
Parametric Programming Approach for More Powerful and General Lasso Selective Inference

Vo Nguyen Le Duy

Nagoya Institute of Technology and RIKEN
duy.mllab.nit@gmail.com

Ichiro Takeuchi

Nagoya Institute of Technology and RIKEN
takeuchi.ichiro@nitech.ac.jp

Abstract

Selective Inference (SI) has been actively studied in the past few years for conducting inference on the features of linear models that are adaptively selected by feature selection methods such as Lasso. The basic idea of SI is to make inference conditional on the selection event. Unfortunately, the main limitation of the original SI approach for Lasso is that the inference is conducted not only conditional on the selected features but also on their signs — this leads to loss of power because of over-conditioning. Although this limitation can be circumvented by considering the union of such selection events for all possible combinations of signs, this is only feasible when the number of selected features is sufficiently small. To address this computational bottleneck, we propose a parametric programming-based method that can conduct SI without conditioning on signs even when we have thousands of active features. The main idea is to compute the continuum path of Lasso solutions in the direction of the selected test statistic, and identify the subset of the data space corresponding to the feature selection event by following the solution path. The proposed parametric programming-based method not only avoids the aforementioned computational bottleneck but also improves the performance and practicality of SI for Lasso in various respects. We conduct several experiments to demonstrate the effectiveness and efficiency of our proposed method.

1 Introduction

Reliable machine learning (ML), which is the problem of assessing the reliability of data-driven knowledge obtained by ML algorithms, is one of the most important issues in the ML community. Among various approaches for reliable ML, *selective inference (SI)* has been recognized as a new promising approach for assessing the statistical reliability of data-driven hypotheses selected by complex data analysis algorithms.

SI was first introduced as a statistical inference tool for the features selected by Lasso (Tibshirani, 1996). Although various properties of Lasso have been extensively studied in the past decades (see, e.g., Hastie et al. (2015)), exact statistical inference such as computing p -values or confidence intervals for *adaptively* selected features by Lasso has only recently begun to be actively studied in the context of SI (Lee et al. 2016; Fithian et al. 2014; Liu et al. 2018).

The main idea of SI is to make inference for the selected features *conditional on the selection event*, leading to exact *valid* inference on adaptively selected features by Lasso is possible in the sense that p -values for proper false positive rate control or confidence intervals with proper coverage guarantees can be obtained. After the seminal work (Lee et al. 2016), *conditional inference-based SI* has been actively studied and applied to various problems (Bachoc et al. 2014; Fithian et al. 2014, 2015; Choi et al. 2017; Tian and Taylor, 2018; Chen and Bien, 2019; Hyun et al. 2018; Bachoc et al. 2018; Charkhi and Claeskens, 2018; Loftus and Taylor, 2014; Loftus, 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,a; Sugiyama et al. 2020).

Existing works and their drawbacks. Let \mathcal{A} be a random variable indicating the set of the selected features by applying Lasso on any random data sample and \mathbf{s} be their signs. Then in the seminal work (Lee et al. 2016), the authors showed that the selection event $\{\mathcal{A} = \mathcal{A}_{\text{obs}}, \mathbf{s} = \mathbf{s}_{\text{obs}}\}$ is characterized as

a polytope in the data space, where \mathcal{A}_{obs} and \mathbf{s}_{obs} are the corresponding observations (see §2 for detailed setup), leading to the sampling distribution of the test-statistic in the form of a *truncated Normal distribution*. However, it is well-known that conditioning on the signs leads to low statistical power because of *over-conditioning*, which is widely recognized as a major drawback of the current Lasso SI approach.

Lee et al. (2016) also discussed the solution to overcome the drawback by conducting conditional inferences without sign event $\{\mathcal{A} = \mathcal{A}_{\text{obs}}\}$, which can be characterized by $2^{|\mathcal{A}_{\text{obs}}|}$ polytopes. If the number of selected features $|\mathcal{A}_{\text{obs}}|$ is moderate (e.g., up to 15), it is feasible to consider all affine constraints of all these $2^{|\mathcal{A}_{\text{obs}}|}$ polytopes. However, if $|\mathcal{A}_{\text{obs}}|$ is large, it is *infeasible* to enumerate the whole affine constraints for exponentially increasing number of polytopes.

Recently, Liu et al. (2018) have proposed two approaches, in which the problem settings are different from Lee et al. (2016), to improve the power. However, in their first approach, it is only applicable when the number of features p is smaller than the number of instances n . In the second approach, they still consider an exponentially large number of all possible sign vectors. This paper is motivated by Section 6 of Liu et al. (2018) in which they provide a recipe for constructing more powerful conditional SI methods. In the other direction, Tian and Taylor (2018) and Terada and Shimodaira (2019) proposed methods using randomization. A drawback of these randomization-based approaches including simple data-splitting approach is that further randomness is added in both feature selection and inference stages.

Many machine learning tasks involve careful tuning of a *regularization parameter* λ that controls the balance between an empirical loss term and a regularization term, e.g., commonly by cross-validation (CV). However, most of the current Lasso SI methods assume a pre-specified λ and ignore the fact that λ is selected based on the data because the selection event of cross-validation is difficult to characterize. Loftus (2015) and Markovic et al. (2017) proposed solutions to incorporate CV event. However, the former requires additional conditioning on all intermediate models which leads to loss of power and the latter considers a randomization version of CV instead of the vanilla CV.

Contribution. Our contributions are as follows:

- In this paper, we propose a new SI approach based on *parametric-programming (PP)* (Ritter, 1984; Allgower and George, 1993; Gal, 1995; Best, 1996), which we call *PP-based SI*, for resolving several major limitations of the seminal polytope-based SI proposed by

Lee et al. (2016). The polytope-based SI is applicable when the selection event can be characterized as a polytope in the data space. Otherwise, the only way is to consider extra conditions, e.g., sign conditioning, so that the over-conditioned event is characterized as a polytope, which leads to loss of statistical power. In contrast, with the proposed PP-based SI, it is possible to characterize the selection event even if they cannot be described by a polytope.

- We introduce a method to compute the continuum path of Lasso solutions in the direction of interest, which is sub-sequently used to identify the exact sampling distribution of the test statistic with the minimum amount of conditioning. Therefore, the PP-based SI can fundamentally resolve the over-conditioning problem, which was a major concern in polytope-based SI, to achieve the high statistical power. Although the concept of PP has been used in various problems (Osborne et al., 2000; Efron and Tibshirani, 2004; Hastie et al., 2004; Rosset, 2005; Bach et al., 2006; Rosset and Zhu, 2007; Tsuda, 2007; Lee and Scott, 2007; Garrigues and Ghaoui, 2008; Takeuchi et al., 2009; Karasuyama and Takeuchi, 2010; Hocking et al., 2011; Karasuyama et al., 2012; Lei, 2019), this is the first study that introduces a piecewise-linear PP approach for characterizing the selection events in SI.

- Furthermore, by using PP-based SI, we can perform SI *with minimal conditioning* for regularization parameter selection by cross-validation, which is complicated and was not possible with polytope-based SI. Besides, we show that our proposed method is general and can be applied in several settings as well as other selection models such as elastic net and interaction model.

