
More Powerful and General Selective Inference for Stepwise Feature Selection Using Homotopy Method

Kazuya Sugiyama^{*1} Vo Nguyen Le Duy^{*12} Ichiro Takeuchi¹²

Abstract

Conditional selective inference (SI) has been actively studied as a new statistical inference framework for data-driven hypotheses. The basic idea of conditional SI is to make inferences conditional on the selection event characterized by a set of linear and/or quadratic inequalities. Conditional SI has been mainly studied in the context of feature selection such as stepwise feature selection (SFS). The main limitation of the existing conditional SI methods is the loss of power due to over-conditioning, which is required for computational tractability. In this study, we develop a more powerful and general conditional SI method for SFS using the homotopy method which enables us to overcome this limitation. The homotopy-based SI is especially effective for more complicated feature selection algorithms. As an example, we develop a conditional SI method for forward-backward SFS with AIC-based stopping criteria, and show that it is not adversely affected by the increased complexity of the algorithm. We conduct several experiments to demonstrate the effectiveness and efficiency of the proposed method.

1. Introduction

As machine learning (ML) is being applied to a greater variety of practical problems, ensuring the reliability of ML is becoming increasingly important. Among several potential approaches to reliable ML, *conditional selective inference* (SI) is recognized as a promising approach for evaluating the statistical reliability of data-driven hypotheses selected by ML algorithms. The basic idea of conditional SI is to make inference on a data-driven hypothesis conditional on the selection event that the hypothesis is selected. Conditional SI has been actively studied especially in the context of

feature selection. Notably, Lee et al. (2016) and Tibshirani et al. (2016) proposed conditional SI methods for selected features using Lasso and stepwise feature selection (SFS), respectively. Their basic idea is to characterize the selection event by a polytope, i.e., a set of linear inequalities, in a sample space. When a selection event can be characterized by a polytope, practical computational methods developed by these authors can be used for making inferences of the selected hypotheses conditional on the selection events.

Unfortunately, however, such *polytope-based SI* has several limitations because it can only be used when the characterization of all relevant selection events is represented by a polytope. In fact, in most of the existing polytope-based SI studies, *extra-conditioning* is required in order for the selection event to be characterized as a polytope. For example, in SI for SFS by Tibshirani et al. (2016), the authors needed to consider conditioning not only on selected features but also on additional events regarding the *history* of the feature selection algorithmic process and the *signs* of the features. Such extra-conditioning leads to loss of power in the inference (Fithian et al., 2014).

In this study, we go beyond polytope-based SI and propose a novel conditional SI method using the *homotopy continuation* approach for SFS. We call the proposed method *homotopy-based SI*. In contrast to polytope-based SI for SFS in Tibshirani et al. (2016), the proposed method is more powerful and more general. The basic idea of homotopy-based SI is to use the homotopy continuation approach to keep track of the hypothesis selection event when the dataset changes in the direction of the selected test statistic, which enables efficient identification of the subspace of the sample space in which the same hypothesis is selected. Furthermore, we show that the proposed homotopy-based SI is more advantageous for more complicated feature selection algorithms. As an example, we develop a homotopy-based conditional SI for forward-backward SFS (FB-SFS) algorithm with AIC-based stopping criteria, and show that it is not adversely affected by the increased complexity of the algorithm and that a sufficiently high power is retained.

Related works Traditional statistical inference presumes that a statistical model and a statistical target for which we

^{*}Equal contribution ¹Nagoya Institute of Technology, Japan ²RIKEN, Japan. Correspondence to: Ichiro Takeuchi <takeuchi.ichiro@nitech.ac.jp>.

seek to conduct inference are determined before observing the dataset. Therefore, if we apply traditional methods to hypotheses selected after observing the dataset, the inferential results are no longer valid. This problem has been extensively discussed in the context of feature selection. In fact, even in commonly used feature selection methods such as SFS, correct assessment of the statistical reliability of selected features has long been difficult. Several approaches have been suggested in the literature toward addressing this problem (Benjamini & Yekutieli, 2005; Leeb & Pötscher, 2005; Leeb et al., 2006; Benjamini et al., 2009; Pötscher et al., 2010; Berk et al., 2013; Lockhart et al., 2014; Taylor et al., 2014).

In recent years, Lee et al. (2016) proposed a practical *conditional* SI framework to perform exact (non-asymptotic) inference for a set of features selected by lasso. In their work, the authors showed that, for a set of features selected by Lasso, the selection event can be characterized as a polytope by conditioning on the selected set of features as well as additional events on their signs. Furthermore, Tibshirani et al. (2016) showed that polytope-based SI is also applicable to SFS by additionally conditioning on the *history* and *signs* of sequential feature selection. Conditional SI has been actively studied in the past few years and extended to various directions (Fithian et al., 2015; Choi et al., 2017; Tian et al., 2018; Chen & Bien, 2019; Hyun et al., 2018; Loftus & Taylor, 2014, 2015; Panigrahi et al., 2016; Tibshirani et al., 2016; Yang et al., 2016; Suzumura et al., 2017; Tanizaki et al., 2020; Duy et al., 2020b; Tsukurimichi et al., 2021; Das et al., 2021).

In conditional SIs, it is typically preferred to condition on as little as possible so that the inference can be more powerful (Fithian et al., 2014). Namely, the main limitations of current polytope-based SI methods are that excessive conditioning is required to represent the selection event with a single polytope. In the seminal paper by Lee et al. (2016), the authors already discussed the problem of over-conditioning, and explained how extra conditioning on signs can be omitted by an exhaustive enumeration of all possible signs and by taking the union over the resulting polyhedra. However, such an exhaustive enumeration of exponentially increasing number of sign combinations is feasible only when the number of selected features is fairly small.

Several other approaches were proposed to circumvent the drawbacks and restrictions of polytope-based SI. Loftus & Taylor (2015) extended polytope-based SI such that selection events characterized by quadratic inequalities can be handled, but this inherits the same over-conditioning problem. To improve the power, Tian et al. (2018) proposed an approach to randomize the algorithm in order to condition on less. Terada & Shimodaira (2019) proposed to use bootstrap re-sampling to characterize the selection event

more generally. The disadvantage of these approaches is that additional randomness is introduced into the algorithm and/or the inference.

This study is motivated by Liu et al. (2018) and Duy & Takeuchi (2021). The former studied Lasso SI for full-model parameters, whereas the latter extended the basic idea of the former so that it can be also applied to Lasso SI in more general settings. These two studies go beyond polytope-based SI for more powerful Lasso SI without conditioning on signs. Because Lasso is formulated as a convex optimization problem, the selection event for conditional SI can be characterized based on the optimality conditions of the convex optimization problem. These two methods rely on the optimality conditions to characterize the minimum conditions for the Lasso feature selection event.

In contrast, the SFS algorithm cannot be formulated as an optimization problem. Therefore, existing conditional SI methods for SFS are based on conditioning on the *history* rather than the *output* of the SFS algorithm. Namely, existing SI methods actually consider conditions not only on which features are selected but also in what order they are selected. Moreover, when the SFS algorithm is extended to more complex procedures such as FB-SFS, the over-conditioning problem has more pronounced negative effects. Because features can be added and removed in various orders in FB-SFS, if we condition on the history rather than the output of the algorithm, we will suffer from extensive over-conditioning on the entire history of feature addition and removal.

