

Exact Post Model Selection Inference for Marginal Screening

Jason D. Lee and Jonathan E. Taylor

February 23, 2014

Abstract

We develop a framework for post model selection inference, via marginal screening, in linear regression. At the core of this framework is a result that characterizes the exact distribution of linear functions of the response y , conditional on the model being selected (“condition on selection” framework). This allows us to construct valid confidence intervals and hypothesis tests for regression coefficients that account for the selection procedure. In contrast to recent work in high-dimensional statistics, our results are exact (non-asymptotic) and require no eigenvalue-like assumptions on the design matrix X . Furthermore, the computational cost of marginal regression, constructing confidence intervals and hypothesis testing is negligible compared to the cost of linear regression, thus making our methods particularly suitable for extremely large datasets. Although we focus on marginal screening to illustrate the applicability of the condition on selection framework, this framework is much more broadly applicable. We show how to apply the proposed framework to several other selection procedures including orthogonal matching pursuit, non-negative least squares, and marginal screening+Lasso.

1 Introduction

Consider the model

$$y_i = \mu(x_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma^2 I), \quad (1)$$

where $\mu(x)$ is an arbitrary function, and $x_i \in \mathbf{R}^p$. Our goal is to perform inference on $(X^T X)^{-1} X^T \mu$, which is the best linear predictor of μ . In the classical setting of $n > p$, the least squares estimator

$$\hat{\beta} = (X^T X)^{-1} X^T y \quad (2)$$

is a commonly used estimator for $(X^T X)^{-1} X^T \mu$. Under the linear model assumption $\mu = X\beta^0$, the exact distribution of $\hat{\beta}$ is

$$\hat{\beta} \sim \mathcal{N}(\beta^0, \sigma^2(X^T X)^{-1}). \quad (3)$$

Using the normal distribution, we can test the hypothesis $H_0 : \beta_j^0 = 0$ and form confidence intervals for β_j^0 using the z-test.

However in the high-dimensional $p > n$ setting, the least squares estimator is an underdetermined problem, and the predominant approach is to perform variable selection or model selection [3]. There are many approaches to variable selection including AIC/BIC, greedy algorithms such as forward stepwise regression, orthogonal matching pursuit, and regularization methods such as the Lasso. The focus of this paper will be on the model selection procedure known as marginal screening, which selects the k most correlated features x_j with the response y .

Marginal screening is the simplest and most commonly used of the variable selection procedures [13, 28, 20]. Marginal screening requires only $O(np)$ computation and is several orders of magnitude faster than regularization methods such as the Lasso; it is extremely suitable for extremely large datasets where the Lasso may be computationally intractable to apply. Furthermore, the selection properties are comparable to the Lasso [12]. In the ultrahigh dimensional setting $p = O(e^{n^k})$, marginal screening is shown to have the SURE screening property, $P(S \subset \hat{S})$, that is marginal screening selects a superset of the truly relevant variables [9, 11, 10]. Marginal screening can also be combined with a second variable selection procedure such as the Lasso to further reduce the dimensionality; our statistical inference methods extend to the Marginal Screening+Lasso method.

Since marginal screening utilizes the response variable y , the confidence intervals and statistical tests based on the distribution in (3) are not valid; confidence intervals with nominal $1 - \alpha$ coverage may no longer cover at the advertised level:

$$\Pr(\beta_j^0 \in C_{1-\alpha}(x)) < 1 - \alpha.$$

Several authors have previously noted this problem including recent work in [17, 18, 19, 2]. A major line of work [17, 18, 19] has described the difficulty of inference post model selection: the distribution of post model selection estimates is complicated and cannot be approximated in a uniform sense by their asymptotic counterparts.

In this paper, we describe how to form exact confidence intervals for linear regression coefficients *post model selection*. We assume the model

(1), and operate under the fixed design matrix X setting. The linear regression coefficients constrained to a subset of variables S is linear in μ , $e_j^T (X_S^T X_S)^{-1} X_S^T \mu = \eta^T \mu$ for some η . We derive the conditional distribution of $\eta^T y$ for any vector η , so we are able to form confidence intervals and test regression coefficients.