Figure 1 shows the schematic illustration of the proposed method. For reproducibility, our implementation is available at

https://github.com/vonguyenleduy/parametric.lasso_selective.inference

2 Problem Statement

To formulate the problem, we consider a random response vector

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

where n is the number of instances, $\boldsymbol{\mu}$ is modeled as a linear function of p features $\mathbf{x}_1, \dots, \mathbf{x}_p \in \mathbb{R}^n$, and $\Sigma \in \mathbb{R}^{n \times n}$ is a covariance matrix which is known or estimable from independent data. The goal is to statistically quantify the significance of the relation between the features and response while properly controlling the false positive rate. To achieve the goal,

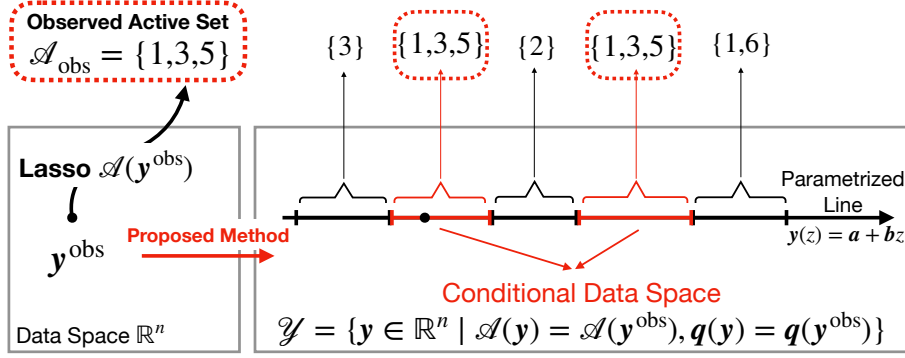


Figure 1: Schematic illustration of the proposed method. By applying Lasso on the observed data \mathbf{y}^{obs} , we obtain the observed active set \mathcal{A}_{obs} . The statistical inference for each selected feature is conducted conditional on the subspace \mathcal{Y} whose data has the same active set as \mathbf{y}^{obs} . We introduce a parametric programming method for characterizing the conditional data space \mathcal{Y} by searching on the parametrized line.

the authors in [Lee et al. \(2016\)](#) have proposed a practical SI framework, in which a subset of features is first “selected” by the Lasso, and the inferences are then conducted for each selected feature.

Feature selection and its selection event. Given an observed response vector $\mathbf{y}^{\text{obs}} \in \mathbb{R}^n$ sampled from the model [\(1\)](#), the Lasso optimization problem is given by

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{y}^{\text{obs}} - X\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1, \quad (2)$$

where $X \in \mathbb{R}^{n \times p}$ is a feature matrix, and $\lambda \geq 0$ is a regularization parameter. Since the Lasso produces sparse solutions, the active set selected by applying the Lasso to \mathbf{y}^{obs} is defined as

$$\mathcal{A}_{\text{obs}} = \mathcal{A}(\mathbf{y}^{\text{obs}}) = \{j : \hat{\beta}_j \neq 0\}. \quad (3)$$

Then, the event that the Lasso active set for a random vector \mathbf{Y} is the same as \mathbf{y}^{obs} is written as

$$\{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}})\}. \quad (4)$$

Statistical inference for the selected feature. The selected j^{th} coefficient is written as $\hat{\beta}_j = \boldsymbol{\eta}_j^\top \mathbf{y}^{\text{obs}}$ by defining

$$\boldsymbol{\eta}_j = X_{\mathcal{A}_{\text{obs}}} (X_{\mathcal{A}_{\text{obs}}}^\top X_{\mathcal{A}_{\text{obs}}})^{-1} \mathbf{e}_j, \quad (5)$$

where $\mathbf{e}_j \in \mathbb{R}^{|\mathcal{A}_{\text{obs}}|}$ is a basis vector with a 1 at position j^{th} . For the inference on the j^{th} selected feature, we consider the following statistical test

$$H_{0,j} : \boldsymbol{\eta}_j^\top \boldsymbol{\mu} = 0 \quad \text{vs.} \quad H_{1,j} : \boldsymbol{\eta}_j^\top \boldsymbol{\mu} \neq 0. \quad (6)$$

Since the hypothesis is generated from the data, selection bias exists. In order to correct the selection bias, we have to remove the information that has been

used for initial hypothesis generating process. This is achieved by considering the sampling distribution of the test statistic $\boldsymbol{\eta}_j^\top \mathbf{Y}$ conditional on the selection event, i.e.,

$$\boldsymbol{\eta}_j^\top \mathbf{Y} \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}, \quad (7)$$

where $\mathbf{q}(\mathbf{Y}) = (I_n - \mathbf{c}\boldsymbol{\eta}_j^\top) \mathbf{Y}$ with $\mathbf{c} = \Sigma \boldsymbol{\eta}_j (\boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j)^{-1}$. The second condition $\mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})$ indicates the component that is independent of the test statistic for a random vector \mathbf{Y} is the same as the one for \mathbf{y}^{obs} . The $\mathbf{q}(\mathbf{Y})$ corresponds to the component \mathbf{z} in the seminal paper (see [Lee et al. \(2016\)](#), Sec 5, Eq 5.2 and Theorem 5.2).

Once the selection event is identified, we can easily compute the pivotal quantity

$$F_{\boldsymbol{\eta}_j^\top \boldsymbol{\mu}, \boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j}^{\mathcal{Z}} (\boldsymbol{\eta}_j^\top \mathbf{Y}) \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}, \quad (8)$$

which is the c.d.f. of the truncated Normal distribution with mean $\boldsymbol{\eta}_j^\top \boldsymbol{\mu}$, variance $\boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j$, and the truncation region \mathcal{Z} which is calculated based on the selection event. The pivotal quantity is crucial for calculating p -value and confidence interval. Based on the pivotal quantity, we can consider *selective type I error* or *selective p -value* ([Fithian et al., 2014](#)) in the form of

$$P_j^{\text{selective}} = 2 \min\{\pi_j, 1 - \pi_j\}, \quad (9)$$

where $\pi_j = 1 - F_{0, \boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j}^{\mathcal{Z}} (\boldsymbol{\eta}_j^\top \mathbf{Y})$, which is *valid* in the sense that

$$\text{Prob}_{H_{0,j}} (P_j^{\text{selective}} < \alpha) = \alpha, \forall \alpha \in [0, 1].$$

Furthermore, to obtain $1-\alpha$ confidence interval for any $\alpha \in [0, 1]$, by inverting the pivotal quantity in Equation [\(8\)](#), we can find the smallest and largest values of

$\boldsymbol{\eta}_j^\top \boldsymbol{\mu}$ such that the value of pivotal quantity remains in the interval $[\frac{\alpha}{2}, 1 - \frac{\alpha}{2}]$ (Lee et al., 2016).