Our contributions Our contributions are as follows:

- We propose a novel conditional SI method for the SFS algorithm using the homotopy continuation approach and show that the proposed method overcomes the over-conditioning problem in existing methods and enables us to conduct minimally conditioned more powerful conditional SI on the selected features by the SFS algorithm.
- We develop a conditional SI method for the FB-SFS algorithm with AIC-based stopping criteria by using the homotopy continuation approach. Although the conventional polytope-based SI method is more adversely affected by the over-conditioning problem when SFS is extended to FB-SFS, we show that our homotopy-based method still enables us to conduct minimally-conditioned SI and retains sufficiently high power.
- We demonstrate the validity and power of the proposed homotopy-based SI methods through numerical experiments on synthetic and real benchmark datasets.

For reproducibility, our implementation is available at https://github.com/takeuchi-lab/selective_inference_stepwise_feature_selection

Notations For a natural number n , we denote $[n]$ to be the set $\{1, \dots, n\}$. The $n \times n$ identity matrix is denoted by I_n . For a matrix X and a set of column indices \mathcal{M} , $X_{\mathcal{M}}$ indicates the matrix with the set of columns corresponding to \mathcal{M} .

2. Problem Statement

We consider the forward SFS for a regression problem. Let n be the number of instances and p be the number of original features. We denote the observed dataset as (X, \mathbf{y}) where $X \in \mathbb{R}^{n \times p}$ is the design matrix and $\mathbf{y} \in \mathbb{R}^n$ is the response vector. Following the problem setup from the existing literature on conditional SI such as Lee et al. (2016) and Tibshirani et al. (2016), we assume that the observed response \mathbf{y} is a realization of the following random response vector

$$\mathbf{Y} = (Y_1, \dots, Y_n)^\top \sim \mathbb{N}(\boldsymbol{\mu}, \Sigma), \quad (1)$$

where $\boldsymbol{\mu} \in \mathbb{R}^n$ is the unknown mean vector and $\Sigma \in \mathbb{R}^{n \times n}$ is the covariance matrix, which is known or estimable from independent data. The design matrix X is assumed to be non-random. For notational simplicity, we assume that each column vector of X is normalized to have unit length.

Stepwise feature selection We consider the standard forward SFS method as studied in Tibshirani et al. (2016) in §2 and §3. At each step of the SFS method, the feature that most improves the fit is newly added. When each feature has unit length, it is equivalent to the feature which is most correlated with the residual of the least-square regression model fitted with the previously selected features. For a response vector $\mathbf{y} \in \mathbb{R}^n$ and a set of features $\mathcal{M} \subseteq [p]$, let $\mathbf{r}(\mathbf{y}, X_{\mathcal{M}})$ be the residual vector obtained by regressing \mathbf{y} onto $X_{\mathcal{M}}$ for a set of features \mathcal{M} , i.e.,

$$\mathbf{r}(\mathbf{y}, X_{\mathcal{M}}) = P_{X_{\mathcal{M}}}^\perp \mathbf{y} = (I_n - P_{X_{\mathcal{M}}})\mathbf{y}$$

where $P_{X_{\mathcal{M}}} = X_{\mathcal{M}}(X_{\mathcal{M}}^\top X_{\mathcal{M}})^{-1}X_{\mathcal{M}}^\top$. Let K be the number of the selected features by the SFS method¹. We denote the index of the feature selected at step k as j_k and the set of selected features up to step k as $\mathcal{M}_k = \{j_1, \dots, j_k\}$ for $k \in [K]$ (Note, however, that, if there is no ambiguity, we simply denote the final set of the selected features as $\mathcal{M} (= \mathcal{M}_K)$). Feature j_k is selected as

$$\begin{aligned} j_k &= \arg \min_{j \in [p] \setminus \mathcal{M}_{k-1}} \|\mathbf{r}(\mathbf{y}, X_{\mathcal{M}_{k-1} \cup \{j\}})\|_2^2 \\ &= \arg \max_{j \in [p] \setminus \mathcal{M}_{k-1}} |\mathbf{x}_j^\top \mathbf{r}(\mathbf{y}, X_{\mathcal{M}_{k-1}})|, \end{aligned} \quad (2)$$

where $\mathcal{M}_0(\mathbf{y}) := \emptyset$ and $\mathbf{r}(\mathbf{y}, \mathcal{M}_0) := \mathbf{y}$.

¹We discuss a situation where the number of features K is selected by cross-validation in Appendix E. In other parts of the paper, we assume that K is determined before looking at the data.

Statistical inference In order to quantify the statistical significance of the relation between the selected features and the response, we consider a statistical test for each coefficient of the selected model parameters

$$\hat{\beta}_{\mathcal{M}} = (X_{\mathcal{M}}^\top X_{\mathcal{M}})^{-1} X_{\mathcal{M}}^\top \mathbf{y}.$$

Note that the j^{th} coefficient is written as $\hat{\beta}_{\mathcal{M},j} = \boldsymbol{\eta}^\top \mathbf{y}$ by defining

$$\boldsymbol{\eta} = X_{\mathcal{M}} (X_{\mathcal{M}}^\top X_{\mathcal{M}})^{-1} \mathbf{e}_j, \quad (3)$$

where $\mathbf{e}_j \in \mathbb{R}^{|\mathcal{M}|}$ is a unit vector, j^{th} element of which is 1 and 0 otherwise. Note that $\boldsymbol{\eta}$ depends on \mathcal{M} and j , but we omit the dependence for notational simplicity. We consider the following statistical test for each coefficient $\beta_{\mathcal{M},j} = \boldsymbol{\eta}^\top \boldsymbol{\mu}$

$$H_0 : \beta_{\mathcal{M},j} = 0 \quad \text{vs.} \quad H_1 : \beta_{\mathcal{M},j} \neq 0, \quad (4)$$

where $\beta_{\mathcal{M},j}$ is the j^{th} element of the population least squares $\boldsymbol{\beta}_{\mathcal{M}} = P_{\mathcal{M}} \boldsymbol{\mu}$, i.e., the projection of $\boldsymbol{\mu}$ onto the column space of $X_{\mathcal{M}}$. Throughout the paper, we do not assume that the linear model is correctly specified, i.e., it is possible that $\boldsymbol{\mu} \neq X_{\mathcal{M}} \boldsymbol{\beta}$ for any $\boldsymbol{\beta} \in \mathbb{R}^p$. Even when the linear model is not correctly specified, $\boldsymbol{\beta}_{\mathcal{M}}$ is still a well-defined best linear approximation.

Conditional Selective Inference Because the target of the inference is selected by observing the data \mathbf{y} , if we naively apply a traditional statistical test to the problem in (4) as if the inference target is pre-determined, the result will be invalid (type-I error cannot be controlled at the desired significance level) owing to *selection bias*. To address the selection bias problem, we consider *conditional SI* introduced in Lee et al. (2016) and Tibshirani et al. (2016). Let us write the SFS algorithm as the following function:

$$A : \mathbf{y} \mapsto \mathcal{M},$$

which maps a response vector $\mathbf{y} \in \mathbb{R}^n$ to the set of the selected feature \mathcal{M} by the SFS algorithm.

