

SIHR: Statistical Inference in High-Dimensional Linear and Logistic Regression Models

by Prabrisha Rakshit, Zhenyu Wang, T. Tony Cai, and Zijian Guo

Abstract We introduce the R package **SIHR** for statistical inference in high-dimensional generalized linear models with continuous and binary outcomes. The package provides functionalities for constructing confidence intervals and performing hypothesis tests for low-dimensional objectives in both one-sample and two-sample regression settings. We illustrate the usage of **SIHR** through numerical examples and present real data applications to demonstrate the package's performance and practicality.

Introduction

In many applications, it is common to encounter regression problems where the number of covariates p exceeds the sample size n . Much progress has been made in point estimation and support recovery in high-dimensional generalized linear models (GLMs), as evidenced by works such as Bühlmann and van de Geer (2011); Negahban et al. (2009); Huang and Zhang (2012); Tibshirani (1996); Fan and Li (2011); Zhang (2010); Sun and Zhang (2012); Belloni et al. (2011); Meinshausen and Yu (2009). In particular, van de Geer et al. (2014); Javanmard and Montanari (2014); Zhang and Zhang (2014) have proposed methods to correct the bias of penalized regression estimators and construct confidence intervals (CIs) for individual regression coefficients of the high-dimensional linear model. Furthermore, Cai and Guo (2017) studied the minimaxity and adaptivity of CIs for linear functionals of the regression vector in high-dimensional linear models, while Cai et al. (2021b) proposed CIs and simultaneous hypothesis tests for individual regression coefficients in high-dimensional binary GLMs with general link functions. This debiased approach has sparked a rapidly growing research area focused on CI construction and hypothesis testing for low-dimensional objectives in high-dimensional GLMs.

In the current paper, we present the R package **SIHR**, which builds on the debiasing inference method and targets a wide range of inference problems in high-dimensional GLMs for both continuous and binary outcomes. We consider the high-dimensional GLMs: for $1 \leq i \leq n$,

$$\mathbb{E}(y_i | X_i) = f(X_i^\top \beta), \quad \text{with } f(z) = \begin{cases} z & \text{for linear model;} \\ \exp(z) / [1 + \exp(z)] & \text{for logistic model;} \end{cases} \quad (1)$$

where $\beta \in \mathbb{R}^p$ denotes the high-dimensional regression vector, $y_i \in \mathbb{R}$ and $X_i \in \mathbb{R}^p$ denote respectively the outcome and the measured covariates of the i -th observation. Throughout the paper, define $\Sigma = \mathbb{E}X_i X_i^\top$ and assume β to be a sparse vector with its sparsity level denoted as $\|\beta\|_0$. In addition to the one-sample setting, we examine the statistical inference methods for the two-sample regression models. Particularly, we generalize the regression model in (1) and consider:

$$\mathbb{E}(y_i^{(k)} | X_i^{(k)}) = f(X_i^{(k)\top} \beta^{(k)}) \quad \text{with } k = 1, 2 \text{ and } 1 \leq i \leq n_k \quad (2)$$

where $f(\cdot)$ is the pre-specified link function defined as (1), $\beta^{(k)} \in \mathbb{R}^p$ denotes the high-dimensional regression vector in k -th sample, $y_i^{(k)} \in \mathbb{R}$ and $X_i^{(k)} \in \mathbb{R}^p$ denote respectively the outcome and the measured covariates in the k -th sample.

The R package **SIHR** consists of five main functions `LF()`, `QF()`, `CATE()`, `InnProd()`, and `Dist()` implementing the statistical inferences for five different quantities correspondingly, under the one-sample model (1) or two-sample model (2).

1. `LF()`, abbreviated for linear functional, implements the inference approach for $x_{\text{new}}^\top \beta$ proposed in Cai et al. (2021a,b), with $x_{\text{new}} \in \mathbb{R}^p$ denoting a loading vector. With $x_{\text{new}} = e_j$ as a special case, `LF()` infers the regression coefficient β_j (van de Geer et al., 2014; Javanmard and Montanari, 2014; Zhang and Zhang, 2014, e.g.). When x_{new} denotes a future observation's covariates, `LF()` makes inferences for the conditional mean of the outcome for the individual. See the usage of `LF()` in the section [Linear functional](#).
2. `QF()`, abbreviated for quadratic functional, makes inferences for $\beta_G^\top A \beta_G$, following the proposal in Guo et al. (2019, 2021b); Cai and Guo (2020). $A \in \mathbb{R}^{|\mathcal{G}| \times |\mathcal{G}|}$ is either a pre-specified submatrix or the unknown $\Sigma_{\mathcal{G}, \mathcal{G}}$ and $\mathcal{G} \in \{1, \dots, p\}$ denotes the index set of interest; $\beta_G^\top A \beta_G$ can be viewed

as a total measure of all effects of variables in the group G . See the section [Quadratic functional](#) for the usage.

3. `CATE()`, abbreviated for conditional average treatment effect, is to make inference for $f(x_{\text{new}}^\top \beta^{(2)}) - f(x_{\text{new}}^\top \beta^{(1)})$, see [Cai et al. \(2021a\)](#) for detailed discussion. This difference measures the discrepancy between conditional means, closely related to the conditional average treatment effect for the new observation with covariates x_{new} . We demonstrate its usage in the section [Conditional average treatment effect](#).
4. `InnProd()`, abbreviated for inner products, implements the statistical inference for $\beta_G^{(1)\top} A \beta_G^{(2)}$ with $A \in \mathbb{R}^{|G| \times |G|}$, which was proposed in [Guo et al. \(2019\)](#); [Ma et al. \(2022\)](#). The inner products measure the similarity between the high-dimensional vectors $\beta^{(1)}$ and $\beta^{(2)}$, which is useful in capturing the genetic relatedness in the GWAS applications ([Guo et al., 2019](#); [Ma et al., 2022](#)). The usage is detailed in the section [Inner Product](#).
5. `Dist()`, short-handed for distance, makes inferences for the weighted distances $\gamma_G^\top A \gamma_G$ with $\gamma = \beta^{(2)} - \beta^{(1)}$. The distance measure is useful in comparing different high-dimensional regression vectors and constructing a generalizable model in the multisource learning problem [Guo et al. \(2023\)](#). See the section [Distance](#) for its usage.

There are a few other R packages for high-dimensional inference. The package `hdi` and the package `SSLASSO` implement the coordinate debiased Lasso estimators proposed in [van de Geer et al. \(2014\)](#) and [Javanmard and Montanari \(2014\)](#) respectively. These functions provide debiased estimators of β along with their standard error estimators, enabling confidence interval construction and hypothesis testing. However, these packages may suffer from heavy computational burden as multiple debiased optimizations are needed for inference on the linear functional. In contrast, our R package `SIHR` is computationally efficient as it performs debiasing only once. Additionally, `SIHR` targets a broader range of inference targets, including the single regression coefficient as a special case. The `DoubleML` package focuses on estimating low-dimensional parameters of interest, such as causal or treatment effect parameters, in the presence of high-dimensional nuisance parameters that can be estimated using machine learning methods, while our package aims to estimate arbitrary linear and weighted quadratic combinations of the coefficient vector in high-dimensional regression. Selective inference is implemented by the R package `selectiveInference`. They focus on parameters based on the selected model, while we focus on fixed parameters independent of the selected models. The method proposes a one-step estimator starting from the initial LASSO estimator based on the KKT conditions for the submodel selected by LASSO.

In the remainder of this paper, we provide a review of the inference methods in Section [Methodological Background](#), and introduce the main functions of the package in Section [Usage of the package](#), accompanied by illustrative examples. Finally, we demonstrate the application of our proposed methods to real data in Section [Applications](#).

Methodological Background

We briefly review the penalized maximum likelihood estimator of β in the high-dimensional GLM (1), defined as:

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \ell(\beta) + \lambda_0 \sum_{j=2}^p \frac{\|X_{\cdot j}\|_2}{\sqrt{n}} |\beta_j| \quad (3)$$

with $X_{\cdot j}$ denoting the j -th column of X , the first column of X set as the constant 1, and

$$\ell(\beta) = \begin{cases} \frac{1}{n} \sum_{i=1}^n (y_i - X_i^\top \beta)^2 & \text{for linear model} \\ -\frac{1}{n} \sum_{i=1}^n y_i \log \left[\frac{f(X_i^\top \beta)}{1 - f(X_i^\top \beta)} \right] - \frac{1}{n} \sum_{i=1}^n \log (1 - f(X_i^\top \beta)) & \text{for GLM with binary outcome} \end{cases} \quad (4)$$

The tuning parameter $\lambda_0 \asymp \sqrt{\log p/n}$ is chosen by cross-validation. In the penalized regression (3), we do not penalize the intercept coefficient β_1 . The penalized estimators have been shown to achieve the optimal convergence rates and satisfy desirable variable selection properties ([Meinshausen and Bühlmann, 2006](#); [Bickel et al., 2009](#); [Zhao and Yu, 2006](#); [Wainwright, 2009](#)). However, these estimators are not ready for statistical inference due to the non-negligible estimation bias induced by the penalty term ([van de Geer et al., 2014](#); [Javanmard and Montanari, 2014](#); [Zhang and Zhang, 2014](#)).