However, the main challenge is that characterizing $\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}})$ in Equation (7) is intractable because we have to consider $2^{|\mathcal{A}(\mathbf{y}^{\text{obs}})|}$ possible sign vectors. To overcome this issue, Lee et al. (2016) consider inference conditional not only on the selected features but also on their signs. Unfortunately, additionally considering the signs leads to low statistical power because of *over-conditioning*.

In the next section, we will introduce a method for identifying the minimum amount of conditioning $\{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}$, which leads to high statistical power. The main idea is to compute the path of Lasso solutions in the direction of interest $\boldsymbol{\eta}_j$. By focusing on the line along $\boldsymbol{\eta}_j$, we can skip majority of the polytopes that do not affect the truncated Normal sampling distribution because they do not intersect with this line. In other words, we can skip majority of combinations of signs that never appear when applying Lasso to the data on the line.

3 Proposed Method

In this section, we propose a parametric programming approach for characterizing conditioning event in (7). The schematic illustration is shown in Figure 1.

3.1 Conditional Data Space Characterization

Let us define the set of $\mathbf{y} \in \mathbb{R}^n$ which satisfies the conditions in Equation (7) as

$$\mathcal{Y} = \{\mathbf{y} \in \mathbb{R}^n \mid \mathcal{A}(\mathbf{y}) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \mathbf{q}(\mathbf{y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}. \quad (10)$$

According to the second condition, the data in \mathcal{Y} is restricted to a line (see Sec 6 in 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

$$\mathcal{Y} = \{\mathbf{y}(z) = \mathbf{a} + \mathbf{b}z \mid z \in \mathcal{Z}\}, \quad (11)$$

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

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

Now, let us consider a random variable $Z \in \mathbb{R}$ and its observation $z^{\text{obs}} \in \mathbb{R}$, which satisfy $\mathbf{Y} = \mathbf{a} + \mathbf{b}Z$ and $\mathbf{y}^{\text{obs}} = \mathbf{a} + \mathbf{b}z^{\text{obs}}$. The conditional inference in (7) is re-written as the problem of characterizing the sampling distribution of

$$Z \mid \{Z \in \mathcal{Z}\}. \quad (13)$$

Since $Z \sim \mathcal{N}(0, \boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j)$ under the null hypothesis, the law of $Z \mid \{Z \in \mathcal{Z}\}$ follows a truncated Normal

distribution. Once the truncation region \mathcal{Z} is identified, the pivotal quantity in Equation (8) is equal to $F_{0, \boldsymbol{\eta}_j^\top \Sigma \boldsymbol{\eta}_j}^{\mathcal{Z}}(Z)$, and can be easily obtained. Thus, the remaining task is to characterize \mathcal{Z} .

Characterization of truncation region \mathcal{Z} . Let us introduce the optimization problem (2) with parametrized response vector $\mathbf{y}(z)$ for $z \in \mathbb{R}$ as

$$\hat{\boldsymbol{\beta}}(z) = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{y}(z) - X\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1. \quad (14)$$

The subdifferential of the ℓ_1 -norm at $\hat{\boldsymbol{\beta}}(z)$ is defined as follows:

$$\partial \|\hat{\boldsymbol{\beta}}(z)\|_1 = \mathbf{s}(z) : \begin{cases} s_j(z) = \text{sign}(\hat{\beta}_j(z)) & \text{if } \hat{\beta}_j(z) \neq 0 \\ s_j(z) \in [-1, 1] & \text{if } \hat{\beta}_j(z) = 0 \end{cases},$$

where we denote $\mathbf{s}(z) = \text{sign}(\hat{\boldsymbol{\beta}}(z))$. Then, for any $z \in \mathbb{R}$, the optimality condition is given by

$$X^\top (X\hat{\boldsymbol{\beta}}(z) - \mathbf{y}(z)) + \lambda \mathbf{s}(z) = 0, \quad (15)$$

$\mathbf{s}(z) \in \partial \|\hat{\boldsymbol{\beta}}(z)\|_1$. To construct the truncation region \mathcal{Z} in Equation (12), we have to 1) compute the entire path of $\hat{\boldsymbol{\beta}}(z)$, and 2) identify the set of intervals of z on which $\mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y}^{\text{obs}})$. However, it seems intractable to compute $\hat{\boldsymbol{\beta}}(z)$ for infinitely many values of $z \in \mathbb{R}$. Our main idea to overcome this difficulty is to propose a parametric programming method for efficiently computing a finite number of ‘‘transition points’’ at which the active set changes.

3.2 A Piecewise Linear Homotopy

We now derive the main technique. We show that $\hat{\boldsymbol{\beta}}(z)$ is a piecewise linear function of z . To make the notation lighter, we write $\mathcal{A}_z = \mathcal{A}(\mathbf{y}(z))$, and we denote the set of inactive features as \mathcal{A}_z^c .

Lemma 1. Consider two real values z' and z ($z' > z$). Suppose $|s_j(z)| < 1$ for all $j \in \mathcal{A}_z^c$, $|s_j(z')| < 1$ for all $j \in \mathcal{A}_{z'}^c$, and $X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z}$ is invertible. If $\hat{\boldsymbol{\beta}}_{\mathcal{A}_z}(z)$ and $\hat{\boldsymbol{\beta}}_{\mathcal{A}_{z'}}(z')$ have the same active set and the same signs, then we have

$$\hat{\boldsymbol{\beta}}_{\mathcal{A}_z}(z') - \hat{\boldsymbol{\beta}}_{\mathcal{A}_z}(z) = \boldsymbol{\psi}_{\mathcal{A}_z}(z) \times (z' - z), \quad (16)$$

$$\lambda \mathbf{s}_{\mathcal{A}_z^c}(z') - \lambda \mathbf{s}_{\mathcal{A}_z^c}(z) = \boldsymbol{\gamma}_{\mathcal{A}_z^c}(z) \times (z' - z), \quad (17)$$

where $\boldsymbol{\psi}_{\mathcal{A}_z}(z) = (X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z})^{-1} X_{\mathcal{A}_z}^\top \mathbf{b}$, and $\boldsymbol{\gamma}_{\mathcal{A}_z^c}(z) = X_{\mathcal{A}_z^c}^\top \mathbf{b} - X_{\mathcal{A}_z^c}^\top X_{\mathcal{A}_z} \boldsymbol{\psi}_{\mathcal{A}_z}(z)$.