In conditional SI, the inference is conducted based on the following conditional sampling distribution of the test-statistic:

$$\boldsymbol{\eta}^\top \mathbf{Y} \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}), \mathcal{Q}(\mathbf{Y}) = \mathcal{Q}(\mathbf{y})\}, \quad (5)$$

where

$$\mathcal{Q}(\mathbf{Y}) = (I_n - \mathbf{c}\boldsymbol{\eta}^\top)\mathbf{Y} \quad \text{with} \quad \mathbf{c} = \Sigma\boldsymbol{\eta}(\boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta})^{-1}$$

is the nuisance parameter, which is independent of the test statistic. The first condition $\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y})$ in (5) indicates the event that the set of the selected features by the K -step SFS method with a random response vector \mathbf{Y} is \mathcal{M} , i.e., the same as those selected with the observed response vector

\mathbf{y} . The second condition $\mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y})$ indicates that the nuisance parameter for a random response vector \mathbf{Y} is the same as that for the observed vector \mathbf{y}^2 .

To conduct the conditional inference for (5), the main task is to identify the conditional data space

$$\mathcal{Y} = \{\mathbf{Y} \in \mathbb{R}^n \mid \mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}), \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y})\}. \quad (6)$$

Once \mathcal{Y} is identified, we can easily compute the pivotal quantity

$$F_{\boldsymbol{\eta}^\top \boldsymbol{\mu}, \boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta}}^{\mathcal{Z}}(\boldsymbol{\eta}^\top \mathbf{y}) \mid \mathbf{y} \in \mathcal{Y}, \quad (7)$$

where $F_{m, s^2}^{\mathcal{Z}}$ is the c.d.f. of the truncated Normal distribution with mean m , variance s^2 , and truncation region \mathcal{Z} . Later, we will explain how \mathcal{Z} is defined in (7). The pivotal quantity is crucial for calculating p -value or obtaining confidence interval. Based on the pivotal quantity, we can obtain *selective p-value* (Fithian et al., 2014) in the form of

$$p_j^{\text{selective}} = 2 \min\{\pi_j, 1 - \pi_j\}, \quad (8)$$

where $\pi_j = 1 - F_{0, \boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta}}^{\mathcal{Z}}(\boldsymbol{\eta}^\top \mathbf{y})$. Furthermore, to obtain $1 - \alpha$ confidence interval for any $\alpha \in [0, 1]$, by inverting the pivotal quantity in (7), we can find the smallest and largest values of $\boldsymbol{\eta}^\top \boldsymbol{\mu}$ such that the value of the pivotal quantity remains in the interval $[\frac{\alpha}{2}, 1 - \frac{\alpha}{2}]$ (Lee et al., 2016).

Characterization of the conditional data space \mathcal{Y} Using the second condition in (6), the data in \mathcal{Y} are restricted to a line (see Liu et al. (2018) and Fithian et al. (2014)). Therefore, the set \mathcal{Y} can be re-written, using a scalar parameter $z \in \mathbb{R}$, as

$$\mathcal{Y} = \{\mathbf{y}(z) \in \mathbb{R}^n \mid \mathbf{y}(z) = \mathbf{a} + \mathbf{b}z, z \in \mathcal{Z}\} \quad (9)$$

where $\mathbf{a} = \mathbf{q}(\mathbf{y})$, $\mathbf{b} = \Sigma \boldsymbol{\eta}(\boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta})^{-1}$, and

$$\mathcal{Z} = \{Z \in \mathbb{R} \mid \mathcal{A}(\mathbf{y}(Z)) = \mathcal{A}(\mathbf{y})\}. \quad (10)$$

Here, \mathbf{a} , \mathbf{b} and \mathcal{Z} depend on \mathcal{M} and j , but we omit the subscripts for notational simplicity. Now, let us consider a random variable $Z \in \mathbb{R}$ and its observation $z \in \mathbb{R}$ such that they respectively satisfy $\mathbf{Y} = \mathbf{a} + \mathbf{b}Z$ and $\mathbf{y} = \mathbf{a} + \mathbf{b}z$. The conditional inference in (5) is re-written as the problem of characterizing the sampling distribution of

$$Z \mid Z \in \mathcal{Z}. \quad (11)$$

Because $Z \sim \mathbb{N}(\boldsymbol{\eta}^\top \boldsymbol{\mu}, \boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta})$, $Z \mid Z \in \mathcal{Z}$ follows a truncated Normal distribution. Once the truncation region \mathcal{Z} is identified, the pivotal quantity in (7) is obtained as $F_{\boldsymbol{\eta}^\top \boldsymbol{\mu}, \boldsymbol{\eta}^\top \Sigma \boldsymbol{\eta}}^{\mathcal{Z}}(z)$. Thus, the remaining task is reduced to the characterization of \mathcal{Z} .

²The $\mathbf{q}(\mathbf{Y})$ corresponds to the component \mathbf{z} in the seminal paper (see Lee et al. (2016), §5, Eq. 5.2, and Theorem 5.2).

Over-conditioning in existing conditional SI methods

Unfortunately, full identification of the truncation region \mathcal{Z} in conditional SI for the SFS algorithm is considered computationally infeasible. Therefore, in existing conditional SI studies such as Tibshirani et al. (2016), the authors circumvent the computational difficulty by *over-conditioning*. Note that over-conditioning is not harmful for selective type-I error control, but it leads to the loss of power (Fithian et al., 2014). In fact, the decrease in the power due to over-conditioning is not unique problem for SFS in Tibshirani et al. (2016) but is a common major problem in many existing conditional SIs (Lee et al., 2016). In the next section, we propose a method to overcome this difficulty.

3. Proposed Homotopy-based SI for SFS

As we discussed in §2, to conduct conditional SI, the truncation region $\mathcal{Z} \subseteq \mathbb{R}$ in (10) should be identified. To construct \mathcal{Z} , our idea is 1) computing $\mathcal{A}(\mathbf{y}(z))$ for all $z \in \mathbb{R}$ and 2) identifying the set of intervals of $z \in \mathbb{R}$ on which $\mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y})$. However, it seems intractable to obtain $\mathcal{A}(\mathbf{y}(z))$ for infinitely many values of $z \in \mathbb{R}$.

To overcome the difficulty, we combine two approaches called *extra-conditioning* and *homotopy continuation*. Our idea is motivated by the regularization paths of Lasso (Osborne et al., 2000; Efron & Tibshirani, 2004), SVM (Hastie et al., 2004) and other similar methods (Rosset & Zhu, 2007; Bach et al., 2006; Rosset & Zhu, 2007; Tsuda, 2007; Lee & Scott, 2007; Takeuchi et al., 2009; Takeuchi & Sugiyama, 2011; Karasuyama & Takeuchi, 2010; Hocking et al., 2011; Karasuyama et al., 2012; Ogawa et al., 2013; Takeuchi et al., 2013; Duverle et al., 2013), in which the solution path along the regularization parameter can be computed by analyzing the KKT optimality conditions of parametrized convex optimization problems.