In section [Linear functional for GLM](#), we propose a unified inference method for $x_{\text{new}}^\top \beta$ under linear and logistic outcome models. We also discuss inferences for quadratic functionals $\beta_G^\top A \beta_G$ and $\beta_G^\top \Sigma_{G,G} \beta_G$ in section [Quadratic functional for GLM](#). In the case of the two-sample high-dimensional regression model (2), we develop the inference method for conditional treatment effect $\Delta(x_{\text{new}}) =$

$f(x_{\text{new}}^T \beta^{(2)}) - f(x_{\text{new}}^T \beta^{(1)})$ in section [Conditional average treatment effects](#); we consider inference for $\beta_G^{(1)T} A \beta_G^{(2)}$ and $\beta_G^{(1)T} \Sigma_{G,G} \beta_G^{(2)}$ in section [Inner product of regression vectors](#) and $\gamma_G^T A \gamma_G$ and $\gamma_G^T \Sigma_{G,G} \gamma_G$ with $\gamma = \beta^{(2)} - \beta^{(1)}$ in section [Distance of regression vectors](#).

Linear functional for linear model

To illustrate the idea of constructing the inference method, we start with the linear functional for the linear model, which will be extended to a unified version in the section [Linear functional for GLM](#). For the linear model in (1), we define $\epsilon_i = y_i - X_i^T \beta$ and rewrite the model as $y_i = X_i^T \beta + \epsilon_i$ for $1 \leq i \leq n$. Given the vector $x_{\text{new}} \in \mathbb{R}^p$, we construct the point estimator and the CI for $x_{\text{new}}^T \beta$.

A natural idea for the point estimator is to use the plug-in estimator $x_{\text{new}}^T \hat{\beta}$ with the penalized estimator $\hat{\beta}$ defined in (3). However, the bias $x_{\text{new}}^T (\hat{\beta} - \beta)$ is not negligible. The work [Cai et al. \(2021a\)](#) proposed the bias-corrected estimator as,

$$\widehat{x_{\text{new}}^T \beta} = x_{\text{new}}^T \hat{\beta} + \hat{u}^T \frac{1}{n} \sum_{i=1}^n X_i \cdot (y_i - X_i^T \hat{\beta}) \tag{5}$$

where the second term on the right hand side in (5) is the estimate of negative bias $-x_{\text{new}}^T (\hat{\beta} - \beta)$, and the projection direction \hat{u} is defined as

$$\hat{u} = \arg \min_{u \in \mathbb{R}^p} u^T \hat{\Sigma} u \quad \text{subject to: } \|\hat{\Sigma} u - x_{\text{new}}\|_{\infty} \leq \|x_{\text{new}}\|_2 \lambda \tag{6}$$

$$\left| x_{\text{new}}^T \hat{\Sigma} u - \|x_{\text{new}}\|_2^2 \right| \leq \|x_{\text{new}}\|_2^2 \lambda \tag{7}$$

where $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X_i \cdot X_i^T$ and $\lambda \asymp \sqrt{\log p/n}$. The bias-corrected estimator $\widehat{x_{\text{new}}^T \beta}$ satisfies the following error decomposition,

$$\widehat{x_{\text{new}}^T \beta} - x_{\text{new}}^T \beta = \underbrace{\hat{u}^T \frac{1}{n} \sum_{i=1}^n X_i^T \epsilon_i}_{\text{asym. normal}} + \underbrace{(\hat{\Sigma} \hat{u} - x_{\text{new}})^T (\beta - \hat{\beta})}_{\text{remaining bias}}$$

The first constraint in (6) controls the remaining bias term in the above equation while the second constraint in (7) is crucial to ensuring the asymptotic normality of $\widehat{x_{\text{new}}^T \beta} - x_{\text{new}}^T \beta$ for any vector x_{new} such that the variance of the ‘‘asym. normal’’ term always dominates the ‘‘remaining bias’’ term. Based on the asymptotic normality, we construct the CI for $x_{\text{new}}^T \beta$ as

$$\text{CI} = \left(\widehat{x_{\text{new}}^T \beta} - z_{\alpha/2} \sqrt{\hat{V}}, \widehat{x_{\text{new}}^T \beta} + z_{\alpha/2} \sqrt{\hat{V}} \right) \quad \text{with } \hat{V} = \frac{\hat{\sigma}^2}{n} \hat{u}^T \hat{\Sigma} \hat{u}$$

where $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (y_i - X_i^T \hat{\beta})^2$ and $z_{\alpha/2}$ denotes the upper $\alpha/2$ quantile for the standard normal distribution.

Linear functional for GLM

In this subsection, we generalize the inference method specifically for the linear model in [Linear functional for linear model](#) to GLM in (1). Given the initial estimator $\hat{\beta}$, the key step is to estimate the bias $x_{\text{new}}^T (\hat{\beta} - \beta)$. We can propose a unified version of the bias-corrected estimator for $x_{\text{new}}^T \beta$ as

$$\widehat{x_{\text{new}}^T \beta} = x_{\text{new}}^T \hat{\beta} + \hat{u}^T \frac{1}{n} \sum_{i=1}^n \omega(X_i^T \hat{\beta}) \left(y_i - f(X_i^T \hat{\beta}) \right) X_i. \tag{8}$$

with the second term on the right hand side of (8) being the estimate of $-x_{\text{new}}^T (\hat{\beta} - \beta)$. In consideration of different link functions $f(\cdot)$ in (1), we shall specify in the following how to construct the projection direction \hat{u} and the weight function $\omega : \mathbb{R} \mapsto \mathbb{R}$ in (8). In [Table 1](#), we consider different GLM models and present the corresponding functions $f(\cdot)$ and $\omega(\cdot)$, together with the derivative $f'(\cdot)$. Note that there are two ways of specifying the weights $w(z)$ for logistic regression. The linearization weighting is proposed in [Guo et al. \(2021b\)](#) specifically for logistic regression; while [Cai et al. \(2021b\)](#) constructed the link-specific weighting method for general link function $f(\cdot)$. The projection direction $\hat{u} \in \mathbb{R}^p$ in

Model	Outcome Type	$f(z)$	$f'(z)$	$\omega(z)$	Weighting
linear	Continuous	z	1	1	
logistic	Binary	$\frac{e^z}{1+e^z}$	$\frac{e^z}{(1+e^z)^2}$	$\frac{(1+e^z)^2}{e^z}$	Linearization
logistic_alter	Binary	$\frac{e^z}{1+e^z}$	$\frac{e^z}{(1+e^z)^2}$	1	Link-specific

Table 1: Definitions of the functions ω and f for different GLMs.

(8) is constructed as follows:

$$\begin{aligned} \hat{u} = \arg \min_{u \in \mathbb{R}^p} u^\top & \left[\frac{1}{n} \sum_{i=1}^n \omega(X_i^\top \hat{\beta}) f'(X_i^\top \hat{\beta}) X_i X_i^\top \right] u \quad \text{subject to:} \\ \left\| \frac{1}{n} \sum_{i=1}^n \omega(X_i^\top \hat{\beta}) f'(X_i^\top \hat{\beta}) X_i X_i^\top u - x_{\text{new}} \right\|_\infty & \leq \|x_{\text{new}}\|_2 \lambda \\ \left| x_{\text{new}}^\top \frac{1}{n} \sum_{i=1}^n \omega(X_i^\top \hat{\beta}) f'(X_i^\top \hat{\beta}) X_i X_i^\top u - \|x_{\text{new}}\|_2^2 \right| & \leq \|x_{\text{new}}\|_2^2 \lambda. \end{aligned} \tag{9}$$

It has been established that $\widehat{x_{\text{new}}^\top \beta}$ in (8) is asymptotically unbiased and normal for the linear model (Cai et al., 2021a), the logistic model (Guo et al., 2021a; Cai et al., 2021b), and the probit model (Cai et al., 2021b). The variance of $\widehat{x_{\text{new}}^\top \beta}$ can be estimated by \widehat{V} , defined as

$$\widehat{V} = \widehat{u}^\top \left[\frac{1}{n^2} \sum_{i=1}^n \left(\omega(X_i^\top \hat{\beta}) \right)^2 \widehat{\sigma}_i^2 X_i X_i^\top \right] \widehat{u} \quad \text{with :} \tag{10}$$

$$\widehat{\sigma}_i^2 = \begin{cases} \frac{1}{n} \sum_{j=1}^n (y_j - X_j^\top \hat{\beta})^2, & \text{for linear model} \\ f(X_i^\top \hat{\beta})(1 - f(X_i^\top \hat{\beta})), & \text{for GLM with binary outcome} \end{cases} \tag{11}$$