Proof. From the optimality conditions of the Lasso, we have

$$X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z} \hat{\boldsymbol{\beta}}_{\mathcal{A}_z}(z) - X_{\mathcal{A}_z}^\top \mathbf{y}(z) + \lambda \mathbf{s}_{\mathcal{A}_z}(z) = 0, \quad (18)$$

$$X_{\mathcal{A}_{z'}}^\top X_{\mathcal{A}_{z'}} \hat{\boldsymbol{\beta}}_{\mathcal{A}_{z'}}(z') - X_{\mathcal{A}_{z'}}^\top \mathbf{y}(z') + \lambda \mathbf{s}_{\mathcal{A}_{z'}}(z') = 0. \quad (19)$$

Then, by subtracting (18) from (19) and $\mathcal{A}_z = \mathcal{A}_{z'}$, we have

$$\begin{aligned}\hat{\beta}_{\mathcal{A}_z}(z') - \hat{\beta}_{\mathcal{A}_z}(z) &= (X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z})^{-1} X_{\mathcal{A}_z}^\top (\mathbf{y}(z') - \mathbf{y}(z)) \\ &= (X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z})^{-1} X_{\mathcal{A}_z}^\top (\mathbf{a} + \mathbf{b}z' - \mathbf{a} - \mathbf{b}z) \\ &= (X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z})^{-1} X_{\mathcal{A}_z}^\top \mathbf{b} \times (z' - z).\end{aligned}$$

Thus, we achieve Equation (16). Next, from the optimality conditions of the Lasso, we also have

$$-X_{\mathcal{A}_z^c}^\top X_{\mathcal{A}_z} \hat{\beta}_{\mathcal{A}_z}(z) + X_{\mathcal{A}_z^c}^\top \mathbf{y}(z) = \lambda \mathbf{s}_{\mathcal{A}_z^c}(z), \quad (20)$$

$$-X_{\mathcal{A}_z^c}^\top X_{\mathcal{A}_{z'}} \hat{\beta}_{\mathcal{A}_{z'}}(z') + X_{\mathcal{A}_z^c}^\top \mathbf{y}(z') = \lambda \mathbf{s}_{\mathcal{A}_z^c}(z'). \quad (21)$$

Similarly, by subtracting (20) from (21) and $\mathcal{A}_z = \mathcal{A}_{z'}$, we can easily achieve Equation (17). \square

Remark 1. In this paper, we assume the uniqueness of the Lasso solution $\hat{\beta}(z)$ for all $z \in \mathbb{R}$ as well as $|s_j(z)| < 1$ for all $j \in \mathcal{A}_z^c$ and the invertibility of $X_{\mathcal{A}_z}^\top X_{\mathcal{A}_z}$. These assumptions are justified by assuming the columns of X are in general position (Tibshirani, 2013). Parametric programming methods for handling the rare cases where these assumptions are not satisfied have been studied, e.g., in Best (1996), and can be applied to our problem setup. In practice, when the design matrix is not in general position, it is also common to introduce an additional ridge penalty term, resulting in the elastic net (Zou and Hastie, 2005). Our proposed method can be extended to the elastic net case (see Appendix 6.2.1 for the details).

Computation of the transition point. From Lemma 1, the solution $\hat{\beta}(z)$ is a linear function of z until z reaches a transition point at which either an element of $\hat{\beta}(z)$ becomes zero or a component of $\mathbf{s}(z)$ becomes one in absolute value. We now introduce how the transition point is identified.

Lemma 2. *Let z be a real value such that $\max_{j \in \mathcal{A}_z^c} |s_j(z)| < 1$. Then, $\mathcal{A}_{z'} = \mathcal{A}_z$, $\max_{j \in \mathcal{A}_{z'}^c} |s_j(z')| < 1$, and $\mathbf{s}(z) = \mathbf{s}(z')$ for any real value z' in the interval $[z, z + t_z)$, where $z + t_z$ is the value of transition point,*

$$t_z = \min \{t_z^1, t_z^2\}, \quad (22)$$

$$t_z^1 = \min_{j \in \mathcal{A}_z} \left(-\frac{\hat{\beta}_j(z)}{\psi_j(z)} \right)_{++}, \quad (23)$$

$$t_z^2 = \min_{j \in \mathcal{A}_z^c} \left(\lambda \frac{\text{sign}(\gamma_j(z)) - s_j(z)}{\gamma_j(z)} \right)_{++}. \quad (24)$$

Here, we use the convention that for any $m \in \mathbb{R}$, $(m)_{++} = m$ if $m > 0$, and $(m)_{++} = \infty$ otherwise.

Proof. From Equation (16), we can see that $\hat{\beta}_{\mathcal{A}_z}(z)$ is a function of z . For a real value z , there exists t_z^1 such

Algorithm 1 parametric_lasso_SI

Input: $X, \mathbf{y}^{\text{obs}}, \lambda, [z_{\min}, z_{\max}]$

- 1: Compute Lasso solution and obtain observed \mathcal{A}_{obs} for data $(X, \mathbf{y}^{\text{obs}})$
- 2: **for** each selected feature $j \in \mathcal{A}_{\text{obs}}$ **do**
- 3: Compute $\boldsymbol{\eta}_j \leftarrow$ Equation (5)
- 4: Compute \mathbf{a} and $\mathbf{b} \leftarrow$ Equation (11)
- 5: $\hat{\beta}(z), \mathcal{A}_z \leftarrow$ `compute_solution_path` ($X, \lambda, \mathbf{a}, \mathbf{b}, [z_{\min}, z_{\max}]$)
- 6: Truncation region $\mathcal{Z} \leftarrow \{z : \mathcal{A}_z = \mathcal{A}_{\text{obs}}\}$
- 7: $P_j^{\text{selective}} \leftarrow$ Equation (9) (and/or selective confidence interval of β_j)
- 8: **end for**

Output: $\{P_j^{\text{selective}}\}_{j \in \mathcal{A}_{\text{obs}}}$ (and/or selective confidence intervals of $\beta_j, j \in \mathcal{A}_{\text{obs}}$)

that for any real value z' in $[z, z + t_z^1)$, all elements of $\hat{\beta}_{\mathcal{A}_{z'}}(z')$ remain the same signs with $\hat{\beta}_{\mathcal{A}_z}(z)$. Similarly, from Equation (17), we can see that $\mathbf{s}_{\mathcal{A}_z^c}(z)$ is a function of z . Then, for a real value z , there exists t_z^2 such that for any real value z' in $[z, z + t_z^2)$, all elements of $\mathbf{s}_{\mathcal{A}_z^c}(z')$ are smaller than 1 in absolute value. Finally, by taking $t_z = \min\{t_z^1, t_z^2\}$, we obtain the interval in which the active set and signs of Lasso solution remain the same. The remaining task is how to compute t_z^1 and t_z^2 . We defer the detailed derivations of t_z^1 and t_z^2 to the Appendix 6.1 \square

3.3 Algorithm

In this section, we show the detailed algorithm of our proposed parametric programming method. In Algorithm 1, for feature selection step, we just simply apply Lasso to the data $(X, \mathbf{y}^{\text{obs}})$, and obtain the active set \mathcal{A}_{obs} . Then, we conduct SI for each selected feature. For testing $\beta_j, j \in \mathcal{A}_{\text{obs}}$, we first obtain the direction of interest $\boldsymbol{\eta}_j$, which can be easily computed as in Equation (5). Second, the main task is to compute the solution path of $\hat{\beta}(z)$ in Equation (14) for the parametrized response vector $\mathbf{y}(z)$, where, note that, the parametrized solution $\hat{\beta}(z)$ are different among different $j \in \mathcal{A}_{\text{obs}}$ since the direction of interest $\boldsymbol{\eta}_j$ depends on j . This task can be done by Algorithm 2. Finally, after having the path, we can easily obtain truncation region \mathcal{Z} which is used to compute selective p -value or selective confidence interval.