In Section 2 we discuss related work on high-dimensional statistical inference, and Section 3 introduces the marginal screening algorithm and shows how z intervals may fail to have the correct coverage properties. Section 4 and 5 show how to represent the marginal screening selection event as constraints on y , and construct pivotal quantities for the truncated Gaussian. Section 6 uses these tools to develop valid hypothesis tests and confidence intervals.

Although the focus of this paper is on marginal screening, the “condition on selection” framework, first proposed for the Lasso in [16], is much more general; we use marginal screening as a simple and clean illustration of the applicability of this framework. In Section 7, we discuss several extensions including how to apply the framework to other variable/model selection procedures and to nonlinear regression problems. Section 7 covers

1. marginal screening+Lasso, a screen and clean procedure that first uses marginal screening and cleans with the Lasso,
2. orthogonal matching pursuit (OMP)
3. non-negative least squares (NNLS).

2 Related Work

Most of the theoretical work on high-dimensional linear models focuses on *consistency*. Such results establish, under restrictive assumptions on X , the Lasso $\hat{\beta}$ is close to the unknown β^0 [24] and selects the correct model [33, 30, 15]. We refer to the reader to [3] for a comprehensive discussion about the theoretical properties of the Lasso.

There is also recent work on obtaining confidence intervals and significance testing for penalized M-estimators such as the Lasso. One class of methods uses sample splitting or subsampling to obtain confidence intervals and p-values [31, 23]. In the post model selection literature, the recent work of [2] proposed the POSI approach, a correction to the usual t-test confidence intervals by controlling the familywise error rate for all parameters in any possible submodel. The POSI approach will produce valid confidence intervals for any possible model selection procedure; however for a given model

selection procedure such as marginal regression, it will be conservative. In addition, the POSI methodology is extremely computationally intensive and currently only applicable for $p \leq 30$.

A separate line of work establishes the asymptotic normality of a corrected estimator obtained by “inverting” the KKT conditions [29, 32, 14]. The corrected estimator \hat{b} has the form $\hat{b} = \hat{\beta} + \lambda \hat{\Theta} \hat{z}$, where \hat{z} is a subgradient of the penalty at $\hat{\beta}$ and $\hat{\Theta}$ is an approximate inverse to the Gram matrix $X^T X$. The two main drawbacks to this approach are 1) the confidence intervals are valid only when the M-estimator is consistent, and thus require restricted eigenvalue conditions on X , 2) obtaining $\hat{\Theta}$ is usually much more expensive than obtaining $\hat{\beta}$, and 3) the method is specific to regularized estimators, and does not extend to marginal screening, forward stepwise, and other variable selection methods.

Most closely related to our work is the “condition on selection” framework laid out in [16] for the Lasso. Our work extends this methodology to other variable selection methods such as marginal screening, marginal screening followed by the Lasso (marginal screening+Lasso), orthogonal matching pursuit, and non-negative least squares. The primary contribution of this work is the observation that many model selection methods, including marginal screening and Lasso, lead to “selection events” that can be represented as a set of constraints on the response variable y . By conditioning on the selection event, we can characterize the exact distribution of $\eta^T y$. This paper focuses on marginal screening, since it is the simplest of variable selection methods, and thus the applicability of the “conditioning on selection event” framework is most transparent. However, this framework is not limited to marginal screening and can be applied to a wide a class of model selection procedures including greedy algorithms such as matching pursuit and orthogonal matching pursuit. We discuss some of these possible extensions in Section 7, but leave a thorough investigation to future work.

A remarkable aspect of our work is that we only assume X is in general position, and the test is exact, meaning the distributional results are true even under finite samples. By extension, we do not make any assumptions on n and p , which is unusual in high-dimensional statistics [3]. Furthermore, the computational requirements of our test are negligible compared to computing the linear regression coefficients.