Based on the asymptotic normality, the CI for $x_{\text{new}}^\top \beta$ is:

$$\text{CI} = \left(\widehat{x_{\text{new}}^\top \beta} - z_{\alpha/2} \sqrt{\widehat{V}}, \widehat{x_{\text{new}}^\top \beta} + z_{\alpha/2} \sqrt{\widehat{V}} \right).$$

Subsequently, for the binary outcome case, we estimate the case probability $\mathbb{P}(y_i = 1 \mid X_i = x_{\text{new}})$ by $f(\widehat{x_{\text{new}}^\top \beta})$ and construct the CI for $f(x_{\text{new}}^\top \beta)$ as:

$$\text{CI} = \left(f \left(\widehat{x_{\text{new}}^\top \beta} - z_{\alpha/2} \sqrt{\widehat{V}} \right), f \left(\widehat{x_{\text{new}}^\top \beta} + z_{\alpha/2} \sqrt{\widehat{V}} \right) \right).$$

Quadratic functional for GLM

We now move our focus to inference for the quadratic functional $Q_A = \beta_G^\top A \beta_G$, where $G \subset \{1, \dots, p\}$ and $A \in \mathbb{R}^{|G| \times |G|}$ denotes a pre-specified matrix of interest. Without loss of generality, we set $G = \{1, 2, \dots, |G|\}$. In the following, we propose a unified version of the point estimator and CI under the GLM (1). With the initial estimator $\hat{\beta}$ defined in (3), the plug-in estimator $\widehat{\beta}_G^\top A \widehat{\beta}_G$ suffers from the following error,

$$\widehat{\beta}_G^\top A \widehat{\beta}_G - \beta_G^\top A \beta_G = 2\widehat{\beta}_G^\top A (\widehat{\beta}_G - \beta_G) - (\widehat{\beta}_G - \beta_G)^\top A (\widehat{\beta}_G - \beta_G).$$

The last term in the above decomposition $(\widehat{\beta}_G - \beta_G)^\top A (\widehat{\beta}_G - \beta_G)$ is the higher-order approximation error under regular conditions; thus the bias mainly comes from the term $2\widehat{\beta}_G^\top A (\widehat{\beta}_G - \beta_G)$, which can be expressed as $2x_{\text{new}}^\top (\widehat{\beta} - \beta)$ with $x_{\text{new}} = (\widehat{\beta}_G^\top A, \mathbf{0})^\top$. Hence the term can be estimated directly by applying the linear functional approach in section [Linear functional for GLM](#). Utilizing this idea, Guo et al. (2021b, 2019) proposed the following estimator of Q_A ,

$$\widehat{Q}_A = \widehat{\beta}_G^\top A \widehat{\beta}_G + 2\widehat{u}_A^\top \left[\frac{1}{n} \sum_{i=1}^n \omega(X_i^\top \hat{\beta}) (y_i - f(X_i^\top \hat{\beta})) X_i \right]$$

with the second term being the estimate of $-2\widehat{\beta}_G^\top A (\widehat{\beta}_G - \beta_G)$, where \widehat{u}_A is the projection direction defined in (9) with $x_{\text{new}} = (\widehat{\beta}_G^\top A, \mathbf{0}^\top)^\top$. Since Q_A is non-negative if A is positive semi-definite, we

truncate \widehat{Q}_A at 0 and define $\widehat{Q}_A = \max(\widehat{Q}_A, 0)$. We further estimate the variance of the \widehat{Q}_A by

$$\widehat{V}_A(\tau) = 4\widehat{u}^\top \left[\frac{1}{n^2} \sum_{i=1}^n \omega^2(X_i^\top \widehat{\beta}) \widehat{\sigma}_i^2 X_i X_i^\top \right] \widehat{u} + \frac{\tau}{n} \tag{12}$$

where the term τ/n with $\tau > 0$ (default value $\tau = 1$) is introduced as an upper bound for the term $(\widehat{\beta}_G - \beta_G)^\top A(\widehat{\beta}_G - \beta_G)$, and $\widehat{\sigma}_i^2$ is defined in (11). Then given a fixed value of τ , we construct the CI as $CI(\tau) = \left(\max\left(\widehat{Q}_A - z_{\alpha/2} \sqrt{\widehat{V}_A(\tau)}, 0\right), \widehat{Q}_A + z_{\alpha/2} \sqrt{\widehat{V}_A(\tau)} \right)$.

Now we turn to the estimation of $Q_\Sigma = \beta_G^\top \Sigma_{G,G} \beta_G$ where the matrix $\Sigma_{G,G}$ is unknown and estimated by $\widehat{\Sigma}_{G,G} = \frac{1}{n} \sum_{i=1}^n X_{iG} X_{iG}^\top$. Decompose the error of the plug-in estimator $\widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} \widehat{\beta}_G$:

$$\widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} \widehat{\beta}_G - \beta_G^\top \Sigma_{G,G} \beta_G = 2 \widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} (\widehat{\beta}_G - \beta_G) + \beta_G^\top (\widehat{\Sigma}_{G,G} - \Sigma_{G,G}) \beta_G - (\widehat{\beta}_G - \beta_G)^\top \widehat{\Sigma}_{G,G} (\widehat{\beta}_G - \beta_G).$$

The first term $\widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} (\widehat{\beta}_G - \beta_G)$ is estimated by applying linear functional approach in **Linear functional for GLM** with $x_{\text{new}} = (\widehat{\beta}_G^\top \widehat{\Sigma}_{G,G}, \mathbf{0})^\top$; the second term $\beta_G^\top (\widehat{\Sigma}_{G,G} - \Sigma_{G,G}) \beta_G$ can be controlled asymptotically by central limit theorem; and the last term $(\widehat{\beta}_G - \beta_G)^\top \widehat{\Sigma}_{G,G} (\widehat{\beta}_G - \beta_G)$ is negligible due to high-order bias. [Guo et al. \(2021b\)](#) proposed the following estimator of Q_Σ

$$\widehat{Q}_\Sigma = \widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} \widehat{\beta}_G + 2 \widehat{u}_\Sigma^\top \left[\frac{1}{n} \sum_{i=1}^n \omega(X_i^\top \widehat{\beta}) (y_i - f(X_i^\top \widehat{\beta})) X_i \right]$$

where \widehat{u}_Σ is the projection direction constructed in (9) with $x_{\text{new}} = (\widehat{\beta}_G^\top \widehat{\Sigma}_{G,G}, \mathbf{0})^\top$. We introduce the estimator $\widehat{Q}_\Sigma = \max(\widehat{Q}_\Sigma, 0)$ and estimate its variance as

$$\widehat{V}_\Sigma(\tau) = 4\widehat{u}_\Sigma^\top \left[\frac{1}{n^2} \sum_{i=1}^n \omega^2(X_i^\top \widehat{\beta}) \widehat{\sigma}_i^2 X_i X_i^\top \right] \widehat{u}_\Sigma + \frac{1}{n^2} \sum_{i=1}^n \left(\widehat{\beta}_G^\top X_{iG} X_{iG}^\top \widehat{\beta}_G - \widehat{\beta}_G^\top \widehat{\Sigma}_{G,G} \widehat{\beta}_G \right)^2 + \frac{\tau}{n} \tag{13}$$

where $\tau > 0$, the term τ/n is introduced as an upper bound for the term $(\widehat{\beta}_G - \beta_G)^\top \widehat{\Sigma}_{G,G} (\widehat{\beta}_G - \beta_G)$, and $\widehat{\sigma}_i^2$ is defined in (11). Then, thanks to the asymptotic normality, for a fixed value of τ , we can construct the CI as

$$CI(\tau) = \left(\max\left(\widehat{Q}_\Sigma - z_{\alpha/2} \sqrt{\widehat{V}_\Sigma(\tau)}, 0\right), \widehat{Q}_\Sigma + z_{\alpha/2} \sqrt{\widehat{V}_\Sigma(\tau)} \right)$$

Conditional average treatment effects

The inference methods proposed for one sample can be generalized to make inferences for conditional average treatment effects, which can be expressed as the difference between two linear functionals. Let $A_i \in \{1, 2\}$ denote the treatment assignment for i -th observation. Consider the two-sample GLMs as

$$\mathbb{E}(y_i | X_i, A_i = 1) = f(X_i^\top \beta^{(1)}) \quad \text{and} \quad \mathbb{E}(y_i | X_i, A_i = 2) = f(X_i^\top \beta^{(2)})$$

where f is the link function listed in table 1. Then, for a future individual $X_i = x_{\text{new}}$, we define $\Delta(x_{\text{new}}) = \mathbb{E}(y_i | X_i, A_i = 2) - \mathbb{E}(y_i | X_i, A_i = 1)$, that measures the difference of the conditional mean of assignment of treatment for the individual with covariates x_{new} .