In Algorithm 2, a sequence of transition points are computed one by one. The algorithm is initialized at $z_k = z_{\min}, k = 0$. At each z_k , the task is to find the next transition point z_{k+1} , where the active set changes. This task can be done by computing the step size in Algorithm 3. This step is repeated until

Algorithm 2 `compute_solution_path`

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

- 1: Initialization: $k = 0, z_k = z_{\min}, \mathcal{T} = z_k$
- 2: **while** $z_k < z_{\max}$ **do**
- 3: $\mathbf{y}(z_k) = \mathbf{a} + \mathbf{b}z_k$
- 4: $t_{z_k}, \hat{\beta}(z_k), \mathcal{A}_{z_k} \leftarrow \text{compute_step_size}(X, \mathbf{y}(z_k), \lambda)$
- 5: $z_{k+1} = z_k + t_{z_k}, \mathcal{T} = \mathcal{T} \cup \{z_{k+1}\}$
 (z_{k+1} is the value of the next transition point)
- 6: $k = k + 1$
- 7: **end while**

Output: $\{\hat{\beta}(z_k)\}_{z_k \in \mathcal{T}}, \{\mathcal{A}_{z_k}\}_{z_k \in \mathcal{T}}$

Algorithm 3 `compute_step_size`

Input: $X, \mathbf{y}(z), \lambda$

- 1: Compute primal/dual Lasso solution $\hat{\beta}(z), \hat{\mathbf{s}}(z)$ for data $(X, \mathbf{y}(z))$
- 2: Obtain active set $\mathcal{A}_z = \{j : \hat{\beta}_j(z) \neq 0\}$
- 3: Compute $\psi_{\mathcal{A}_z}(z), \gamma_{\mathcal{A}_z^c}(z) \leftarrow$ Lemma 1
- 4: $t_z^1, t_z^2 \leftarrow$ Equations (23) and (24) in Lemma 2
- 5: $t_z = \min\{t_z^1, t_z^2\}$

Output: $t_z, \hat{\beta}(z), \mathcal{A}_z$

$z_k > z_{\max}$. The algorithm returns the sequences of Lasso solutions and transition points.

Choice of $[z_{\min}, z_{\max}]$. Under the normality, very positive and negative values of z does not affect the inference. Therefore, it is reasonable to consider range of values, e.g., $[-20\sigma, 20\sigma]$ (Liu et al., 2018), where σ is the standard deviation of the sampling distribution of test statistic.

3.4 Characterization of CV-based Tuning Parameter Selection Event

In this section, we introduce a new way to characterize the *minimal* selection event that λ is chosen based on the data, e.g., via cross-validation, which is complicated and thus none of the currently available Lasso SI methods can handle. Given a set of regularization parameter candidates Λ , we denote $\mathcal{V}(\mathbf{y}^{\text{obs}}) = \lambda^{\text{obs}} \in \Lambda$ is the event that λ^{obs} is selected when performing validation on \mathbf{y}^{obs} . The conditional inference on selected feature j when applying Lasso on $\{X, \mathbf{y}^{\text{obs}}\}$ is then defined as

$$\begin{aligned} \eta_j^\top \mathbf{Y} \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \\ \mathcal{V}(\mathbf{Y}) = \mathcal{V}(\mathbf{y}^{\text{obs}}), \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}. \end{aligned} \quad (25)$$

The conditional data space in (11) with validation selection event is re-defined as

$$\mathcal{Y} = \{\mathbf{y}(z) = \mathbf{a} + \mathbf{b}z \mid z \in \mathcal{Z}_{\text{CV}}\}, \quad (26)$$

where $\mathcal{Z}_{\text{CV}} = \{z \in \mathbb{R} \mid \mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y}^{\text{obs}}), \mathcal{V}(\mathbf{y}(z)) = \mathcal{V}(\mathbf{y}^{\text{obs}})\}$. We now can easily construct $\mathcal{Z}_1 = \{z \in \mathbb{R} \mid \mathcal{A}(\mathbf{y}(z)) = \mathcal{A}(\mathbf{y}^{\text{obs}})\}$ by using the proposed method in previous parts. The remaining task is to identify $\mathcal{Z}_2 = \{z \in \mathbb{R} \mid \mathcal{V}(\mathbf{y}(z)) = \mathcal{V}(\mathbf{y}^{\text{obs}})\}$. Finally, $\mathcal{Z}_{\text{CV}} = \mathcal{Z}_1 \cap \mathcal{Z}_2$.

For notational simplicity, we consider the case where the data is divided into training and validation sets, and the latter is used for selecting λ . The following discussion can be easily extended to cross-validation scenario. Let us re-write $\{X, \mathbf{y}^{\text{obs}}\} = \{(X_{\text{train}} \ X_{\text{val}})^\top \in \mathbb{R}^{n \times p}, (\mathbf{y}_{\text{train}}^{\text{obs}} \ \mathbf{y}_{\text{val}}^{\text{obs}})^\top \in \mathbb{R}^n\}$. For $\lambda \in \Lambda$, the Lasso problem on parametrized training response vector is written as

$$\hat{\beta}_\lambda(z) \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{y}_{\text{train}}(z) - X_{\text{train}}\beta\|_2^2 + \lambda \|\beta\|_1.$$

The validation error is defined as $E_\lambda(z) = \frac{1}{2} \|\mathbf{y}_{\text{val}}(z) - X_{\text{val}}\hat{\beta}_\lambda(z)\|_2^2$. Then, we can re-defined $\mathcal{Z}_2 = \{z \in \mathbb{R} \mid E_{\lambda^{\text{obs}}}(z) \leq E_\lambda(z) \text{ for any } \lambda \in \Lambda\}$. Since $\hat{\beta}_\lambda(z)$ is a piecewise-linear function of z and $\mathbf{y}_{\text{val}}(z)$ is a linear function of z , the validation error $E_\lambda(z)$ is a piecewise-quadratic function of z . Now, for each $\lambda \in \Lambda$, we have a corresponding piecewise-quadratic function of z . Finally, we can identify \mathcal{Z}_2 by finding the intervals of z in which the validation error $E_{\lambda^{\text{obs}}}(z)$ corresponding to λ^{obs} is minimum among a set of piecewise-quadratic functions.

3.5 The Generality of the Proposed Method

Since we can efficiently compute the path of Lasso solutions, our proposed method is flexible and can be easily extended to various respects. In Liu et al. (2018), the main limitations are their method can not be applied when $p > n$, or requires huge computation time. With our method, all these limitations are resolved. We provide detailed discussions and solutions in Appendices 6.2.2 and 6.2.3. Besides, we also apply the proposed method to other respects, which can not be solved by the methods in Lee et al. (2016) and Liu et al. (2018), including characterizing the minimum amount of conditioning in elastic net (Zou and Hastie, 2005) (Appendix 6.2.1), marginal model (Appendix 6.2.4), and interaction model (Appendix 6.2.5).