Our test assumes that the noise variance σ^2 is known. However, there are many methods for estimating σ^2 in high dimensions. A data splitting technique is used in [8], while [27] proposes a method that computes the regression estimate and an estimate of the variance simultaneously. We refer the reader to [25] for a survey and comparison of the various methods,

and assume σ^2 is known for the remainder of the paper.

3 Marginal Screening

Let $X \in \mathbb{R}^{n \times p}$ be the design matrix, $y \in \mathbb{R}^n$ the response variable, and assume the model

$$y_i = \mu(x_i) + \epsilon_i, \epsilon_i \sim \mathcal{N}(0, \sigma^2 I).$$

We will assume that X is in general position and has unit norm columns. The algorithm estimates $\hat{\beta}$ via Algorithm 1. The marginal screening algo-

Algorithm 1 Marginal screening algorithm

- 1: **Input:** Design matrix X , response y , and model size k .
 - 2: Compute $|X^T y|$.
 - 3: Let \hat{S} be the index of the k largest entries of $|X^T y|$.
 - 4: Compute $\hat{\beta}_{\hat{S}} = (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T y$
-

rithm chooses the k variables with highest absolute dot product with y , and then fits a linear model over those k variables. We will assume $k \leq \min(n, p)$. For any fixed subset of variables S , the distribution of $\hat{\beta}_S = (X_S^T X_S)^{-1} X_S^T y$ is

$$\hat{\beta}_S \sim \mathcal{N}(\beta_S^*, \sigma^2 (X_S^T X_S)^{-1}) \quad (4)$$

$$\beta_S^* := (X_S^T X_S)^{-1} X_S^T \mu. \quad (5)$$

We will use the notation $\beta_{j \in S}^* := (\beta_S^*)_j$, where j is indexing a variable in the set S . The z-test intervals for a regression coefficient are

$$C(\alpha, j, S) := \left(\hat{\beta}_{j \in S} - \sigma z_{1-\alpha/2} (X_S^T X_S)_{jj}, \hat{\beta}_{j \in S} + \sigma z_{1-\alpha/2} (X_S^T X_S)_{jj} \right) \quad (6)$$

and each interval has $1 - \alpha$ coverage, meaning $\Pr(\beta_{j \in S}^* \in C(\alpha, j, S)) = 1 - \alpha$.

However if \hat{S} is chosen using a model selection procedure that depends on y , the distributional result (5) no longer holds and the z-test intervals will not cover at the $1 - \alpha$ level. It is possible that

$$\Pr(\beta_{j \in \hat{S}}^* \in C(\alpha, j, \hat{S})) < 1 - \alpha.$$

Similarly, the test of the hypothesis $H_0 : \beta_{j \in \hat{S}}^* = 0$ will not control type I error at level α , meaning $\Pr(\text{reject } H_0 | H_0) > \alpha$.

3.1 Failure of z-test confidence intervals

We will illustrate empirically that the z-test intervals do not cover at $1 - \alpha$ when \hat{S} is chosen by marginal screening in Algorithm 1. For this experiment




Figure 1: Plots of the coverage proportion across a range of SNR (log-scale). We see that the coverage proportion of the z intervals can be far below the nominal level of $1 - \alpha = .9$, even at $\text{SNR} = 5$. The adjusted intervals always have coverage proportion $.9$.

we generated X from a standard normal with $n = 20$ and $p = 200$. The signal vector is 2 sparse with $\beta_1^0, \beta_2^0 = \text{SNR}$, $y = X\beta^0 + \epsilon$, and $\epsilon \sim N(0, 1)$. The confidence intervals were constructed for the $k = 2$ variables selected by the marginal screening algorithm. The z-test intervals were constructed via (6) with $\alpha = .1$, and the adjusted intervals were constructed using Algorithm 3. The results are described in Figure 1. The y-axis plots the coverage proportion or the fraction of times the true parameter value fell in the confidence interval. Each point represents 500 independent trials. The x-axis varies the SNR parameter over the values 0.1, .2, .5, 1, 2, 5, 10. From the figure, we see that the z intervals can have coverage proportion drastically less than the nominal level of $1 - \alpha = .9$, and only for $\text{SNR} = 10$ does the coverage tend to $.9$. This motivates the need for intervals that have the correct coverage proportion after model selection.

