Trans-GLM and Pooled-Trans-GLM with naïve-Lasso, SVM, trees, random forests (RF), boosting trees (Boosting) as well as their pooled version. Average test errors and the standard deviations of various methods are summarized in Table 1. The best method and other top three methods for each target are highlighted in bold and italics, respectively.

Table 1 shows that in four out of eight scenarios, Trans-GLM performs the best among all approaches. Moreover, Trans-GLM is always ranked in the top three except in the case of target state MS. This verifies the effectiveness of our GLM transfer learning algorithm. Besides, Pooled-Trans-GLM can always improve the performance of naïve-Lasso, while for other methods, pooling the data can sometimes lead to worse performance than that of the model fitted on only the target data. This marks the success of our two-step transfer learning framework and the importance of the debiasing step. Combining the results with Figure 4, it can be seen that the performance improvement of Trans-GLM (compared to naïve-Lasso) for the target states with more connections (share more similarities with other states) are larger. For example, Trans-GLM outperforms naïve-Lasso a lot on Michigan, Minnesota and North Carolina, while it performs similarly to naïve-Lasso on Mississippi.

We also try to identify significant variables by Algorithm 3. Due to the space limit, we put the results and analysis in Section

Table 1. The average test error rate (in percentage) of various methods with different targets among 500 replications.

Methods	Target states							
	AR	GA	IL	MI	MN	MS	NC	VA
naïve-Lasso	4.79 _{3,36}	6.98 _{3,90}	5.73 _{4,14}	11.49 _{2,44}	12.46 _{2,70}	7.53 _{6,57}	15.60 _{6,73}	9.48 _{4,88}
Pooled-Lasso	3.59 _{4,71}	9.98 _{4,22}	7.89 _{5,56}	7.04 _{5,80}	10.38 _{5,18}	22.01 _{7,18}	12.73 _{5,35}	21.44 _{5,46}
Pooled-Trans-GLM	1.83 _{3,12}	4.86 _{3,60}	2.52 _{3,55}	5.62 _{4,54}	10.75 _{5,60}	7.23 _{6,65}	9.71 _{5,75}	7.15 _{4,23}
Trans-GLM	1.54 _{2,94}	4.74 _{3,54}	2.51 _{3,45}	5.53 _{4,73}	10.34 _{5,73}	7.24 _{6,81}	9.34 _{5,57}	7.18 _{4,67}
SVM	6.71 _{1,70}	17.09 _{3,89}	7.00 _{5,40}	12.59 _{1,87}	13.29 _{2,29}	23.92 _{8,90}	12.66 _{6,86}	10.78 _{5,29}
Pooled-SVM	7.84 _{6,32}	13.47 _{4,73}	7.75 _{5,24}	7.58 _{6,40}	13.01 _{5,69}	27.32 _{8,72}	12.30 _{5,75}	17.31 _{5,46}
Tree	2.23 _{3,58}	8.37 _{4,40}	4.62 _{5,27}	10.05 _{5,53}	10.97 _{8,42}	5.97 _{5,26}	18.29 _{8,01}	14.46 _{6,88}
Pooled-tree	7.81 _{6,89}	7.68 _{4,59}	4.63 _{4,26}	7.42 _{6,18}	10.53 _{5,91}	16.73 _{7,30}	14.76 _{7,26}	17.43 _{5,85}
RF	3.60 _{3,57}	6.04 _{3,59}	4.08 _{3,98}	6.42 _{4,79}	10.51 _{5,10}	7.27 _{5,72}	11.29 _{6,29}	7.73 _{4,77}
Pooled-RF	3.73 _{4,82}	7.49 _{3,90}	4.35 _{3,63}	5.34 _{4,99}	10.86 _{4,96}	12.56 _{6,88}	11.04 _{6,03}	10.40 _{5,18}
Boosting	2.23 _{3,58}	4.65 _{3,77}	2.55 _{3,82}	7.79 _{5,52}	10.64 _{6,51}	5.28 _{5,16}	10.88 _{6,47}	7.53 _{5,10}
Pooled-boosting	3.10 _{4,84}	5.71 _{3,53}	3.82 _{3,85}	5.81 _{5,27}	11.21 _{5,13}	14.31 _{7,42}	10.82 _{5,99}	11.95 _{5,25}

NOTE: The cutoff for all binary classification methods is set to be 1/2. Numbers in the subscript indicate the standard deviations.

S.1.3.4 of supplementary materials. Interested readers can find the details there. Furthermore, since we have considered all main effects and their interactions, one reviewer pointed out that besides the classical Lasso penalty, there are other variants like group Lasso (Yuan and Lin 2006) or Lasso with hierarchy restriction (Bien, Taylor, and Tibshirani 2013), which may bring better practical performance and model interpretation. To be consistent with our theories, we only consider the Lasso penalty here and leave other options for future study.

5. Discussions

In this work, we study the GLM transfer learning problem. To the best of our knowledge, this is the first article to study high-dimensional GLM under a transfer learning framework, which can be seen as an extension to Bastani (2021) and Li, Cai, and Li (2021). We propose GLM transfer learning algorithms, and derive bounds for ℓ_1/ℓ_2 -estimation error and a prediction error measure with fast and slow rates under different conditions. In addition, to avoid the negative transfer, an algorithm-free transferable source detection algorithm is developed and its theoretical properties are presented in detail. Moreover, we accommodate the two-step transfer learning method to construct confidence intervals of each coefficient component with theoretical guarantees. Finally, we demonstrate the effectiveness of our algorithms via simulations and a real-data study.