Following (8), we construct the bias-corrected point estimators of $x_{\text{new}}^\top \widehat{\beta}^{(1)}$ and $x_{\text{new}}^\top \widehat{\beta}^{(2)}$, together with their corresponding variance $\widehat{V}_{(1)}$ and $\widehat{V}_{(2)}$ as (10). The paper [Cai et al. \(2021a\)](#) proposed to estimate $\Delta(x_{\text{new}})$ by $\widehat{\Delta}(x_{\text{new}})$ as:

$$\widehat{\Delta}(x_{\text{new}}) = f(x_{\text{new}}^\top \widehat{\beta}^{(2)}) - f(x_{\text{new}}^\top \widehat{\beta}^{(1)})$$

Its variance can be estimated with delta method by:

$$\widehat{V}_\Delta = \left(f'(x_{\text{new}}^\top \widehat{\beta}^{(1)}) \right)^2 \widehat{V}_{(1)} + \left(f'(x_{\text{new}}^\top \widehat{\beta}^{(2)}) \right)^2 \widehat{V}_{(2)}$$

Then we construct the CI as $CI = \left(\widehat{\Delta}(x_{\text{new}}) - z_{\alpha/2} \sqrt{\widehat{V}_\Delta}, \widehat{\Delta}(x_{\text{new}}) + z_{\alpha/2} \sqrt{\widehat{V}_\Delta} \right)$.

Inner product of regression vectors

The paper Guo et al. (2019); Ma et al. (2022) have carefully investigated the CI construction for $\beta_G^{(1)\top} A \beta_G^{(2)}$, provided with a pre-specified submatrix $A \in \mathbb{R}^{|\mathcal{G}| \times |\mathcal{G}|}$ and the set of indices $\mathcal{G} \in \{1, \dots, p\}$. Let $\hat{\beta}^{(1)}$ and $\hat{\beta}^{(2)}$ respectively be the initial estimators for their corresponding sample in (2), the plug-in but biased estimator is $\hat{\beta}_G^{(1)\top} A \hat{\beta}_G^{(2)}$. Its bias can be decomposed as:

$$\begin{aligned} \hat{\beta}_G^{(1)\top} A \hat{\beta}_G^{(2)} - \beta_G^{(1)\top} A \beta_G^{(2)} &= \hat{\beta}_G^{(2)\top} A (\hat{\beta}_G^{(1)} - \beta_G^{(1)}) + \hat{\beta}_G^{(1)\top} A (\hat{\beta}_G^{(2)} - \beta_G^{(2)}) \\ &\quad - (\hat{\beta}_G^{(1)} - \beta_G^{(1)})^\top A (\hat{\beta}_G^{(2)} - \beta_G^{(2)}). \end{aligned}$$

The key step is to estimate the error components $\hat{\beta}_G^{(2)\top} A (\hat{\beta}_G^{(1)} - \beta_G^{(1)})$ and $\hat{\beta}_G^{(1)\top} A (\hat{\beta}_G^{(2)} - \beta_G^{(2)})$. Then the following procedures can be interpreted as applying Linear Functional twice on two independent samples. To be specific, we propose the following bias-corrected estimator for $\beta_G^{(1)\top} A \beta_G^{(2)}$

$$\begin{aligned} \widehat{\beta}_G^{(1)\top} A \widehat{\beta}_G^{(2)} &= \hat{\beta}_G^{(1)\top} A \hat{\beta}_G^{(2)} + \hat{u}_1^\top \frac{1}{n_1} \sum_{i=1}^{n_1} \omega(X_i^{(1)\top} \hat{\beta}^{(1)}) (y_i^{(1)} - f(X_i^{(1)\top} \hat{\beta}^{(1)})) X_i^{(1)} \\ &\quad + \hat{u}_2^\top \frac{1}{n_2} \sum_{i=1}^{n_2} \omega(X_i^{(2)\top} \hat{\beta}^{(2)}) (y_i^{(2)} - f(X_i^{(2)\top} \hat{\beta}^{(2)})) X_i^{(2)} \end{aligned} \tag{14}$$

with the second term and the third term in right-hand-side of (14) estimating $-\hat{\beta}_G^{(2)\top} A (\hat{\beta}_G^{(1)} - \beta_G^{(1)})$ and $-\hat{\beta}_G^{(1)\top} A (\hat{\beta}_G^{(2)} - \beta_G^{(2)})$ respectively, where \hat{u}_1 is the projection direction defined in (9) with $x_{\text{new}} = (\hat{\beta}_G^{(2)\top} A, \mathbf{0})^\top$ and \hat{u}_2 is the projection direction defined in (9) with $x_{\text{new}} = (\hat{\beta}_G^{(1)\top} A, \mathbf{0})^\top$. The corresponding variance of $\widehat{\beta}_G^{(1)\top} A \widehat{\beta}_G^{(2)}$, when A is a known positive definite matrix, is estimated as

$$\widehat{V}_A(\tau) = \widehat{V}^{(1)} + \widehat{V}^{(2)} + \frac{\tau}{\min(n_1, n_2)}$$

where $\widehat{V}^{(k)}$ is computed as (10) for the k -th regression model ($k = 1, 2$) in (2) and $\tau > 0$, the term $\tau / \min(n_1, n_2)$ is introduced as an upper bound for the term $(\hat{\beta}_G^{(1)} - \beta_G^{(1)})^\top A (\hat{\beta}_G^{(2)} - \beta_G^{(2)})$.

When A is not specified, we treat $A = \Sigma_{\mathcal{G}, \mathcal{G}}$, which is unknown. As a natural generalization, the quantity $\beta_G^{(1)\top} \Sigma_{\mathcal{G}, \mathcal{G}} \beta_G^{(2)}$ is well defined if the two regression models in (2) share the design covariance matrix $\Sigma = \mathbb{E} X_i^{(1)} X_i^{(1)\top} = \mathbb{E} X_i^{(2)} X_i^{(2)\top}$. We follow the above procedures replacing A by $\widehat{\Sigma}_{\mathcal{G}, \mathcal{G}} = \frac{1}{n_1 + n_2} \sum_{i=1}^{n_1 + n_2} X_{i, \mathcal{G}} X_{i, \mathcal{G}}^\top$ where X is the row-combined matrix of $X^{(1)}$ and $X^{(2)}$. The variance of $\widehat{\beta}_G^{(1)\top} \widehat{\Sigma}_{\mathcal{G}, \mathcal{G}} \widehat{\beta}_G^{(2)}$ is now estimated as

$$\widehat{V}_\Sigma(\tau) = \widehat{V}^{(1)} + \widehat{V}^{(2)} + \frac{1}{(n_1 + n_2)^2} \sum_{i=1}^{n_1 + n_2} \left(\widehat{\beta}_G^{(1)\top} X_{i, \mathcal{G}} X_{i, \mathcal{G}}^\top \widehat{\beta}_G^{(2)} - \widehat{\beta}_G^{(1)\top} \widehat{\Sigma}_{\mathcal{G}, \mathcal{G}} \widehat{\beta}_G^{(2)} \right)^2 + \frac{\tau}{\min(n_1, n_2)}$$

Depending on whether the submatrix A is specified or not, the CI is

$$\text{CI}(\tau) = \begin{cases} \left(\widehat{\beta}_G^{(1)\top} A \widehat{\beta}_G^{(2)} - z_{\alpha/2} \widehat{V}_A(\tau), \widehat{\beta}_G^{(1)\top} A \widehat{\beta}_G^{(2)} + z_{\alpha/2} \widehat{V}_A(\tau) \right) & \text{if } A \text{ is specified} \\ \left(\widehat{\beta}_G^{(1)\top} \widehat{\Sigma}_{\mathcal{G}, \mathcal{G}} \widehat{\beta}_G^{(2)} - z_{\alpha/2} \widehat{V}_\Sigma(\tau), \widehat{\beta}_G^{(1)\top} \widehat{\Sigma}_{\mathcal{G}, \mathcal{G}} \widehat{\beta}_G^{(2)} + z_{\alpha/2} \widehat{V}_\Sigma(\tau) \right) & \text{otherwise} \end{cases}$$