Although SFS cannot be formulated as a convex optimization problem, by introducing the notion of extra-conditioning, we note that a conceptually similar approach as homotopy continuation can be used to keep track all possible changes in the selected features when the response vector \mathbf{y} changes along the direction of the test-statistic. A conceptually similar idea has recently been used for evaluating the reliability of deep learning representations (Duy et al., 2020a).

3.1. Extra-Conditioning

First, let us consider the *history* \mathcal{H} and *signs* \mathcal{S} of the SFS algorithm. The *history* of the SFS algorithm is defined as

$$\mathcal{H} := (\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_K),$$

i.e., the sequence of the sets of the selected features in K steps. The *signs* of the SFS algorithm are defined as a vector

$\mathcal{S} := (\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K)$, k^{th} element of which is defined as

$$\mathcal{S}_k = \text{sgn}(\mathbf{x}_{j_k}^\top \mathbf{r}(\mathbf{y}, X_{\mathcal{M}_{k-1}}))$$

for $k \in [K]$, which indicates the sign of the j_k^{th} feature when it is first entered to the model at step k .

We interpret the function $\mathcal{A} : \mathbf{y} \mapsto M$ as a composite function $\mathcal{A} = \mathcal{A}_1 \circ \mathcal{A}_2$ where

$$\mathcal{A}_1 : (\mathcal{H}, \mathcal{S}) \mapsto \mathcal{M}, \text{ and } \mathcal{A}_2 : \mathbf{y} \mapsto (\mathcal{H}, \mathcal{S}),$$

i.e., the following relationships hold:

$$\begin{aligned} \mathcal{M} &= \mathcal{A}(\mathbf{y}) = \mathcal{A}_1(\mathcal{A}_2(\mathbf{y})), \\ \mathcal{M} &= \mathcal{A}_1((\mathcal{H}, \mathcal{S})), \\ (\mathcal{H}, \mathcal{S}) &= \mathcal{A}_2(\mathbf{y}). \end{aligned}$$

Let us now consider conditional SI not on $\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y})$ but on $\mathcal{A}_2(\mathbf{Y}) = \mathcal{A}_2(\mathbf{y})$. The next lemma indicates that, by conditioning on the history and the signs $(\mathcal{H}, \mathcal{S})$ rather than the set of the selected features \mathcal{M} , the truncation region can be simply represented as an interval in the line

$$\mathbf{y}(z) = \mathbf{a} + \mathbf{b}z, \quad z \in \mathbb{R}.$$

Lemma 1. Consider a response vector $\mathbf{y}' \in \mathcal{Y}$. Let $(\mathcal{H}', \mathcal{S}') = \mathcal{A}_2(\mathbf{y}')$ be the history and the signs obtained by applying the K -step SFS algorithm to the response vector \mathbf{y}' , and their elements are written as

$$\mathcal{H}' = (\mathcal{M}'_1, \dots, \mathcal{M}'_K) \text{ and } \mathcal{S}' = (\mathcal{S}'_1, \dots, \mathcal{S}'_K).$$

Then, the over-conditioned truncation region defined as

$$\mathcal{Z}^{\text{oc}}(\mathbf{y}') = \{z \in \mathbb{R} \mid \mathcal{A}_2(\mathbf{y}(z)) = \mathcal{A}_2(\mathbf{y}')\} \quad (12)$$

is an interval

$$z \in \left[\begin{array}{cc} \max_{\substack{k \in [K], \\ j \in [p] \setminus \mathcal{M}'_{k-1}, \\ d_{(k,j)} > 0}} \frac{e_{(k,j)}}{d_{(k,j)}}, & \min_{\substack{k \in [K], \\ j \in [p] \setminus \mathcal{M}'_{k-1}, \\ d_{(k,j)} < 0}} \frac{e_{(k,j)}}{d_{(k,j)}} \end{array} \right], \quad (13)$$

where

$$\begin{aligned} e_{(k,j)} &= (\mathbf{x}_j - \mathbf{x}_{j_k} \mathcal{S}'_k)^\top P_{X_{\mathcal{M}'_{k-1}}}^\perp \mathbf{a}, \\ d_{(k,j)} &= (\mathbf{x}_{j_k} \mathcal{S}'_k - \mathbf{x}_j)^\top P_{X_{\mathcal{M}'_{k-1}}}^\perp \mathbf{b}, \end{aligned}$$

for $k \in [K]$, $j \in [p] \setminus \mathcal{M}'_{k-1}$.

The proof is presented in Appendix A.

Note that the condition $\mathcal{A}_2(\mathbf{Y}) = \mathcal{A}_2(\mathbf{y})$ is redundant for our goal of making the inference conditional on $\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y})$. This over-conditioning is undesirable because it decreases the power of conditional inference (Fithian et al.,

Algorithm 1 SFS_conditional_SI

Input: $X, \mathbf{y}, K, [z_{\min}, z_{\max}]$

- 1: $\mathcal{M} \leftarrow$ Applying the K -step SFS algorithm to (X, \mathbf{y})
- 2: **for** each selected feature $j \in \mathcal{M}$ **do**
- 3: Compute $\boldsymbol{\eta} \leftarrow$ Equation (3)
- 4: Compute \mathbf{a} and $\mathbf{b} \leftarrow$ Equation (9)
- 5: $\mathcal{Z} \leftarrow$ compute_truncation_region($X, K, \mathbf{a}, \mathbf{b}, [z_{\min}, z_{\max}], \mathcal{M}$)
- 6: $p_j^{\text{selective}} \leftarrow$ Equation (8) (and/or selective confidence interval)
- 7: **end for**

Output: $\{p_j^{\text{selective}}\}_{j \in \mathcal{M}}$ (and/or selective confidence intervals)

2014), but the original conditional SI method for SFS in Tibshirani et al. (2016) performs conditional inference under exactly this over-conditioning case because otherwise the selection event cannot be formulated as a polytope. In the following subsection, we overcome this computational difficulty by utilizing the homotopy continuation approach.

3.2. Homotopy Continuation

In order to obtain the optimal truncation region \mathcal{Z} in (10), our idea is to enumerate all the over-conditioned regions in (12) and to consider their union

$$\mathcal{Z} = \bigcup_{\mathbf{y}' \in \mathcal{Y} \mid \mathcal{A}(\mathbf{y}') = \mathcal{A}(\mathbf{y})} \mathcal{Z}^{\text{oc}}(\mathbf{y}')$$

by homotopy continuation. To this end, let us consider all possible pairs of history and signs $(\mathcal{H}(z), \mathcal{S}(z)) = \mathcal{A}_2(\mathbf{y}(z))$ for $z \in \mathbb{R}$. Clearly, the number of such pairs is finite. With a slight abuse of notation, we index each pair by $t = 1, 2, \dots, T$ and denote it such as $(\mathcal{H}^{(t)}, \mathcal{S}^{(t)})$, where T is the number of the pairs. Lemma 1 indicates that each pair corresponds to an interval in the line. Without loss of generality, we assume that the left-most interval corresponds to $t = 1$, the second left-most interval corresponds to $t = 2$, and so on. Then, using an increasing sequence of real numbers $z_1 < z_2 < \dots < z_T < z_{T+1}$, we can write these intervals as $[z_1, z_2], [z_2, z_3], \dots, [z_T, z_{T+1}]$. In practice, we do not have to consider the entire line $z \in (-\infty, \infty)$, but it suffices to consider a sufficiently wide range. Specifically, in the experiments in §5, we set $z_{\min} = z_1 = -(|z| + 10)\sigma$ and $z_{\max} = z_{T+1} = (|z| + 10)\sigma$ where σ is the standard error of the test-statistic. With this choice of the range, the probability mass outside the range is negligibly small.