4 Representing the selection event

Since Equation (5) does not hold for a selected \hat{S} when the selection procedure depends on y , the z-test intervals are not valid. Our strategy will be to understand the conditional distribution of y and contrasts (linear functions of y) $\eta^T y$, then construct inference conditional on the selection event \hat{E} . We will use $\hat{E}(y)$ to represent a random variable, and E to represent an element of the range of $\hat{E}(y)$. In the case of marginal screening, the selection event $\hat{E}(y)$ corresponds to the set of selected variables \hat{S} and signs s :

$$\begin{aligned}\hat{E}(y) &= \left\{ y : \text{sign}(x_i^T y) x_i^T y > \pm x_j^T y \text{ for all } i \in \hat{S} \text{ and } j \in \hat{S}^c \right\} \\ &= \left\{ y : \hat{s}_i x_i^T y > \pm x_j^T y \text{ and } \hat{s}_i x_i^T y \geq 0 \text{ for all } i \in \hat{S} \text{ and } j \in \hat{S}^c \right\} \\ &= \left\{ y : A(\hat{S}, \hat{s})y \leq b(\hat{S}, \hat{s}) \right\}\end{aligned}\tag{7}$$

for some matrix $A(\hat{S}, \hat{s})$ and vector $b(\hat{S}, \hat{s})$ ¹. We will use the selection event \hat{E} and the selected variables/signs pair (\hat{S}, \hat{s}) interchangeably since they are in bijection.

The space \mathbf{R}^n is partitioned by the selection events,

$$\mathbf{R}^n = \bigsqcup_{(S,s)} \{y : A(S, s)y \leq b(S, s)\}.$$

The vector y can be decomposed with respect to the partition as follows

$$y = \sum_{S,s} y \mathbb{1}(A(S, s)y \leq b(S, s))\tag{8}$$

The previous equation establishes that y is a different constrained Gaussian for each element of the partition, where the partition is specified by a possible subset of variables and signs (S, s) . The above discussion can be summarized in the following theorem.

Theorem 4.1. *The distribution of y conditional on the selection event is a constrained Gaussian,*

$$y|\{\hat{E}(y) = E\} \stackrel{d}{=} z|\{A(S, s)z \leq b\}, \quad z \sim \mathcal{N}(\mu, \sigma^2 I).$$

Proof. The event E is in bijection with a pair (S, s) , and y is unconditionally Gaussian. Thus the conditional $y|\{A(S, s)y \leq b(S, s)\}$ is a Gaussian constrained to the set $\{A(S, s)y \leq b(S, s)\}$. \square

¹ b can be taken to be 0 for marginal screening, but this extra generality is needed for other model selection methods

5 Truncated Gaussian test

This section summarizes the recent tools developed in [16] for testing contrasts² $\eta^T y$ of a constrained Gaussian y . The results are stated without proof and the proofs can be found in [16].

The distribution of a constrained Gaussian $y \sim N(\mu, \Sigma)$ conditional on affine constraints $\{Ay \leq b\}$ has density $\frac{1}{\Pr(Ay \leq b)} f(y; \mu, \Sigma) \mathbb{1}(Ay \leq b)$, involves the intractable normalizing constant $\Pr(Ay \leq b)$. In this section, we derive a one-dimensional pivotal quantity for $\eta^T \mu$. This pivot relies on characterizing the distribution of $\eta^T y$ as a truncated normal. The key step to deriving this pivot is the following lemma:

Lemma 5.1. *The conditioning set can be rewritten in terms of $\eta^T y$ as follows:*

$$\{Ay \leq b\} = \{\mathcal{V}^-(y) \leq \eta^T y \leq \mathcal{V}^+(y), \mathcal{V}^0(y) \geq 0\}$$

where

$$\alpha = \frac{A\Sigma\eta}{\eta^T\Sigma\eta} \quad (9)$$

$$\mathcal{V}^- = \mathcal{V}^-(y) = \max_{j: \alpha_j < 0} \frac{b_j - (Ay)_j + \alpha_j \eta^T y}{\alpha_j} \quad (10)$$

$$\mathcal{V}^+ = \mathcal{V}^+(y) = \min_{j: \alpha_j > 0} \frac{b_j - (Ay)_j + \alpha_j \eta^T y}{\alpha_j}. \quad (11)$$

$$\mathcal{V}^0 = \mathcal{V}^0(y) = \min_{j: \alpha_j = 0} b_j - (Ay)_j \quad (12)$$

Moreover, $(\mathcal{V}^+, \mathcal{V}^-, \mathcal{V}^0)$ are independent of $\eta^T y$.

The geometric picture gives more intuition as to why \mathcal{V}^+ and \mathcal{V}^- are independent of $\eta^T y$. Without loss of generality, we assume $\|\eta\|_2 = 1$ and $y \sim N(\mu, I)$ (otherwise we could replace y by $\Sigma^{-\frac{1}{2}}y$). Now we can decompose y into two independent components, a 1-dimensional component $\eta^T y$ and an $(n-1)$ -dimensional component orthogonal to η :

$$y = \eta^T y + P_{\eta^\perp} y.$$

The case of $n = 2$ is illustrated in Figure 2. Since the two components are independent, the distribution of $\eta^T y$ is the same as $\eta^T y | \{P_{\eta^\perp} y\}$. If we condition on $P_{\eta^\perp} y$, it is clear from Figure 2 that in order for y to lie in the set, it is necessary for $\mathcal{V}^- \leq \eta^T y \leq \mathcal{V}^+$, where \mathcal{V}^- and \mathcal{V}^+ are functions of $P_{\eta^\perp} y$.

²A contrast of y is a linear function of the form $\eta^T y$.




Figure 2: A picture demonstrating that the set $\{Ay \leq b\}$ can be characterized by $\{\mathcal{V}^- \leq \eta^T y \leq \mathcal{V}^+\}$. Assuming $\Sigma = I$ and $\|\eta\|_2 = 1$, \mathcal{V}^- and \mathcal{V}^+ are functions of $P_{\eta^\perp} y$ only, which is independent of $\eta^T y$.

Corollary 5.2. *The distribution of $\eta^T y$ conditioned on $\{Ay \leq b, \mathcal{V}^+(y) = v^+, \mathcal{V}^-(y) = v^-\}$ is a (univariate) Gaussian truncated to fall between \mathcal{V}^- and \mathcal{V}^+ , i.e.*

$$\eta^T y \mid \{Ay \leq b, \mathcal{V}^+(y) = v^+, \mathcal{V}^-(y) = v^-\} \stackrel{d}{=} W$$

where $W \sim TN(\eta^T \mu, \eta^T \Sigma \eta, v^-, v^+)$. $TN(\mu, \sigma, a, b)$ is the normal distribution truncated to lie between a and b .

In Figure 3, we plot the density of the truncated Gaussian, noting that its shape depends on the location of μ relative to $[a, b]$ as well as the width relative to σ .

The following pivotal quantity³ follows from Corollary 5.2 via the probability integral transform.

Theorem 5.3. *Let $\Phi(x)$ denote the CDF of a $N(0, 1)$ random variable, and let $F_{\mu, \sigma^2}^{[a, b]}$ denote the CDF of $TN(\mu, \sigma, a, b)$, i.e.:*

$$F_{\mu, \sigma^2}^{[a, b]}(x) = \frac{\Phi((x - \mu)/\sigma) - \Phi((a - \mu)/\sigma)}{\Phi((b - \mu)/\sigma) - \Phi((a - \mu)/\sigma)}. \quad (13)$$

³The distribution of a pivotal quantity does not depend on unobserved parameters.