Distance of regression vectors

We denote $\gamma = \beta^{(2)} - \beta^{(1)}$ and its initial estimator $\hat{\gamma} = \hat{\beta}^{(2)} - \hat{\beta}^{(1)}$. The quantity of interest is the distance between two regression vectors $\gamma_G^\top A \gamma_G$, given a pre-specified submatrix $A \in \mathbb{R}^{|\mathcal{G}| \times |\mathcal{G}|}$ and the set of indices $\mathcal{G} \in \{1, \dots, p\}$. The bias of the plug-in estimator $\hat{\gamma}_G^\top A \hat{\gamma}_G$ is:

$$\hat{\gamma}_G^\top A \hat{\gamma}_G - \gamma_G^\top A \gamma_G = 2 \hat{\gamma}_G^\top A (\hat{\beta}_G^{(2)} - \beta_G^{(2)}) - 2 \hat{\gamma}_G^\top A (\hat{\beta}_G^{(1)} - \beta_G^{(1)}) - (\hat{\gamma}_G - \gamma_G)^\top A (\hat{\gamma}_G - \gamma_G)$$

The key step is to estimate the error components $\hat{\gamma}_G^\top A (\hat{\beta}_G^{(1)} - \beta_G^{(1)})$ and $\hat{\gamma}_G^\top A (\hat{\beta}_G^{(2)} - \beta_G^{(2)})$ in the above decomposition. We apply linear functional techniques twice here, and propose the bias-corrected

estimator:

$$\begin{aligned} \widehat{\gamma_G^T A \gamma_G} &= \widehat{\gamma_G^T} A \widehat{\gamma_G} - 2 \widehat{u}_1^T \frac{1}{n_1} \sum_{i=1}^{n_1} \omega(X_i^{(1)T} \widehat{\beta}^{(1)}) \left(y_i^{(1)} - f(X_i^{(1)T} \widehat{\beta}^{(1)}) \right) X_i^{(1)} \\ &\quad + 2 \widehat{u}_2^T \frac{1}{n_2} \sum_{i=1}^{n_2} \omega(X_i^{(2)T} \widehat{\beta}^{(2)}) \left(y_i^{(2)} - f(X_i^{(2)T} \widehat{\beta}^{(2)}) \right) X_i^{(2)} \end{aligned} \quad (15)$$

Then by non-negative distance, we define $\widehat{\gamma_G^T A \gamma_G} = \max \left\{ \widehat{\gamma_G^T} A \widehat{\gamma_G}, 0 \right\}$. The second term on right-hand-side of (15) is to estimate $-2 x_{\text{new}}^T (\widehat{\beta}_G^{(1)} - \beta_G^{(1)})$ with $x_{\text{new}} = (\widehat{\gamma}_G^T A, \mathbf{0})^T$; and the third term on right-hand-side of (15) is to estimate $-2 x_{\text{new}}^T (\widehat{\beta}_G^{(2)} - \beta_G^{(2)})$ with $x_{\text{new}} = (\widehat{\gamma}_G^T A, \mathbf{0})^T$ as well. The corresponding asymptotic variance for the bias-corrected estimator is

$$\widehat{V}_A(\tau) = 4 \widehat{V}^{(1)} + 4 \widehat{V}^{(2)} + \frac{\tau}{\min(n_1, n_2)}$$

where $\widehat{V}^{(k)}$ is computed as (10) for the k -th regression model ($k = 1, 2$) and $\tau > 0$, the term $\tau / \min(n_1, n_2)$ is introduced as an upper bound for the term $(\widehat{\gamma}_G - \gamma_G)^T A (\widehat{\gamma}_G - \gamma_G)$. With asymptotic normality, we construct the CI

$$\text{CI}(\tau) = \left(\max \left(\widehat{\gamma_G^T} A \widehat{\gamma_G} - z_{\alpha/2} \sqrt{\widehat{V}_A(\tau)}, 0 \right), \widehat{\gamma_G^T} A \widehat{\gamma_G} + z_{\alpha/2} \sqrt{\widehat{V}_A(\tau)} \right).$$

When the submatrix A is not specified, we treat $A = \Sigma_{G,G}$, which is unknown. The point estimator $\widehat{\gamma_G^T \Sigma_{G,G} \gamma_G}$ can be computed similarly as outlined in (15). In this case, the submatrix A is substituted with $\widehat{\Sigma}_{G,G}$ and the resulting value is truncated at 0, where $\widehat{\Sigma}_{G,G} = \frac{1}{n_1+n_2} \sum_{i=1}^{n_1+n_2} X_{i,G} X_{i,G}^T$ with X as the row-combined matrix of $X^{(1)}$ and $X^{(2)}$. Its corresponding asymptotic variance is

$$\widehat{V}_\Sigma = 4 \widehat{V}^{(1)} + 4 \widehat{V}^{(2)} + \frac{1}{(n_1 + n_2)^2} \sum_{i=1}^{n_1+n_2} \left(\widehat{\gamma}_G^T X_{i,G} X_{i,G}^T \widehat{\gamma}_G - \widehat{\gamma}_G^T \widehat{\Sigma}_{G,G} \widehat{\gamma}_G \right)^2 + \frac{\tau}{\min(n_1, n_2)}$$

Next we present its CI

$$\text{CI}(\tau) = \left(\max \left(\widehat{\gamma_G^T \Sigma} \widehat{\gamma}_G - z_{\alpha/2} \sqrt{\widehat{V}_\Sigma(\tau)}, 0 \right), \widehat{\gamma_G^T \Sigma} \widehat{\gamma}_G + z_{\alpha/2} \sqrt{\widehat{V}_\Sigma(\tau)} \right)$$

Usage of the package

The **SIHR** package contains a set of functions for inference methods of various low-dimensional objectives, such as linear and quadratic functions. See the table 2 for each function and its corresponding objective.

Function	Objective	Description
LF()	$x_{\text{new}}^T \beta$	Generate an LF object.
QF()	$\beta_G^T A \beta_G$	Generate a QF object.
CATE()	$f(x_{\text{new}}^T \beta^{(2)}) - f(x_{\text{new}}^T \beta^{(1)})$	Generate a CATE object.
InnProd()	$\beta_G^{(1)T} A \beta_G^{(2)}$	Generate an InnProd object.
Dist()	$\gamma_G^T A \gamma_G$ with $\gamma = \beta^{(1)} - \beta^{(0)}$	Generate a Dist object.
ci()		Input object, return CIs.
summary()		Input object, compute and return a list of summary statistics, including bias-corrected point estimators, standard error and so on.

Table 2: Functions of **SIHR**

Linear functional

The function `LF()`, shorthand for Linear Functional, performs inference for $x_{\text{new}}^T \beta$ under the high-dimensional model (1). A typical `LF()` code snippet looks like:

```
LF(X,y,loading.mat,model=c("linear","logistic","logistic_alter"),intercept=TRUE,
  intercept.loading=FALSE,beta.init=NULL,lambda=NULL,mu=NULL,prob.filter=0.05,
  rescale=1.1,alpha=0.05,verbose=FALSE)
```

The argument `loading.mat` takes values of x_{new} as a matrix, which allows for multiple x_{new} as the input, with each column representing a new future observation $x_{\text{new}} \in \mathbb{R}^p$. The argument `model` specifies what regression model the algorithm is working on, which can take "linear", "logistic", "logistic_alter", corresponding to the Table 1. The argument `intercept.loading` is logical, specifying whether the intercept term should be included or not for defining the objective $x_{\text{new}}^T \beta$ and the default value is FALSE. More detailed descriptions of each input argument in `LF()` function can be found in Table 3. In the following code, we make inference for $x_{\text{new}}^T \beta$ with simulated data when the outcome y_i is continuous.