Our simple idea is to compute all possible pairs $\{(\mathcal{H}^{(t)}, \mathcal{S}^{(t)})\}_{t=1}^T$ by keeping track of the intervals $[z_1, z_2], [z_2, z_3], \dots, [z_T, z_{T+1}]$ one by one, and then to

Algorithm 2 `compute_truncation_region`

Input: $X, K, \mathbf{a}, \mathbf{b}, [z_{\min}, z_{\max}], \mathcal{M}$

- 1: Initialization: $t = 1, z_t = z_{\min}, \mathcal{Z} = \emptyset$
- 2: **while** $z_t < z_{\max}$ **do**
- 3: $\mathbf{y}(z_t + \Delta z) = \mathbf{a} + \mathbf{b}(z_t + \Delta z)$
- 4: $(\mathcal{H}^{(t)}, \mathcal{S}^{(t)}) \leftarrow$ Applying K -step SFS to $(X, \mathbf{y}(z_t + \Delta z))$
- 5: $z_{t+1} \leftarrow$ Equation (13)
- 6: **if** $\mathcal{M}^{(t)} = \mathcal{M}$ **then**
- 7: $\mathcal{Z} \leftarrow \mathcal{Z} \cup [z_t, z_{t+1}]$
- 8: **end if**
- 9: $t \leftarrow t + 1$
- 10: **end while**

Output: $\mathcal{Z} = \{z \in [z_{\min}, z_{\max}] \mid \mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y})\}$

compute the truncation region \mathcal{Z} by collecting the intervals in which the set of the selected features $\mathcal{M}^{(t)} = \mathcal{A}_1((\mathcal{H}^{(t)}, \mathcal{S}^{(t)}))$ is the same as the set of the actually selected features from the observed data $\mathcal{M} = \mathcal{A}(\mathbf{y})$, i.e.,

$$\mathcal{Z} = \bigcup_{t \in [T] \mid \mathcal{A}_1((\mathcal{H}^{(t)}, \mathcal{S}^{(t)})) = \mathcal{A}(\mathbf{y})} [z_t, z_{t+1}].$$

We call z_1, z_2, \dots, z_{T+1} *breakpoints*. We start from applying the SFS algorithm to the response vector $\mathbf{y}(z_1)$ and obtain the first pair $(\mathcal{H}^{(1)}, \mathcal{S}^{(1)})$. Next, using Lemma 1 with $\mathbf{y}' = \mathbf{y}(z_1)$, the next breakpoint z_2 is obtained as the right end of the interval in (13). Next, the second pair $(\mathcal{H}^{(2)}, \mathcal{S}^{(2)})$ is obtained by applying the SFS method to the response vector $\mathbf{y}(z_2 + \Delta z)$, where Δz is a small value such that $z_t + \Delta z < z_{t+1}$ for all $t \in [T]$. This process is repeated until the next breakpoint becomes greater than z_{\max} . The pseudo-code of the entire method is presented in Algorithm 1 and the computation of the truncation region \mathcal{Z} is described in Algorithm 2.

4. Forward Backward Stepwise Feature Selection

In this section, we present a conditional SI method for the *FB-SFS* algorithm using the homotopy method. At each step of the FB-SFS algorithm, a feature is considered for addition to or subtraction from the current set of the selected features based on some pre-specified criterion. As an example of commonly used criterion, we study the AIC-based criterion, which is used as the default option in well-known `stepAIC` package in R.

Under the assumption of Normal linear model $\mathbf{Y} \sim \mathcal{N}(X\boldsymbol{\beta}, \Sigma)$, for a set of the selected features \mathcal{M} , the AIC is

written as

$$\begin{aligned} \text{AIC} &:= \max_{\boldsymbol{\beta}} (\mathbf{y} - X_{\mathcal{M}}\boldsymbol{\beta})^\top \Sigma^{-1} (\mathbf{y} - X_{\mathcal{M}}\boldsymbol{\beta}) + 2|\mathcal{M}| \\ &= \mathbf{y}^\top A_{\mathcal{M}} \mathbf{y} + 2|\mathcal{M}|, \end{aligned} \quad (14)$$

where $A_{\mathcal{M}} = \Sigma^{-1} - \Sigma^{-1} X_{\mathcal{M}} (X_{\mathcal{M}}^\top \Sigma^{-1} X_{\mathcal{M}})^{-1} X_{\mathcal{M}}^\top \Sigma^{-1}$ and $|\mathcal{M}|$ is the number of the selected features (the irrelevant constant term is omitted here). The goal of AIC-based FB-SFS is to find a model with the smallest AIC while adding and removing features step by step. The algorithm starts either from the null model (the model with only a constant term) or the full model (the model with all the p features). At each step, among all possible choices of adding one feature to or deleting one feature from the current model, the one with the smallest AIC is selected. The algorithm terminates when the AIC is no longer reduced.

With a slight abuse of notations, let us use the following notations, the same as in §3.

$$\mathcal{A} : \mathbf{y} \mapsto \mathcal{M}, \mathcal{A}_1 : \mathcal{H} \mapsto \mathcal{M}, \mathcal{A}_2 : \mathbf{y} \mapsto \mathcal{H},$$

where \mathcal{M} is the set of the selected features and \mathcal{H} is the history of the FB-SFS algorithm written as

$$\mathcal{H} = (\mathcal{M}_1, \dots, \mathcal{M}_K),$$

where \mathcal{M}_k is the set of the selected features at step k for $k \in [K]$. Here, the history \mathcal{H} contains the information on feature addition and removal in the FB-SFS algorithm. Therefore, unlike the forward SFS algorithm in §3, the size of \mathcal{M}_k can be increased or decreased depending on whether a feature is added or removed.

Using these notations, the truncation region \mathcal{Z} is similarly obtained as

$$\begin{aligned} \mathcal{Z} &= \{z \in \mathbb{R} \mid \mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y})\} \\ &= \bigcup_{\mathcal{H} \mid \mathcal{A}_1(\mathcal{H}) = \mathcal{A}(\mathbf{y})} \{z \in \mathbb{R} \mid \mathcal{A}_2(\mathbf{y}(z)) = \mathcal{H}\} \end{aligned} \quad (15)$$

The basic idea of the homotopy method is the same as before; we find a range of $z \in \mathbb{R}$ such that the algorithm has a certain history \mathcal{H} and take the union of those for which the output \mathcal{M} obtained from the history \mathcal{H} is equal to the actual selected set of features $\mathcal{M} = \mathcal{A}(\mathbf{y})$. The following lemma indicates that a range of $z \in \mathbb{R}$ in which the algorithm has a certain history \mathcal{H} can be analytically obtained.