4 Experiment

In this section, we will demonstrate the performance of the proposed method. Here, we present the main results. Several additional experiments can be found in Appendix 6.3.

4.1 Experimental Setup

We executed the code on Intel(R) Xeon(R) CPU E5-2687W v4 @ 3.00GHz.

Methods for comparison. We show the false positive rates (FPRs), true positive rates (TPRs) and confidence intervals (CIs) for the following cases of conditional inferences:

- **TN-A:** conditional inference *without* sign conditioning, which is mainly focused in this paper,

$$\eta_j^\top \mathbf{Y} \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}_{\text{obs}}, \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\}.$$

- **TN-As:** conditional inference with additional sign conditioning, which is mainly focused in [Lee et al. \(2016\)](#),

$$\eta_j^\top \mathbf{Y} \mid \{\mathcal{A}(\mathbf{Y}) = \mathcal{A}_{\text{obs}}, \mathbf{s} = \mathbf{s}_{\text{obs}}, \mathbf{q}(\mathbf{Y}) = \mathbf{q}(\mathbf{y}^{\text{obs}})\},$$

where \mathbf{s} is the sign vector of Lasso solutions on \mathbf{Y} , and \mathbf{s}_{obs} is the sign vector of the Lasso solutions on \mathbf{y}^{obs} .

We also show the FPRs, TPRs and CIs of data splitting (**DS**) method ([Cox, 1975](#)), which is the commonly used procedure for the purpose of selection bias correction. In this approach, the data is randomly divided in two halves — one half is used for model selection and the other half is used for inference.

Synthetic data generation. We generated n outcomes as $y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i$, $i = 1, \dots, n$, where $\mathbf{x}_i \sim \mathcal{N}(0, I_p)$ in which $p = 5$, and $\varepsilon_i \sim \mathcal{N}(0, 1)$. Here, we assume that the variance of the noise is known. In practice, the variance can be estimated from *independent* data. We set the regularization parameter $\lambda = 1$ and significance level $\alpha = 0.05$. We used Bonferroni correction to account for the multiplicity in all the experiments. If we test m selected features (hypotheses) at the same time, then the Bonferroni correction would test each individual hypothesis at $\alpha^* = \alpha/m$. For the FPR experiments, all elements of $\boldsymbol{\beta}$ were set to 0 and we set $n \in \{100, 200, 300, 400, 500\}$. For the TPR experiments, the first two elements of $\boldsymbol{\beta}$ were set to 0.25. We ran 100 trials for each $n \in \{50, 100, 150, 200\}$, and we repeated this experiments 10 times. For the experiments of CIs, we set $n = 100, p = 10$, and the first 5 elements of $\boldsymbol{\beta}$ were set to 0.25.

Definition of TPR. In SI, we only conduct statistical testing when there is at least one hypothesis discovered by the algorithm. Therefore, the definition of TPR, which can be also called *conditional power*, is as follows:

$$\text{TPR} = \frac{\# \text{ correctly detected \& rejected}}{\# \text{ correctly detected}},$$

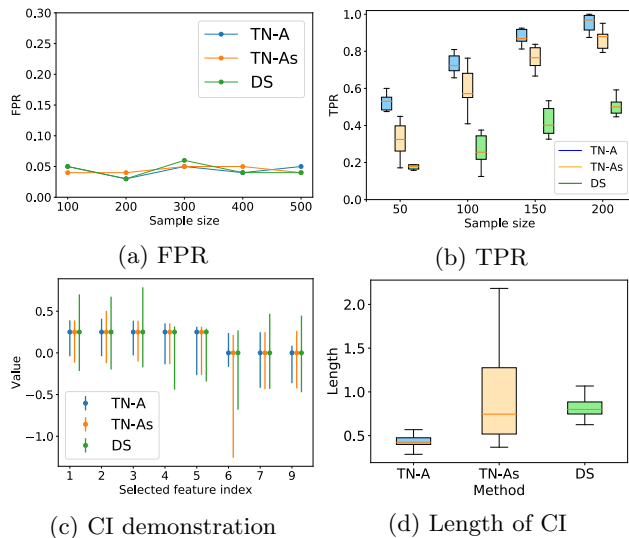


Figure 2: Results of false positive rate (FPR) control, true positive rate (TPR) and confidence interval (CI).

where $\#$ correctly detected is the number of truly positive features selected by the algorithm (e.g., Lasso) and $\#$ rejected is the number of truly positive features whose null hypothesis is rejected by SI.

4.2 Numerical Results

The results of FPRs, TPRs and CIs. The results of FPR and TPR are shown in Figures [2a](#) and [2b](#). In three cases, the FPRs are properly controlled under the significance level α . Regarding the TPR comparison, it is obvious that TN-A has the highest power. In regard to CI experiments, we note that the number of selected features between Lasso and DS can be different. Therefore, for a fair comparison, we only consider the features that are selected in both methods. In our experiments, since 9 features were selected by the Lasso in the cases of TN-A and TN-As while only 8 features were selected in the case of DS, we only show the 95% CI of the features that are selected in both cases in Figure [2c](#). The lengths of CI obtained by TN-A are almost the shortest. We repeated this experiment 100 times and showed the boxplot of the lengths of the confidence intervals in Figure [2d](#). In summary, the CI results are consistent with the TPR results, i.e., TN-A has the shortest length of CI which indicates it has the highest power.

The results when accounting CV selection event. We also demonstrate the TPRs and the lengths of CIs between the case when $\lambda = 2^0$ is fixed and λ is selected from the set $\Lambda_1 = \{2^{-1}, 2^0, 2^1\}$ or $\Lambda_2 = \{2^{-10}, 2^{-9}, \dots, 2^9, 2^{10}\}$. We show that the TPR tends to decrease when increasing the size of Λ as

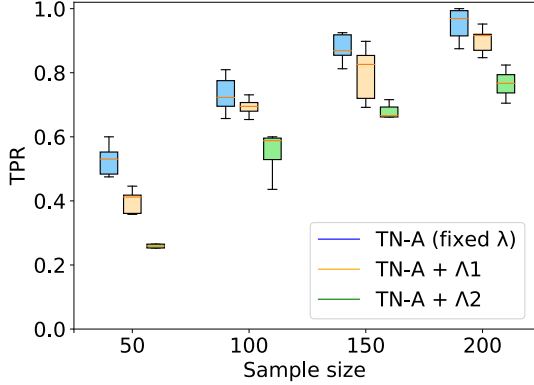


Figure 3: Demonstration of TPR when accounting cross-validation selection event.