Argument	Description	Default
X	Design matrix, of dimension $n \times p$	
y	Outcome vector, of length n	
loading.mat	Loading matrix each column corresponds to a loading of interest x_{new}	
model	The regression model to fit, one of "linear", "logistic" and "logistic_alter"	"linear"
intercept	Should intercept be fitted for the initial estimator	TRUE
intercept.loading	Should intercept term be included for the inference of objective.	FALSE
beta.init	The initial estimator of the regression vector	NULL
lambda	The tuning parameter in fitting initial model. If not specified, it will be picked by cross-validation.	NULL
mu	The dual tuning parameter used in the construction of the projection direction. If not specified, it will be searched automatically.	NULL
prob.filter	The threshold of estimated probabilities for filtering observations for binary outcome.	0.05
rescale	The factor to enlarge the standard error to account for the finite sample bias.	1.1
alpha	Level of significance to construct two-sided CI	0.05
verbose	Should intermediate message(s) be printed, the projection direction be returned.	FALSE

Table 3: Arguments for function `LF()`

Example 1. For $1 \leq i \leq n$, the covariates X_i are independently generated from the multivariate normal distribution with mean $\mu = 0_p$ and covariance $\Sigma = \mathbf{I}_p$. The outcome is generated as $y_i = a_0 + X_i^T \beta + \epsilon_i$ with standard normal noise. Given two further observations $x_{\text{new}}^{(1)}, x_{\text{new}}^{(2)}$, we're going to make inference for $x_{\text{new}}^{(1)T} \beta$ and $x_{\text{new}}^{(2)T} \beta$ simultaneously.

```
## Data Preparation ##
set.seed(0)
n = 100; p = 120
mu = rep(0,p); Cov = diag(p)
a0 = -0.5
beta = rep(0,p); beta[c(1,2)] = c(0.5, 1)
X = MASS::mvrnorm(n, mu, Cov)
y = a0 + X %*% beta + rnorm(n)

## two further observations ##
loading1 = c(1, 1, rep(0, p-2))
loading2 = c(-0.5, -1, rep(0, p-2))
loading.mat = cbind(loading1, loading2)

## Linear Functional ##
Est = LF(X, y, loading.mat, model='linear')
```

Having fitted the model, we have two following functions `ci()` and `summary()`.

```
ci(Est)
#> loading      lower      upper
```



```
#>1      1  1.167873  1.8753934
#>2      2 -1.544138 -0.7995375
```

In the above result, we can find the 95% CI for $x_{\text{new}}^{(1)\top}\beta$ and $x_{\text{new}}^{(2)\top}\beta$. Both true values $x_{\text{new}}^{(1)\top}\beta = 1.5$ and $x_{\text{new}}^{(2)\top}\beta = -1.25$ lie in the corresponding CIs.

```
summary(Est)
#>Call:
#>Inference for Linear Functional
#>
#>Estimators:
#> loading est.plugin est.debias Std. Error z value Pr(>|z|)
#>      1      1.268      1.522      0.1805      8.430 0.000e+00 ***
#>      2     -1.033     -1.172      0.1900     -6.169 6.868e-10 ***
```

`summary()` returns a list of the summary statistics, in which we can find the plugin estimator, bias-corrected estimator, and the standard error for the bias-corrected estimator. The bias-corrected estimators are closer to the true values.

As a second example, we consider the logistic regression where the argument `model` is set as "logistic" or "logistic_alter". To boost computation efficiency, we may specify the argument `beta.init` as the common initial coefficients estimators for all further observations.

Example 2. For $1 \leq i \leq n$, the covariates X_i are independently generated from the multivariate normal distribution with mean $\mu = 0_p$ and covariance $\Sigma = I_p$. We generate the outcome following the model $Y_i \sim \text{Bernoulli}(f(a_0 + X_i^\top \beta))$ with $f(z) = \exp(z)/(1 + \exp(z))$.

```
## Data Preparation ##
set.seed(0)
n = 300; p = 120
mu = rep(0,p); Cov = diag(p)
a0 = -1
beta = rep(0,p); beta[c(1,2)] = c(1, 1)
X = MASS::mvrnorm(n, mu, Cov)
val = a0 + X %*% beta
y = rbinom(n, 1, exp(val)/(1+exp(val)))

## two further observations ##
loading1 = c(1, 1, rep(0, p-2))
loading2 = c(-0.5, -2, rep(0, p-2))
loading.mat = cbind(loading1, loading2)

## obtain initial estimators ##
cv.fit = glmnet::cv.glmnet(X, y, family='binomial', alpha=1, standardize=TRUE)
beta.init = as.vector(coef(cv.fit, s=cv.fit[['lambda.min']]))
Est = LF(X, y, loading.mat, model='logistic', beta.init=beta.init)
```

The corresponding CIs and summary statistics are given below:

```
ci(Est)
#> loading      lower      upper
#>1      1  1.257559  2.492327
#>2      2 -3.186513 -1.605671
```

Consequently, we have two objective values $x_{\text{new}}^{(1)\top}\beta = 2$ and $x_{\text{new}}^{(2)\top}\beta = -2.5$. Both of these values lie within their corresponding 95% CIs.

```
summary(Est)
#> Call:
#> Inference for Linear Functional
#>
#> Estimators:
#> loading est.plugin est.debias Std. Error z value Pr(>|z|)
#>      1      1.340      1.875      0.3150      5.952 2.645e-09 ***
#>      2     -1.741     -2.396      0.4033     -5.941 2.825e-09 ***
```

Note that the plugin estimators $x_{\text{new}}^{(1)\top}\hat{\beta}$ and $x_{\text{new}}^{(2)\top}\hat{\beta}$ are severely biased in such setting, the proposed bias-correction approach significantly saves the bias with $\widehat{x_{\text{new}}^{(1)\top}\beta}$ and $\widehat{x_{\text{new}}^{(2)\top}\beta}$.

Quadratic functional

The function `QF()`, abbreviated for Quadratic Functional, conducts inference for $\beta_G^T A \beta_G$ if A is the submatrix pre-specified or $\beta_G^T \Sigma_{G,G} \beta_G$ under the high-dimensional regression model (1). The function `QF()` can be called with the following arguments.

```
QF(X,y,G,A=NULL,model=c("linear","logistic","logistic_alter"),intercept=TRUE,
  beta.init=NULL,split=TRUE,lambda=NULL,mu=NULL,prob.filter=0.05,rescale=1.1,
  tau=c(0.25,0.5,1),alpha=0.05,verbose=FALSE)
```

The argument `G` is the set of indices of interest. If the argument `A` is specified, it will conduct inference for $\beta_G^T A \beta_G$; otherwise, it will turn to $\beta_G^T \Sigma_{G,G} \beta_G$. The argument `model` specifies what regression model the algorithm is working on, which can take "linear", "logistic", "logistic_alter", corresponding to Table 1. The argument `split` indicates whether we conduct the sample splitting for computing the initial estimator of regression coefficients. When `split=TRUE`, the initial estimator of regression coefficients is computed using half of the available observations while the remaining half is used for bias correction. The option of using sampling splitting might require a larger sample size. The argument `tau.vec` allows the user to supply a vector of possible values for τ in (12) and (13). Table 4 presents the arguments for `QF()` while excluding those that are repeated in Table 3.

Argument	Description	Default
<code>X</code>	Design matrix, of dimension $n \times p$	
<code>y</code>	Outcome vector, of length n	
<code>G</code>	The set of indices in the quadratic form	
<code>A</code>	The matrix A in the quadratic form, of dimension $ G \times G $. If not specified, A would be set as the $ G \times G $ submatrix of the population covariance matrix corresponding to the index set G	NULL
<code>model</code>	The regression model to fit, one of "linear", "logistic" and "logistic_alter"	"linear"
<code>intercept</code>	Should intercept be fitted for the initial estimator	TRUE
<code>split</code>	Sampling splitting or not for computing the initial estimator. It takes effect only when <code>beta.init = NULL</code> .	TRUE
<code>tau</code>	The enlargement factor for asymptotic variance of the bias-corrected estimator to handle super-efficiency. It allows for a scalar or vector.	<code>c(0.25, 0.5, 1)</code>

Table 4: Arguments for function `QF()`, others are repeated in Table 3

In the third example, we illustrate the usage of `QF()` in linear regression model,

Example 3. For $1 \leq i \leq n$, the covariates X_i is generated from multivariate normal distribution with mean $\mu = 0_p$ and covariance $\Sigma \in \mathbb{R}^{p \times p}$ where $\Sigma_{j,k} = 0.5^{|j-k|}$. We generate the outcome following the model $y_i = X_i \beta + \epsilon_i$ with standard normal distributed noise. We're going to make inference for $\beta_G^T \Sigma_{G,G} \beta_G$ with $G = \{40, \dots, 60\}$.

```
## Data Preparation ##
set.seed(0)
n = 200; p = 150
mu = rep(0,p)
Cov = matrix(0, p, p); for(j in 1:p) for(k in 1:p) Cov[j,k] = 0.5^{abs(j-k)}
beta = rep(0, p); beta[25:50] = 0.2
X = MASS::mvrnorm(n,mu,Cov)
y = X%%beta + rnorm(n)

## set G ##
test.set =c(40:60)

## Quadratic Functional ##
Est = QF(X, y, G = test.set, A = NULL, model = "linear", split=FALSE)
```

Continuing running two functions `ci()` and `summary()`:

```
ci(Est)
#> tau lower upper
#>1 0.25 0.8118792 1.466422
#>2 0.50 0.8046235 1.473677
#>3 1.00 0.7905648 1.487736
```

With the default $\tau = c(0.25, 0.5, 1)$, we obtain three different CIs for $\beta_G^T \Sigma_{G,G} \beta_G$. Note that the true value $\beta_G^T \Sigma_{G,G} \beta_G = 1.16$ belongs to all of these constructed CIs.

```
summary(Est)
#> Call:
#> Inference for Quadratic Functional
#>
#> tau est.plugin est.debias Std. Error z value Pr(>|z|)
#> 0.25 0.904 1.139 0.1670 6.822 8.969e-12 ***
#> 0.50 0.904 1.139 0.1707 6.674 2.486e-11 ***
#> 1.00 0.904 1.139 0.1779 6.405 1.504e-10 ***
```

Similarly to the LF() case, our proposed bias-corrected estimator is effective in correcting the bias of plugin estimator.