Lemma 2. *Consider a response vector $\mathbf{y}' \in \mathcal{Y}$. Then, the over-conditioned truncation region defined as*

$$\mathcal{Z}^{\text{oc}}(\mathbf{y}') = \{z \in \mathbb{R} \mid \mathcal{A}_2(\mathbf{y}(z)) = \mathcal{A}_2(\mathbf{y}')\} \quad (16)$$

is characterized by a finite number of quadratic inequalities of $z \in \mathbb{R}$.

The proof is presented in Appendix A.

Unlike the case of polytope-based SI for forward SFS in the previous section, the over-conditioned region $\mathcal{Z}^{oc}(\mathbf{y}')$ in Lemma 2 possibly consists of multiple intervals. The homotopy continuation approach can be similarly used for computing the union of intervals in (15). Starting from $z_1 = z_{\min}$, we first compute $\mathcal{Z}^{oc}(\mathbf{y}')$ with $\mathbf{y}' = \mathbf{y}(z_1)$ using Lemma 2. If $\mathcal{Z}^{oc}(\mathbf{y}')$ consists of multiple intervals, we can consider only the interval containing \mathbf{y}' and set the right end of that interval as the next breakpoint. The computation of the next breakpoint z_t can be written as

$$z_t = \max\{z \in \mathcal{Z}^{oc}(\mathbf{y}(z_{t-1} + \Delta z))\} + \Delta z, \quad (17)$$

for $t = 1, 2, \dots, T$, where, if $\mathcal{Z}^{oc}(\mathbf{y}(z_{t-1} + \Delta z))$ consists of two separated intervals, then the maximum operator in (17) shall take the maximum in the interval containing z_{t-1} .

5. Experiment

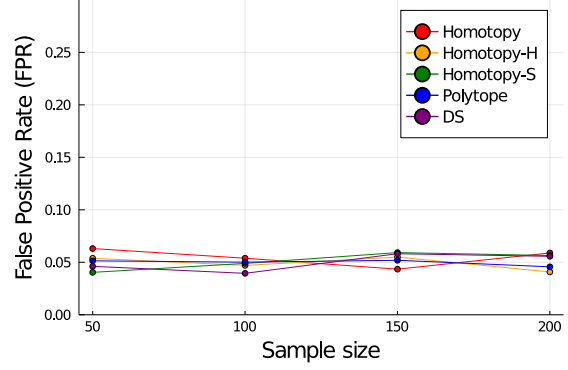
We present the experimental results of forward SFS in §3 and §4 in 5.1 and 5.2, respectively. The details of the experimental setups are described in Appendix B. More experimental results on the computational and robustness aspects of the proposed method are provided in Appendices C and D, respectively. In all the experiments, we set the significance level $\alpha = 0.05$.

5.1. Forward SFS

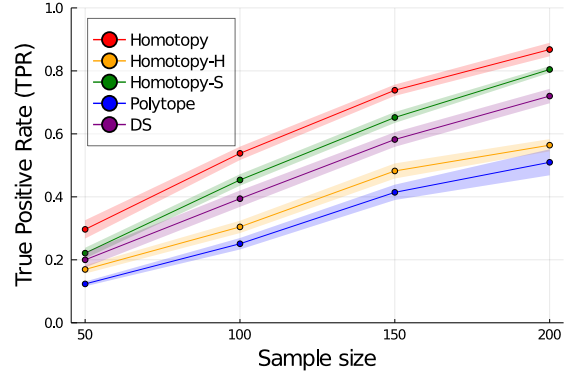
We compared the following five methods:

- **Homotopy**: conditioning on the selected features \mathcal{M} (minimal conditioning);
- **Homotopy-H**: additionally conditioning on the history \mathcal{H} ;
- **Homotopy-S**: additionally conditioning on the signs \mathcal{S} ;
- **Polytope** (Tibshirani et al., 2016): additionally conditioning on the history \mathcal{H} and signs \mathcal{S} ;
- **DS**: split the data and use one for feature selection and the other for inference.

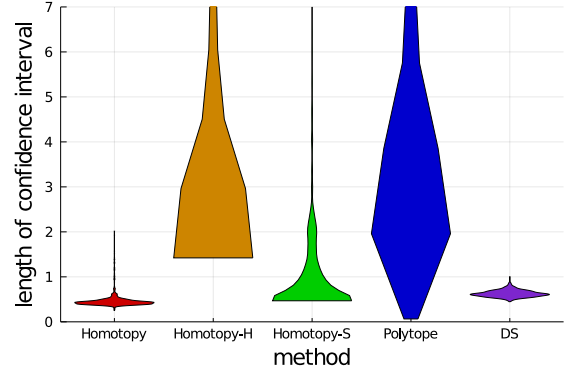
Synthetic data experiments We examined the false positive rates (FPRs), true positive rates (TPRs) and confidence interval (CI) lengths. We generated the dataset $\{(\mathbf{x}_i, y_i)\}_{i \in [n]}$ by $\mathbf{x}_i \sim \mathbb{N}(0, I_p)$ and $y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i$ with $\varepsilon_i \sim \mathbb{N}(0, 1)$. We set $n \in \{50, 100, 150\}$, $p = 5$, $K = 3$ for experiments on FPRs and TPRs. The coefficients $\boldsymbol{\beta}$ was respectively set as $[0, 0, 0, 0, 0]^\top$ and $[0.25, 0.25, 0, 0, 0]^\top$ for FPR and TPR experiments. We set $n = 100$, $p = 10$,



(a) False Positive Rates (FPRs)



(b) True Positive Rates (TPRs)



(c) Confidence Interval (CI) Lengths

Figure 1. Results of forward SFS on synthetic data.

$K = 9$ and $\boldsymbol{\beta} = [0.25, 0.25, 0.25, 0.25, 0.25, 0, 0, 0, 0, 0]^\top$ for experiments on CIs.

The results of FPRs, TPRs and CIs are shown in Fig. 1(a), (b) and (c), respectively. The FPRs and TPRs were estimated by 100 trials, and the plots in Fig. 1(a) and (b) are the averages when the 100 trials were repeated 20 times, while the plots in Fig. 1(c) are the results of 100 trials. In all five methods, the FPRs are properly controlled under the signifi-

cance level α . Regarding the TPR comparison, it is obvious that `Homotopy` method has the highest power because it is minimally conditioned. Note that the additional conditioning on the history \mathcal{H} and/or the signs \mathcal{S} significantly decreases the power. The results of the CIs are consistent with the results of TPRs.

Real data experiments We compared the proposed method (`Homotopy`) and conventional method (`Polytope`) on three real datasets³. From each of the original dataset, we randomly generated sub-sampled datasets with sizes 25, 50 and 100. We computed the lengths of CIs for each sub-sampled dataset by applying the proposed and the conventional method. Fig. 2 shows the results from 10000 trials. For all subsample sizes and datasets, the proposed methods provides shorter lengths of CIs than the conventional method.

5.2. Forward-Backward (FB)-SFS


We compared the following two methods:

- `Homotopy`: conditioning on the selected features \mathcal{M} (minimal conditioning);
- `Quadratic`: additionally conditioning on the history \mathcal{H} (implemented by using quadratic inequality-based conditional SI in (Loftus & Taylor, 2015)).