Figure 3: The density of the truncated Gaussian $TN(\mu, \sigma^2, a, b)$ depends on the width of $[a, b]$ relative to σ as well as the location of μ relative to $[a, b]$. When μ is firmly inside the interval, the distribution resembles a Gaussian. As μ leaves $[a, b]$, the density begins to converge to an exponential distribution with mean inversely proportional to the distance between μ and its projection onto $[a, b]$.

Then $F_{\eta^T \mu, \eta^T \Sigma \eta}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta^T y)$ is a pivotal quantity, conditional on $\{Ay \leq b\}$:

$$F_{\eta^T \mu, \eta^T \Sigma \eta}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta^T y) \mid \{Ay \leq b\} \sim \text{Unif}(0, 1) \quad (14)$$

where \mathcal{V}^- and \mathcal{V}^+ are defined in (10) and (11).

6 Inference for marginal screening

In this section, we apply the theory summarized in Sections 4 and 5 to marginal screening. In particular, we will construct confidence intervals for the selected variables.

To summarize the developments so far, recall that our model (1) says that $y \sim N(\mu, \sigma^2 I)$. The distribution of interest is $y \mid \{\hat{E}(y) = E\}$, and by Theorem 4.1, this is equivalent to $y \mid \{A(S, s)z \leq b(S, s)\}$, where $y \sim N(\mu, \sigma^2 I)$. By applying Theorem 5.3, we obtain the pivotal quantity

$$F_{\eta^T \mu, \sigma^2 \|\eta\|_2^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta^T y) \mid \{\hat{E}(y) = E\} \sim \text{Unif}(0, 1) \quad (15)$$

for any η , where \mathcal{V}^- and \mathcal{V}^+ are defined in (10) and (11).




Figure 4: Histogram and qq plot of $F_{\eta^T \mu, \eta^T \Sigma \eta}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta^T y)$ where y is a constrained Gaussian. The distribution is very close to $\text{Unif}(0, 1)$, which is in agreement with Theorem 5.3.

6.1 Hypothesis tests for selected variables

In this section, we describe how to form confidence intervals for the components of $\beta_{\hat{S}}^* = (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T \mu$. The best linear predictor of μ that uses only the selected variables is $\beta_{\hat{S}}^*$, and $\hat{\beta}_{\hat{S}} = (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T y$ is an unbiased estimate of $\beta_{\hat{S}}^*$. In this section, we propose hypothesis tests and confidence intervals for $\beta_{\hat{S}}^*$. If we choose

$$\eta_j = ((X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T e_j)^T, \quad (16)$$

then $\eta_j^T \mu = \beta_{j \in \hat{S}}^*$, so the above framework provides a method for inference about the j^{th} variable in the model \hat{S} . This choice of η is not fixed before marginal screening selects \hat{S} , but it is measurable with respect to the σ -algebra generated by the partition. Since it is measurable, η is constant on each partition, so the pivot is uniformly distributed on each element of the partition, and thus uniformly distributed for all y .

If we assume the linear model $\mu = X\beta^0$ for some $\beta^0 \in \mathbf{R}^p$, $S^0 := \text{support}(\beta^0) \subset \hat{S}$, and $X_{\hat{S}}$ is full rank, then by the following computation $\beta_{\hat{S}}^* = \beta_{\hat{S}}^0$:

$$\begin{aligned} \beta_{\hat{S}}^* &= (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T X_S \beta_S^0 \\ &= (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T X_{\hat{S}} \beta_{\hat{S}}^0 \\ &= \beta_{\hat{S}}^0 \end{aligned}$$

In [9], the screening property $S^0 \subset \hat{S}$ for the marginal screening algorithm is

established under mild conditions. Thus under the screening property, our method provides hypothesis tests and confidence intervals for $\beta_{\hat{S}}^0$.

By applying Theorem 5.3, we obtain the following (conditional) pivot for $\beta_{j \in \hat{S}}^*$:

$$F_{\beta_{j \in \hat{S}}^*, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) \Big| \{\hat{E}(y) = E\} \sim \text{Unif}(0, 1). \quad (17)$$

The quantities j and η_j are both random through \hat{E} , a quantity which is fixed after conditioning, therefore Theorem 5.3 holds even for this choice of η .