Conditional average treatment effect

The function CATE(), shorthanded for Conditional Average Treatment Effect, conducts inference for $\Delta(x_{\text{new}}) = f(x_{\text{new}}^T \beta^{(2)}) - f(x_{\text{new}}^T \beta^{(1)})$ under the high-dimensional regression model (2). This function can be implemented as follows:

```
CATE(X1,y1,X2,y2,loading.mat,model=c("linear","logistic","logistic_alter"),
intercept=TRUE,intercept.loading=FALSE,beta.init1=NULL,beta.init2=NULL,
lambda=NULL,mu=NULL,prob.filter=0.05,rescale=1.1,alpha=0.05,verbose=FALSE)
```

Here, X1 and y1 denote the design matrix and the response vector for the first sample of data respectively, while X2 and y2 denote those for the second sample of data. beta.init1 and beta.init2 are the initial estimator of the regression vector for the first and second samples. All other arguments are similarly defined as for the function LF().

For the fourth example, we consider the logistic regression case to illustrate CATE() with the argument model='logistic_alter'.

Example 4. In the first group of data, the covariates $X_i^{(1)}$ follows multivariate normal distribution with $\mu = 0_p$ and covariance $\Sigma = \mathbf{I}_p$; in the second group of data, the covariates $X_i^{(2)}$ follows multivariate normal distribution with $\mu = 0_p$ and covariance $\Sigma \in \mathbb{R}^{p \times p}$ with $\Sigma_{j,k} = 0.5^{|j-k|}$. We generate following the model $y_i^{(k)} \sim \text{Bernoulli}(f(X_i^{(k)T} \beta^{(k)}))$ with $f(z) = \exp(z)/(1 + \exp(z))$ for $k = 1, 2$. See the following code for details of $\beta^{(1)}, \beta^{(2)}$ and the further observation x_{new} .

```
## Data Preparation ##
set.seed(0)
n1 = 100; n2 = 180; p = 120
mu1 = mu2 = rep(0,p)
Cov1 = diag(p)
Cov2 = matrix(0, p, p); for(j in 1:p) for(k in 1:p) Cov2[j,k] = 0.5^{abs(j-k)}
beta1 = rep(0, p); beta1[c(1,2)] = c(0.5, 0.5)
beta2 = rep(0, p); beta2[c(1,2)] = c(1.8, 1.8)
X1 = MASS::mvrnorm(n1,mu1,Cov1); val1 = X1%*%beta1
X2 = MASS::mvrnorm(n2,mu2,Cov2); val2 = X2%*%beta2
y1 = rbinom(n1, 1, exp(val1)/(1+exp(val1)))
y2 = rbinom(n2, 1, exp(val2)/(1+exp(val2)))

## further observation ##
loading.mat = c(1, 1, rep(0, p-2))

## CATE ##
Est <- CATE(X1, y1, X2, y2,loading.mat, model="logistic_alter")
```

Having fitted the model, it allows for method ci() and summary() as LF() does.

```
ci(Est)
#> loading lower upper
#>1 1 1.614269 4.514703
```

The true value $x_{\text{new}}^T (\beta^{(2)} - \beta^{(1)}) = 2.6$ is included in the above 95% CI.

```
ci(Est, probability = TRUE)
#> loading      lower      upper
#>1           1 0.1531872 0.5086421
```

If we specify probability as TRUE, for the logistic regression, `ci()` yields the CI for $f(x_{\text{new}}^T \beta^{(2)}) - f(x_{\text{new}}^T \beta^{(1)})$ whose true value is 0.2423.

Inner Product

The function `InnProd()`, shorthand for Inner Product, conducts inference for $\beta_G^{(1)T} A \beta_G^{(2)}$ if A is the submatrix pre-specified or $\beta_G^{(1)T} \Sigma_{G,G} \beta_G^{(2)}$ under the high-dimensional regression models. Here, X_1 and y_1 denote the design matrix and the response vector for the first sample of data respectively, while X_2 and y_2 denote those for the second sample of data. All other arguments are similarly defined as for the function `QF()`.

```
InnProd(X1,y1,X2,y2,G,A = NULL,model=c("linear","logistic","logistic_alter"),intercept=TRUE,
        beta.init1=NULL,beta.init2=NULL,split = TRUE,lambda=NULL,mu=NULL,prob.filter=0.05,
        rescale=1.1,tau = c(0.25,0.5,1),alpha=0.05,verbose=FALSE)
```

In the following code, we demonstrate the use of `InnProd()` in linear regression.

Example 5. See the following code for generating two samples of data and inference for $\beta_G^{(1)T} A \beta_G^{(2)}$.

```
set.seed(0)
n1 = 200; n2 = 260; p = 120
mu1 = mu2 = rep(0,p)
Cov1 = diag(p)
Cov2 = matrix(0, p, p); for(j in 1:p) for(k in 1:p) Cov2[j,k] = 0.5^{abs(j-k)}
beta1 = rep(0, p); beta1[1:10] = 0.5
beta2 = rep(0, p); beta2[3:12] = 0.4
X1 <- MASS::mvrnorm(n1,mu1,Cov1)
X2 <- MASS::mvrnorm(n2,mu2,Cov2)
y1 <- X1*beta1 + rnorm(n1)
y2 <- X2*beta2 + rnorm(n2)

## Specify G and A ##
test.set = c(1:20)
A = diag(length(test.set))

## Inner Product ##
Est <- InnProd(X1, y1, X2, y2, G=test.set, A, model="linear")
```

Having fitted the model, it allows for method `ci()` and `summary()` as `QF()` does.

```
ci(Est)
#>   tau      lower      upper
#> 1 0.25 0.7432061 2.490451
#> 2 0.50 0.7128181 2.520839
#> 3 1.00 0.6520422 2.581615
```

The true value $\beta^{(1)T} A \beta^{(2)} = 1.6$ is included in the above CIs with all default τ values.

Distance

The function `Dist()`, shorthand for Distance, conducts inference for $\gamma_G^T A \gamma_G$, where $\gamma = \beta^{(1)} - \beta^{(0)}$, if A is the submatrix pre-specified or $\gamma_G^T \Sigma_{G,G} \gamma_G$ under the high-dimensional regression models. All arguments are similarly defined as for the function `InnProd()`.

```
Dist(X1,y1,X2,y2,G,A = NULL,model=c("linear","logistic","logistic_alter"),intercept=TRUE,
     beta.init1=NULL,beta.init2=NULL,split = TRUE,lambda=NULL,mu=NULL,prob.filter=0.05,
     rescale=1.1,tau = c(0.25,0.50,1),alpha=0.05,verbose=FALSE)
```

In Example 6 we illustrate the use of `Dist()` in linear regression.

Example 6. See the following code for generating two samples of data and inference for $\gamma_G^{(1)T} \Sigma_{G,G} \gamma_G^{(2)}$.

```
## Data Preparation ##
set.seed(0)
n1 = 220; n2 = 180; p = 100
mu = rep(0,p); Cov = diag(p)
beta1 = rep(0, p); beta1[1:2] = c(0.5, 1)
beta2 = rep(0, p); beta2[1:10] = c(0.3, 1.5, rep(0.08, 8))
X1 <- MASS::mvrnorm(n1,mu,Cov)
X2 <- MASS::mvrnorm(n2,mu,Cov)
y1 = X1%*%beta1 + rnorm(n1)
y2 = X2%*%beta2 + rnorm(n2)

## G ##
test.set = c(1:10)

## A is not specified $$$
Est <- Dist(X1, y1, X2, y2, G=test.set, A=NULL, model="linear", split=FALSE)
```

Having fitted the model, it allows for method `ci()` and `summary()` as `LF()` does.

```
ci(Est)
#>   tau   lower   upper
#>1 0.25 0.028202 0.6831165
#>2 0.50 0.000000 0.7196383
#>3 1.00 0.000000 0.7926819

summary(Est)
#> Call:
#> Inference for Distance
#>
#>   tau est.plugin est.debias Std. Error z value Pr(>|z|)
#> 0.25    0.4265    0.3557    0.1671  2.129 0.03327 *
#> 0.50    0.4265    0.3557    0.1857  1.915 0.05547 .
#> 1.00    0.4265    0.3557    0.2230  1.595 0.11070
```

The true value $\gamma_G^T \Sigma_{G,G} \gamma_G = 0.3412$. Similar to the previous instances, we note that the bias-corrected estimator effectively correct the bias of the plugin estimator. Depending on the τ values, we obtain various CIs, all of which encompass the true value. It is important to mention that in case of negative lower boundaries, they will be truncated at 0 for $\tau = 0.5$ and $\tau = 1$.