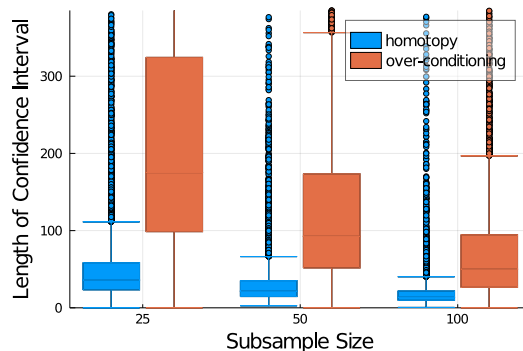
Synthetic data experiments We examined the FPRs, TPRs and CI lengths by generating synthetic datasets in the same way as above. We set $n \in \{50, 100, 150\}$ $p \in \{10, 20, 50\}$ for experiments on FPRs and TPRs. The coefficients β were zero for the case of FPR, whereas the first half of the coefficients were either 0.01, 0.25, 0.5, 1 and the second half were zero for the case of TPR. We set $n = 100$, $p = 10$, and each element of β was randomly set from $\mathbb{N}(0, 1)$ for experiments on CIs.

The results of FPRs, TPRs and CIs are shown in Fig. 3(a), (b) and (c), respectively. The FPRs and TPRs were estimated by 100 trials, and the plots in Fig. 3 (a) and (b) are the averages when the 100 trials were repeated 20 times for FPR and 10 times for TPR, whereas the plots in Fig. 3(c) are the results of 100 trials. In both methods, the FPRs are properly controlled under the significance level α . Regarding the TPR comparison, it is obvious that `Homotopy` method has the higher power than `Quadratic` because it is minimally conditioned. The results of the CIs are consistent with the results of TPRs.


³We used Housing (Dataset 1), Abalone (Dataset 2), and Concrete Compressive Strength (Dataset 3) datasets in UCI machine learning repository.



(a) Housing.




(b) Abalone.



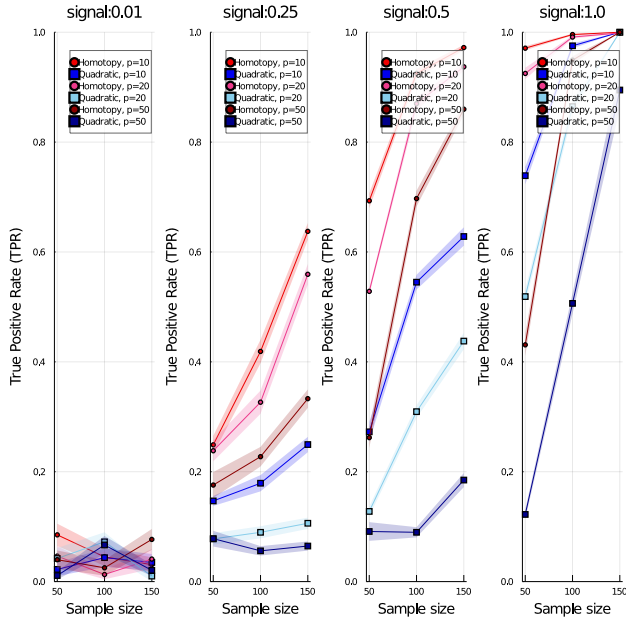
(c) Concrete Compressive Strength.

Figure 2. Results on three real data experiments for forward SFS.


Real data experiments We compared the proposed method (`Homotopy`) and conventional method (`Quadratic`) on the same three real datasets. From each of the original dataset, we randomly generated sub-sampled datasets with sizes 25, 50 and 100. We computed the lengths of CIs for each sub-sampled datasets by applying the proposed and the conventional methods. Fig. 4 shows the results from 10000 trials. For all subsample sizes and datasets, the proposed methods provides shorter lengths of CIs than the conventional method.



(a) False Positive Rates (FPRs)



(b) True Positive Rates (TPRs)

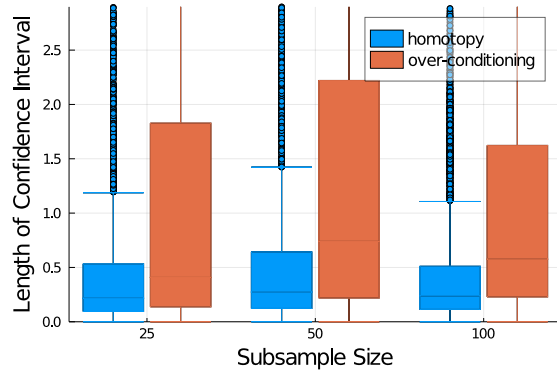


(c) Confidence Interval (CI) Lengths

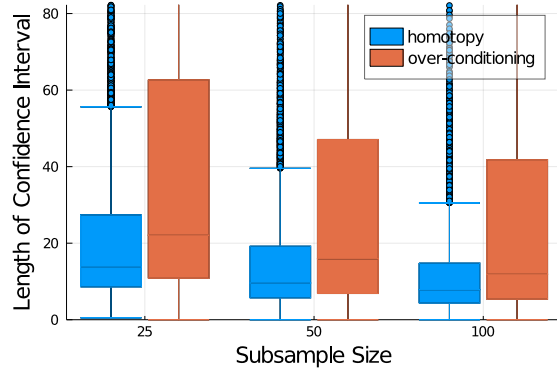
Figure 3. Results of FB-SFS on synthetic data.

6. Conclusion

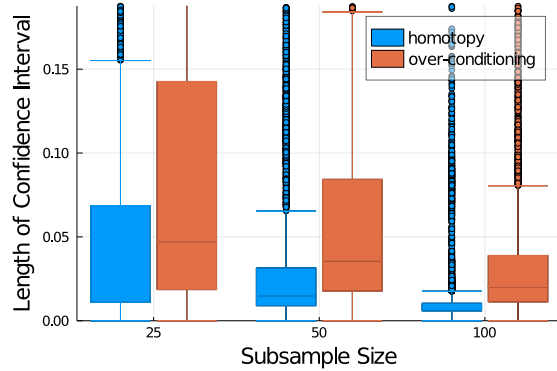
In this paper, we proposed a more powerful and general conditional SI method for SFS algorithm. We resolved the over-conditioning problem in existing approaches by introducing the homotopy continuation approach. The experimental results indicated that the proposed homotopy-based



(a) Housing.



(b) Abalone.



(c) Concrete Compressive Strength.

Figure 4. Results on three real data experiments for FB-SFS.

approach is more powerful and general.

Acknowledgements

This work was partially supported by MEXT KAKENHI (20H00601, 16H06538), JST Moonshot R&D (JPMJMS2033-05), NEDO (JPNP18002, JPNP20006), RIKEN Center for Advanced Intelligence Project, and RIKEN Junior Research Associate Program.