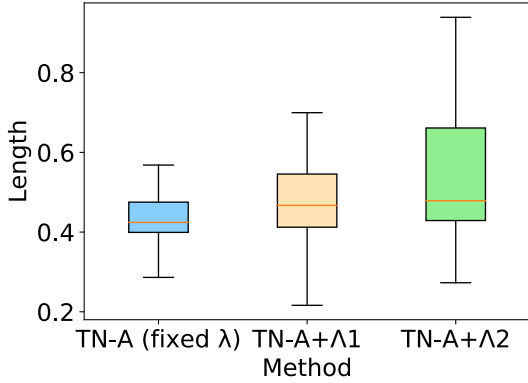


Figure 4: Demonstration of CI length when considering cross-validation selection event.

shown in Figure 3. This is due to the fact that when we increase the size of Λ , we have to condition on more information which leads to shorter truncation interval and results low TPR. The TPR results are consistent with the CI results shown in Figure 4 in which the length of CI is longer when increasing the size of Λ . Besides, we also conducted TPR comparison between our method and the *over-conditioning* version proposed in Loftus (2015). The results are shown in Figure 5. Our method has higher power since we can characterize minimum amount of conditioning.

The efficiency of the proposed method. In Lee et al. (2016), the authors mentioned the *naive* way to remove sign conditioning by enumerating all possible combination of signs $2^{|\mathcal{A}_{\text{obs}}|}$ which is only feasible when $|\mathcal{A}_{\text{obs}}|$ is small. On the left-hand side of Figure 6, we show the efficiency of our method compared to the naive way of removing sign conditioning. On the right-hand side of Figure 6, the Lasso SI without conditioning on signs can be done even when $n = 10,000$, $p = 10,000$ and thousands of features are selected while the naive way can not finish the task in realistic




Figure 5: TPR comparison with the existing method (Loftus, 2015) when accounting CV selection event.




Figure 6: Efficiency of the proposed method. With our method, Lasso SI without conditioning on signs can be done even when thousands of features are selected.

time. We also additionally show the efficiency of our method compare to two methods in Liu et al. (2018), which we call $\text{TN-}\ell_1$ and TN-Custom . The details of these two methods are shown in Appendix 6.2.3. In general, to perform these two methods, we still need to *naively* enumerate all possible combinations of signs. The results are shown in Figure 7.

One might wonder how we can circumvent the computational bottleneck of exponentially increasing number of polytopes. Our experience suggests that, by focusing on the the line along the test-statistic in data space, we can skip majority of the polytopes that do not affect the truncated Normal sampling distribution because they do not intersect with this line. In other words, we can skip majority of combinations of signs that never appear.

In Figure 8, we show the boxplot of the actual number of intervals of z that we encountered on the line when constructing the truncation region \mathcal{Z} . This indicates that the number of polytopes intersecting the line z that we need to consider is much smaller than $2^{|\mathcal{A}_{\text{obs}}|}$, which is considered in Lee et al. (2016)—this is the reason why the proposed approach can resolve all major limitations of the current SI method, making Lasso SI more powerful and practical.

We did not compare the computational time between the proposed method TN-A and the over-conditioning version TN-As because TN-As is obviously faster than TN-A but it has lower power than TN-A . Our main purpose is to demonstrate that the proposed method




Figure 7: Comparison between the proposed method and methods in Liu et al. (2018), in which an exponentially increasing number of all possible sign combinations are still required.

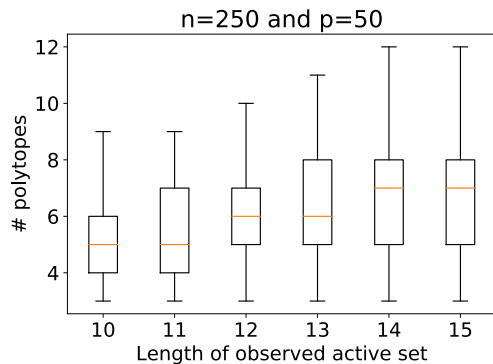


Figure 8: Number of encountered intervals on the line.

not only has high statistical power but also has practically computational costs.

We note that, in the worst-case, the complexity of the proposed method still grows exponentially. This is a common issue in other parametric programming applications such as regularization paths. However, fortunately, it has been well-recognized that this worst case rarely happens in practice, and our experiments suggest that this also applies to PP-based SI.

The robustness of the proposed method in terms of the FPR control. We demonstrate the robustness of our method in terms of the FPR control by considering the following cases:

- **Non-normal noise:** we consider the noise following Laplace distribution, skew normal distribution (skewness coefficient 10), and t_{20} distribution.
- **Unknown σ^2 :** we also consider the case when the variance is estimated from the data.

We generated n outcomes as $y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i$, $i = 1, \dots, n$, where $p = 5$, $\mathbf{x}_i \sim \mathcal{N}(0, I_p)$, and ε_i follows Laplace distribution, skew normal distribution, or t_{20} distribution with zero mean and standard deviation was set to 1. In the case of estimated σ^2 , $\varepsilon_i \sim \mathcal{N}(0, 1)$. We set all elements of $\boldsymbol{\beta}$ to 0, and set $\lambda = 0.5$. For each case, we

Table 1: Results on high-dimensional real-world bioinformatics related datasets.

	n	p	$ \mathcal{A}_{\text{obs}} $	Avg. Time (s)
Dataset 1	89	5787	600	0.374
Dataset 2	76	5144	621	0.344
Dataset 3	133	5787	660	0.342

ran 1,200 trials for each $n \in \{100, 200, 300, 400\}$. We confirmed that our method still maintains good performance on FPR control. The results are shown in Appendix 6.3

4.3 Results on Real-World Datasets

We demonstrate the efficiency of the proposed method by applying it on high-dimensional real-world bioinformatics related datasets, which is available at http://www.coepra.org/CoEPRA_regr.html. In datasets 1 and 3, n is the number of nona-peptides. Each amino acid in a nona-peptide is described by 643 descriptors, for a total of $p = 643 \times 9 = 5787$ descriptors. In dataset 2, n is the number of octa-peptides. Each amino acid in a octa-peptide is described by 643 descriptors, for a total of $p = 643 \times 8 = 5144$ descriptors. For these experiments, we used elastic net instead of Lasso to obtain large \mathcal{A}_{obs} . The extension of the proposed method for elastic net is presented in Appendix 6.2.1. The results are shown in Table 1. The time shown in the table is the average time to compute p -value for a selected feature.

5 Conclusion

In this paper, we have proposed a general method for characterizing the selection event of Lasso SI by introducing piecewise-linear parametric programming approach. With the proposed method, we can conduct a powerful SI by conditioning only on the selected features without the need of enumerating all possible sign vectors. Besides, we also introduced a new way to characterizing the cross-validation based tuning parameter selection. The proposed method not only overcomes the drawbacks of current Lasso SI methods but also improves the performance and practicality of SI for Lasso in various respects. Our idea is general and can be applied to circumvent several drawbacks of all the methods that are based on the current SI framework. We conducted experiments on both synthetic and real-world datasets to demonstrate the effectiveness and efficiency of our proposed method.

Acknowledgements

This work was partially supported by MEXT KAKENHI (20H00601, 16H06538), JST CREST (JP-MJCR1502), RIKEN Center for Advanced Intelligence Project, and RIKEN Junior Research Associate Program.