Applications

Motif Regression

We demonstrate the use of `LF()` function on a motif regression problem for predicting transcription factor binding sites (TFBS, also called ‘motifs’) in DNA sequences. The data set consists of a univariate response variable y measuring the binding intensity of the transcription factor on coarse DNA segments for $n = 2587$ genes. Moreover, for each of the n genes, a score describing the abundance of occurrence, is available for each of the $p = 666$ candidate motifs. This data set has been previously explored in [Yuan et al. \(2007\)](#). To summarize, we have the following data:

y_i :the binding intensity of the transcription factor on coarse DNA segment i
 $X_{i,j}$:the abundance score of candidate motif j in DNA segment i
 $i = 1, \dots, n; \quad j = 1, \dots, p$

Given the real data, we run the following code:

```
p = ncol(X)
loading.mat = diag(p)
## apply LF ##
Est = LF(X, y, loading.mat, model='linear')
## CI for each regression coef ##
ci(Est)
```

We apply the package function `LF()` and obtain 95% CIs for the 666 regression coefficients. The constructed CIs are illustrated in Figure 1. Among the 666 CIs, 25 are marked in red and lie completely

above 0, suggesting a positive relationship between the Motif and the binding intensity. Conversely, 23 of the CIs highlighted in blue lie completely below 0, indicating that the corresponding motif has a negative impact on the binding intensity of the transcription factor. In other words, many genes may be targeted by the transcription factors that bind to these 48 motifs.




Figure 1: Constructed CIs for the 666 regression coefficients.

Fasting Glucose Level Data

The aim is to analyze the effect of polymorphic genetic markers on the glucose level in a stock mice population using the `LF()` with the argument `model="logistic"`. The data set is available at <https://wp.cs.ucl.ac.uk/outbredmice/heterogeneous-stock-mice/>. Since fasting glucose level is an important indicator of type-2 diabetes, the fasting glucose level dichotomized at 11.1 (unit: mmol/L) is taken as the response variable. Specifically, glucose level below 11.1 is considered normal and above 11.1 high (pre-diabetic and diabetic). The covariates consist of 10,346 polymorphic genetic markers, and the sample size is 1,269. We include "gender" and "age" as baseline covariates. The polymorphic markers and baseline covariates are standardized before analysis. However, the number of polymorphic markers is large, and there exists a high correlation among some of them. To address this issue, we select a subset of polymorphic markers such that the maximum of absolute correlation among the markers is below 0.75. Eventually, we select a subset of 2,341 polymorphic markers. To sum up, we have the following data, for $i = 1, \dots, 1269$:

y_i : whether the fasting glucose level is above 11.1 mmol/L for unit i
 $X_{i,j}$: polymorphic marker j for unit i with $j = 1, 2, \dots, 2341$
 $X_{i,2342}$: gender of unit i
 $X_{i,2343}$: age of unit i

Given the real data, we run the following code:

```
p = ncol(X)
loading.mat = diag(p)[-c(2342, 2343)]
## apply LF ##
Est = LF(X, y, loading.mat, model='logistic')
## CI for each regression coef ##
ci(Est)
```

Once more, we utilize the package function `LF()` with `model = "logistic"` to generate CIs for the first 2341 regression coefficients (corresponding to all polymorphic markers). In Figure 2, we observe that 13 genes have CIs that lie entirely above 0 (highlighted in red), while 16 genes have CIs below 0 (highlighted in blue). This indicates their respective associations with the fasting glucose level.

Conclusion

There has been significant recent progress in debiasing inference methods for high-dimensional GLMs. This paper highlights the application of advanced debiasing techniques in high-dimensional GLMs




Figure 2: Constructed CIs for the 2341 regression coefficients.

using the R package **SIHR**. The package provides tools for estimating bias-corrected point estimators and constructing CIs for various low-dimensional objectives in both one- and two-sample regression settings. Through extensive simulations and real-data analyses, we demonstrate the practicality and versatility of the package across diverse fields of study, making it an essential addition to the literature.

Acknowledgement

Prabrisha Rakshit and Zhenyu Wang contributed equally to this work and are considered co-first authors. Dr. Tony Cai's research was supported in part by NSF grant DMS-2015259 and NIH grants R01-GM129781 and R01-GM123056. Dr. Zijian Guo's research was supported in part by NSF grants DMS-1811857 and DMS-2015373 and NIH grants R01-GM140463 and R01-LM013614. Dr. Zijian Guo is grateful to Dr. Lukas Meier for sharing the motif regression data used in this paper.

Bibliography

- A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso : pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011. [p1]
- P. J. Bickel, Y. Ritov, and A. B. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of statistics*, 37(4):1705–1732, 2009. [p2]
- P. Bühlmann and S. van de Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011. [p1]
- T. Cai, T. Tony Cai, and Z. Guo. Optimal statistical inference for individualized treatment effects in high-dimensional models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(4):669–719, 2021a. [p1, 2, 3, 4, 5]
- T. T. Cai and Z. Guo. Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *The Annals of Statistics*, 45(2):615–646, 2017. [p1]
- T. T. Cai and Z. Guo. Semisupervised inference for explained variance in high dimensional linear regression and its applications. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(2):391–419, 2020. [p1]
- T. T. Cai, Z. Guo, and R. Ma. Statistical inference for high-dimensional generalized linear models with binary outcomes. *Journal of the American Statistical Association*, pages 1–14, 2021b. [p1, 3, 4]
- J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2011. [p1]
- Z. Guo, W. Wang, T. T. Cai, and H. Li. Optimal estimation of genetic relatedness in high-dimensional linear models. *Journal of the American Statistical Association*, 114:358–369, 2019. [p1, 2, 4, 6]

- Z. Guo, P. Rakshit, D. S. Herman, and J. Chen. Inference for the case probability in high-dimensional logistic regression. *The Journal of Machine Learning Research*, 22(1):11480–11533, 2021a. [p4]
- Z. Guo, C. Renaux, P. Bühlmann, and T. Cai. Group inference in high dimensions with applications to hierarchical testing. *Electronic Journal of Statistics*, 15(2):6633–6676, 2021b. [p1, 3, 4, 5]
- Z. Guo, X. Li, L. Han, and T. Cai. Robust inference for federated meta-learning. *arXiv preprint arXiv:2301.00718*, 2023. [p2]
- J. Huang and C.-H. Zhang. Estimation and selection via absolute penalized convex minimization and its multistage adaptive applications. *Journal of Machine Learning Research*, 13(Jun):1839–1864, 2012. [p1]
- A. Javanmard and A. Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research*, 15(1):2869–2909, 2014. [p1, 2]
- R. Ma, Z. Guo, T. T. Cai, and H. Li. Statistical inference for genetic relatedness based on high-dimensional logistic regression. *arXiv preprint arXiv:2202.10007*, 2022. [p2, 6]
- N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006. [p2]
- N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics*, 37(1):246–270, 2009. [p1]
- S. Negahban, B. Yu, M. J. Wainwright, and P. K. Ravikumar. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009. [p1]
- T. Sun and C.-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012. [p1]
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 58(1):267–288, 1996. [p1]
- S. van de Geer, P. Bühlmann, Y. Ritov, and R. Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, 42:1166–1202, 2014. [p1, 2]
- M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE transactions on information theory*, 55(5):2183–2202, 2009. [p2]
- Y. Yuan, L. Guo, L. Shen, and J. S. Liu. Predicting gene expression from sequence: a reexamination. *PLoS computational biology*, 3(11):e243, 2007. [p13]
- C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*, 38(2):894–942, 2010. [p1]
- C.-H. Zhang and S. S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014. [p1, 2]
- P. Zhao and B. Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7: 2541–2563, 2006. [p2]

Prabrisha Rakshit
Rutgers, The State University of New Jersey
Address
USA
prabrisha.rakshit@rutgers.edu

Zhenyu Wang
Rutgers, The State University of New Jersey
Address
USA
zw425@stat.rutgers.edu

Tony Cai
University of Pennsylvania

Address
USA
tcai@wharton.upenn.edu

Zijian Guo
Rutgers, The State University of New Jersey
Address
USA
zijguo@stat.rutgers.edu