There are several promising future avenues that are worth further research. The first interesting problem is how to extend the current framework and theoretical analysis to other models, such as multinomial regression and the Cox model. Second, Algorithm 1 is shown to achieve the minimax ℓ_1/ℓ_2 estimation error rate when the homogeneity assumption (Assumption 4) holds. Without homogeneity of predictors between target and source, only suboptimal rates are obtained. This problem exists in the line of most existing high-dimensional transfer learning research (Li, Cai, and Li 2020, 2021; Bastani 2021). It remains unclear how to achieve the minimax rate when source predictors' distribution deviates a lot from the target one. Another promising direction is to explore similar frameworks for other machine learning models, including support vector machines, decision trees, and random forests.

Supplementary Materials

We summarize the additional simulation and real data analysis results, more theoretical details as well as all the proofs in the supplements.

Acknowledgments

We thank the editor, the AE, and anonymous reviewers for their insightful comments which have greatly improved the scope and quality of the article.

Funding

This work was supported by NIH grant 1R21AG074205-01, NYU University Research Challenge Fund, and through the NYU IT High Performance Computing resources, services, and staff expertise.

ORCID

Yang Feng  <http://orcid.org/0000-0001-7746-7598>

References

- Bastani, H. (2021), "Predicting with Proxies: Transfer Learning in High Dimension," *Management Science*, 67, 2964–2984. [2684,2685,2686,2688,2696]
- Bien, J., Taylor, J., and Tibshirani, R. (2013), "A Lasso for Hierarchical Interactions," *The Annals of Statistics*, 41, 1111–1141. [2696]
- Cai, T. T., and Wei, H. (2021), "Transfer Learning for Nonparametric Classification: Minimax Rate and Adaptive Classifier," *The Annals of Statistics*, 49, 100–128. [2684]
- Chen, A., Owen, A. B., Shi, M. (2015), "Data Enriched Linear Regression," *Electronic Journal of Statistics*, 9, 1078–1112. [2685]
- Gross, S. M., and Tibshirani, R. (2016), "Data Shared Lasso: A Novel Tool to Discover Uplift," *Computational Statistics & Data Analysis*, 101, 226–235. [2685]
- Hajiramezanali, E., and Zamani, S. (2018), "Bayesian Multi-Domain Learning for Cancer Subtype Discovery from Next-Generation Sequencing Count Data," in *Advances in Neural Information Processing Systems 31 (NIPS 2018)*, 31. [2684]
- Hanneke, S., and Kpotufe, S. (2020a), "A no-free-lunch Theorem for Multi-task Learning," arXiv preprint arXiv:2006.15785. [2685]
- (2020b), "On the Value of Target Data in Transfer Learning," arXiv preprint arXiv:2002.04747. [2685]
- Li, S., Cai, T. T., and Li, H. (2020), "Transfer Learning in Large-Scale Gaussian Graphical Models with False Discovery Rate Control," arXiv preprint arXiv:2010.11037. [2684,2696]
- (2021), "Transfer Learning for High-Dimensional Linear Regression: Prediction, Estimation and Minimax Optimality," *Journal of the Royal Statistical Society, Series B (to appear)*. [2684,2685,2686,2688,2689,2696]

- McCullagh, P., and Nelder, J. A. (1989), *Generalized Linear Models* (Vol. 37), Boca Raton, FL: CRC Press. [2685]
- Meinshausen, N., and Bühlmann, P. (2006), “High-Dimensional Graphs and Variable Selection with the Lasso,” *The Annals of Statistics*, 34, 1436–1462. [2687]
- Negahban, S., Yu, B., Wainwright, M. J., and Ravikumar, P. K. (2009), “A Unified Framework for High-Dimensional Analysis of m -estimators with Decomposable Regularizers,” in *Advances in Neural Information Processing Systems*, pp. 1348–1356. Citeseer. [2688]
- Ollier, E., and Viallon, V. (2017), “Regression Modelling on Stratified Data with the Lasso,” *Biometrika*, 104, 83–96. [2685]
- Pan, S. J., and Yang, Q. (2009), “A Survey on Transfer Learning,” *IEEE Transactions on Knowledge and Data Engineering*, 22, 1345–1359. [2684,2686]
- Reeve, H. W., Cannings, T. I., and Samworth, R. J. (2021), “Adaptive Transfer Learning,” *The Annals of Statistics*, 49, 3618–3649. [2684]
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288. [2684]
- Torrey, L., and Shavlik, J. (2010), “Transfer Learning,” in *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, eds. E. S. Olivas, J. D. M. Guerrero, M. M. Sober, J. R. M. Benedito, A. J. S. López, pp. 242–264. Hershey, PA: IGI Global. [2684,2685,2686]
- Van de Geer, S., Bühlmann, P., Ritov, Y., and Dezeure, R. (2014), “On Asymptotically Optimal Confidence Regions and Tests for High-dimensional Models,” *The Annals of Statistics*, 42, 1166–1202. [2685,2687,2691,2692]
- 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)*, pp. 617–626. IEEE. [2684]
- Weiss, K., Khoshgoftaar, T. M., and Wang, D. (2016), “A Survey of Transfer Learning,” *Journal of Big Data*, 3, 1–40. [2684,2686]
- Yuan, M., and Lin, Y. (2006), “Model Selection and Estimation in Regression with Grouped Variables,” *Journal of the Royal Statistical Society, Series B*, 68, 49–67. [2696]
- Zheng, C., Dasgupta, S., Xie, Y., Haris, A., and Chen, Y. Q. (2019), “On Data Enriched Logistic Regression,” arXiv preprint arXiv:1911.06380. [2685]