References

- E. L. Allgower and K. George. Continuation and path following. *Acta Numerica*, 2:1–63, 1993.
- F. R. Bach, D. Heckerman, and E. Horvits. Considering cost asymmetry in learning classifiers. *Journal of Machine Learning Research*, 7:1713–41, 2006.
- F. Bachoc, H. Leeb, and B. M. Pötscher. Valid confidence intervals for post-model-selection predictors. *arXiv preprint arXiv:1412.4605*, 2014.
- F. Bachoc, G. Blanchard, and P. Neuvial. On the post selection inference constant under restricted isometry properties. *Electronic Journal of Statistics*, 12(2):3736–3757, 2018.
- M. J. Best. An algorithm for the solution of the parametric quadratic programming problem. *Applied Mathematics and Parallel Computing*, pages 57–76, 1996.
- A. Charkhi and G. Claeskens. Asymptotic post-selection inference for the akaike information criterion. *Biometrika*, 105(3):645–664, 2018.
- S. Chen and J. Bien. Valid inference corrected for outlier removal. *Journal of Computational and Graphical Statistics*, pages 1–12, 2019.
- Y. Choi, J. Taylor, and R. Tibshirani. Selecting the number of principal components: Estimation of the true rank of a noisy matrix. *The Annals of Statistics*, 45(6):2590–2617, 2017.
- D. R. Cox. A note on data-splitting for the evaluation of significance levels. *Biometrika*, 62(2):441–444, 1975.
- V. N. L. Duy, S. Iwazaki, and I. Takeuchi. Quantifying statistical significance of neural network representation-driven hypotheses by selective inference. *arXiv preprint arXiv:2010.01823*, 2020a.
- V. N. L. Duy, H. Toda, R. Sugiyama, and I. Takeuchi. Computing valid p-value for optimal changepoint by selective inference using dynamic programming. *arXiv preprint arXiv:2002.09132*, 2020b.
- B. Efron and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32(2):407–499, 2004.
- W. Fithian, D. Sun, and J. Taylor. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- W. Fithian, J. Taylor, R. Tibshirani, and R. Tibshirani. Selective sequential model selection. *arXiv preprint arXiv:1512.02565*, 2015.
- T. Gal. *Postoptimal Analysis, Parametric Programming, and Related Topics*. Walter de Gruyter, 1995.
- P. Garrigues and L. Ghaoui. An homotopy algorithm for the lasso with online observations. *Advances in neural information processing systems*, 21:489–496, 2008.
- T. Hastie, S. Rosset, R. Tibshirani, and J. Zhu. The entire regularization path for the support vector machine. *Journal of Machine Learning Research*, 5: 1391–415, 2004.
- T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC press, 2015.
- T. Hocking, j. P. Vert, F. Bach, and A. Joulin. Clusterpath: an algorithm for clustering using convex fusion penalties. In *Proceedings of the 28th International Conference on Machine Learning*, pages 745–752, 2011.
- S. Hyun, K. Lin, M. G’Sell, and R. J. Tibshirani. Post-selection inference for changepoint detection algorithms with application to copy number variation data. *arXiv preprint arXiv:1812.03644*, 2018.
- M. Karasuyama and I. Takeuchi. Nonlinear regularization path for quadratic loss support vector machines. *IEEE Transactions on Neural Networks*, 22(10):1613–1625, 2010.
- M. Karasuyama, N. Harada, M. Sugiyama, and I. Takeuchi. Multi-parametric solution-path algorithm for instance-weighted support vector machines. *Machine Learning*, 88(3):297–330, 2012.
- G. Lee and C. Scott. The one class support vector machine solution path. In *Proc. of ICASSP 2007*, pages II521–II524, 2007.
- J. D. Lee and J. E. Taylor. Exact post model selection inference for marginal screening. In *Advances in neural information processing systems*, pages 136–144, 2014.
- J. D. Lee, D. L. Sun, Y. Sun, and J. E. Taylor. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016.
- J. Lei. Fast exact conformalization of the lasso using piecewise linear homotopy. *Biometrika*, 106(4):749–764, 2019.
- K. Liu, J. Markovic, and R. Tibshirani. More powerful post-selection inference, with application to the lasso. *arXiv preprint arXiv:1801.09037*, 2018.
- J. R. Loftus. Selective inference after cross-validation. *arXiv preprint arXiv:1511.08866*, 2015.

- J. R. Loftus and J. E. Taylor. A significance test for forward stepwise model selection. *arXiv preprint arXiv:1405.3920*, 2014.
- J. Markovic, L. Xia, and J. Taylor. Unifying approach to selective inference with applications to cross-validation. *arXiv preprint arXiv:1703.06559*, 2017.
- M. R. Osborne, B. Presnell, and B. A. Turlach. A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis*, 20(20):389–404, 2000.
- S. Panigrahi, J. Taylor, and A. Weinstein. Bayesian post-selection inference in the linear model. *arXiv preprint arXiv:1605.08824*, 28, 2016.
- K. Ritter. On parametric linear and quadratic programming problems. *mathematical Programming: Proceedings of the International Congress on Mathematical Programming*, pages 307–335, 1984.
- S. Rosset. Following curved regularized optimization solution paths. In *Advances in Neural Information Processing Systems 17*, pages 1153–1160, 2005.
- S. Rosset and J. Zhu. Piecewise linear regularized solution paths. *Annals of Statistics*, 35:1012–1030, 2007.
- K. Sugiyama, V. N. L. Duy, and I. Takeuchi. More powerful and general selective inference for stepwise feature selection using the homotopy continuation approach. *arXiv preprint arXiv:2012.13545*, 2020.
- S. Suzumura, K. Nakagawa, Y. Umezu, K. Tsuda, and I. Takeuchi. Selective inference for sparse high-order interaction models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3338–3347. JMLR. org, 2017.
- I. Takeuchi, K. Nomura, and T. Kanamori. Non-parametric conditional density estimation using piecewise-linear solution path of kernel quantile regression. *Neural Computation*, 21(2):539–559, 2009.
- K. Tanizaki, N. Hashimoto, Y. Inatsu, H. Hontani, and I. Takeuchi. Computing valid p-values for image segmentation by selective inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9553–9562, 2020.
- J. Taylor and R. Tibshirani. Post-selection inference for-penalized likelihood models. *Canadian Journal of Statistics*, 46(1):41–61, 2018.
- Y. Terada and H. Shimodaira. Selective inference after variable selection via multiscale bootstrap. *arXiv preprint arXiv:1905.10573*, 2019.
- X. Tian and J. Taylor. Selective inference with a randomized response. *The Annals of Statistics*, 46(2):679–710, 2018.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- R. J. Tibshirani. The lasso problem and uniqueness. *Electronic Journal of statistics*, 7:1456–1490, 2013.
- R. J. Tibshirani, J. Taylor, R. Lockhart, and R. Tibshirani. Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association*, 111(514):600–620, 2016.
- K. Tsuda. Entire regularization paths for graph data. In *In Proc. of ICML 2007*, pages 919–925, 2007.
- F. Yang, R. F. Barber, P. Jain, and J. Lafferty. Selective inference for group-sparse linear models. In *Advances in Neural Information Processing Systems*, pages 2469–2477, 2016.
- H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320, 2005.