References

- Bach, F. R., Heckerman, D., and Horvits, E. Considering cost asymmetry in learning classifiers. *Journal of Machine Learning Research*, 7:1713–41, 2006.
- Benjamini, Y. and Yekutieli, D. False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469):71–81, 2005.
- Benjamini, Y., Heller, R., and Yekutieli, D. Selective inference in complex research. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4255–4271, 2009.
- Berk, R., Brown, L., Buja, A., Zhang, K., Zhao, L., et al. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Chen, S. and Bien, J. Valid inference corrected for outlier removal. *Journal of Computational and Graphical Statistics*, pp. 1–12, 2019.
- Choi, Y., Taylor, J., Tibshirani, R., et al. Selecting the number of principal components: Estimation of the true rank of a noisy matrix. *The Annals of Statistics*, 45(6): 2590–2617, 2017.
- Das, D., Duy, V. N. L., Hanada, H., Tsuda, K., and Takeuchi, I. Fast and more powerful selective inference for sparse high-order interaction model. *arXiv preprint arXiv:2106.04929*, 2021.
- Duverle, D. A., Takeuchi, I., Murakami-Tonami, Y., Kadamatsu, K., and Tsuda, K. Discovering combinatorial interactions in survival data. *Bioinformatics*, 29(23):3053–3059, 2013.
- Duy, V. N. L. and Takeuchi, I. Parametric programming approach for more powerful and general lasso selective inference. In *International Conference on Artificial Intelligence and Statistics*, pp. 901–909. PMLR, 2021.
- Duy, V. N. L., Iwazaki, S., and Takeuchi, I. Quantifying statistical significance of neural network representation-driven hypotheses by selective inference. *arXiv preprint arXiv:2010.01823*, 2020a.
- Duy, V. N. L., Toda, H., Sugiyama, R., and Takeuchi, I. Computing valid p-value for optimal changepoint by selective inference using dynamic programming. In *Advances in Neural Information Processing Systems*, volume 33, pp. 11356–11367, 2020b.
- Efron, B. and Tibshirani, R. Least angle regression. *Annals of Statistics*, 32(2):407–499, 2004.
- Fithian, W., Sun, D., and Taylor, J. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- Fithian, W., Taylor, J., Tibshirani, R., and Tibshirani, R. Selective sequential model selection. *arXiv preprint arXiv:1512.02565*, 2015.
- Hastie, T., Rosset, S., Tibshirani, R., and Zhu, J. The entire regularization path for the support vector machine. *Journal of Machine Learning Research*, 5:1391–415, 2004.
- Hocking, T., j. P. Vert, Bach, F., and Joulin, A. Clusterpath: an algorithm for clustering using convex fusion penalties. In *Proceedings of the 28th International Conference on Machine Learning*, pp. 745–752, 2011.
- Hyun, S., Lin, K., G’Sell, M., and Tibshirani, R. J. Post-selection inference for changepoint detection algorithms with application to copy number variation data. *arXiv preprint arXiv:1812.03644*, 2018.
- Karasuyama, M. and Takeuchi, I. Nonlinear regularization path for quadratic loss support vector machines. *IEEE Transactions on Neural Networks*, 22(10):1613–1625, 2010.
- Karasuyama, M., Harada, N., Sugiyama, M., and Takeuchi, I. Multi-parametric solution-path algorithm for instance-weighted support vector machines. *Machine Learning*, 88(3):297–330, 2012.
- Lee, G. and Scott, C. The one class support vector machine solution path. In *Proc. of ICASSP 2007*, pp. II521–II524, 2007.
- Lee, J. D., Sun, D. L., Sun, Y., Taylor, J. E., et al. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016.
- Leeb, H. and Pötscher, B. M. Model selection and inference: Facts and fiction. *Econometric Theory*, pp. 21–59, 2005.
- Leeb, H., Pötscher, B. M., et al. Can one estimate the conditional distribution of post-model-selection estimators? *The Annals of Statistics*, 34(5):2554–2591, 2006.
- Liu, K., Markovic, J., and Tibshirani, R. More powerful post-selection inference, with application to the lasso. *arXiv preprint arXiv:1801.09037*, 2018.
- Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. A significance test for the lasso. *Annals of statistics*, 42(2):413, 2014.
- Loftus, J. R. and Taylor, J. E. A significance test for forward stepwise model selection. *arXiv preprint arXiv:1405.3920*, 2014.

- Loftus, J. R. and Taylor, J. E. Selective inference in regression models with groups of variables. *arXiv preprint arXiv:1511.01478*, 2015.
- Ogawa, K., Imamura, M., Takeuchi, I., and Sugiyama, M. Infinitesimal annealing for training semi-supervised support vector machines. In *International Conference on Machine Learning*, pp. 897–905, 2013.
- Osborne, M. R., Presnell, B., and Turlach, B. A. A new approach to variable selection in least squares problems. *IMA journal of numerical analysis*, 20(3):389–403, 2000.
- Panigrahi, S., Taylor, J., and Weinstein, A. Bayesian post-selection inference in the linear model. *arXiv preprint arXiv:1605.08824*, 28, 2016.
- Pötscher, B. M., Schneider, U., et al. Confidence sets based on penalized maximum likelihood estimators in gaussian regression. *Electronic Journal of Statistics*, 4:334–360, 2010.
- Rosset, S. and Zhu, J. Piecewise linear regularized solution paths. *Annals of Statistics*, 35:1012–1030, 2007.
- Suzumura, S., Nakagawa, K., Umezu, Y., Tsuda, K., and Takeuchi, I. Selective inference for sparse high-order interaction models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3338–3347. JMLR. org, 2017.
- Takeuchi, I. and Sugiyama, M. Target neighbor consistent feature weighting for nearest neighbor classification. In *Advances in neural information processing systems*, pp. 576–584, 2011.
- Takeuchi, I., Nomura, K., and Kanamori, T. Nonparametric conditional density estimation using piecewise-linear solution path of kernel quantile regression. *Neural Computation*, 21(2):539–559, 2009.
- Takeuchi, I., Hongo, T., Sugiyama, M., and Nakajima, S. Parametric task learning. *Advances in Neural Information Processing Systems*, 26:1358–1366, 2013.
- Tanizaki, K., Hashimoto, N., Inatsu, Y., Hontani, H., and Takeuchi, I. Computing valid p-values for image segmentation by selective inference. 2020.
- Taylor, J., Lockhart, R., Tibshirani, R. J., and Tibshirani, R. Post-selection adaptive inference for least angle regression and the lasso. *arXiv preprint arXiv:1401.3889*, 354, 2014.
- Terada, Y. and Shimodaira, H. Selective inference after variable selection via multiscale bootstrap. *arXiv preprint arXiv:1905.10573*, 2019.
- Tian, X., Taylor, J., et al. Selective inference with a randomized response. *The Annals of Statistics*, 46(2):679–710, 2018.
- Tibshirani, R. J., Taylor, J., Lockhart, R., and Tibshirani, R. Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association*, 111(514):600–620, 2016.
- Tsuda, K. Entire regularization paths for graph data. In *In Proc. of ICML 2007*, pp. 919–925, 2007.
- Tsukurimichi, T., Inatsu, Y., Duy, V. N. L., and Takeuchi, I. Conditional selective inference for robust regression and outlier detection using piecewise-linear homotopy continuation. *arXiv preprint arXiv:2104.10840*, 2021.
- Yang, F., Barber, R. F., Jain, P., and Lafferty, J. Selective inference for group-sparse linear models. In *Advances in Neural Information Processing Systems*, pp. 2469–2477, 2016.