Consider testing the hypothesis $H_0 : \beta_{j \in \hat{S}}^* = \beta_j$. A valid test statistic is given by $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y})$, which is uniformly distributed under the null hypothesis and $y | \{\hat{E}(y) = E\}$. Thus, this test would reject when $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) > 1 - \frac{\alpha}{2}$ or $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) < \frac{\alpha}{2}$.

Theorem 6.1. *The test of $H_0 : \beta_{j \in \hat{S}}^* = \beta_j$ that accepts when*

$$\frac{\alpha}{2} < F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) < 1 - \frac{\alpha}{2}$$

is an α level test of H_0 .

Proof. Under H_0 , we have $\beta_{j \in \hat{S}}^* = \beta_j$, so by (17) $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) | \{\hat{E}(y) = E\}$ is uniformly distributed. Thus

$$\Pr\left(\frac{\alpha}{2} < F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) \leq 1 - \frac{\alpha}{2} \mid \{\hat{E}(y) = E, H_0\}\right) = 1 - \alpha,$$

and the type 1 error is exactly α . Under H_0 , but not conditional on selection event \hat{E} , we have

$$\begin{aligned} & \Pr\left(\frac{\alpha}{2} < F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) \leq 1 - \frac{\alpha}{2} \mid H_0\right) \\ &= \sum_E \Pr\left(\frac{\alpha}{2} < F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T \mathbf{y}) \leq 1 - \frac{\alpha}{2} \mid \{\hat{E}(y) = E, H_0\}\right) \Pr(\hat{E}(y) = E \mid H_0) \\ &= \sum_E (1 - \alpha) \Pr(\hat{E}(y) = E \mid H_0) \\ &= (1 - \alpha) \sum_E \Pr(\hat{E}(y) = E \mid H_0) \\ &= 1 - \alpha. \end{aligned}$$

For each element of the partition E , the conditional (on selection) hypothesis test is level $1 - \alpha$, so by summing over the partition the unconditional test is level $1 - \alpha$. \square

Our hypothesis test is not conservative, in the sense that the type 1 error is exactly α ; also, it is non-asymptotic, since the statement holds for fixed n and p . We summarize the hypothesis test in this section in the following algorithm.

Algorithm 2 Hypothesis test for selected variables

- 1: **Input:** Design matrix X , response y , model size k .
 - 2: Use Algorithm 1 to select a subset of variables \hat{S} and signs $\hat{s} = \text{sign}(X_{\hat{S}}^T y)$.
 - 3: Specify the null hypothesis $H_0 : \beta_{j \in \hat{S}}^* = \beta_j$.
 - 4: Let $A = A(\hat{S}, \hat{s})$ and $b = b(\hat{S}, \hat{s})$ using (7). Let $\eta_j = (X_{\hat{S}}^T)^\dagger e_j$.
 - 5: Compute $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y)$, where \mathcal{V}^- and \mathcal{V}^+ are computed via (10) and (11) using the A , b , and η previously defined.
 - 6: **Output:** Reject if $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) > \frac{\alpha}{2}$ or $F_{\beta_j, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) < 1 - \frac{\alpha}{2}$.
-

6.2 Confidence intervals for selected variables

Next, we discuss how to obtain confidence intervals for $\beta_{j \in \hat{S}}^*$. The standard way to obtain an interval is to invert a pivotal quantity [4]. In other words, since

$$\Pr \left(\frac{\alpha}{2} \leq F_{\beta_{j \in \hat{S}}^*, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) \leq 1 - \frac{\alpha}{2} \mid \{\hat{E} = E\} \right) = \alpha,$$

one can define a $(1 - \alpha)$ (conditional) confidence interval for $\beta_{j, \hat{E}}^*$ as

$$\left\{ x : \frac{\alpha}{2} \leq F_{x, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) \leq 1 - \frac{\alpha}{2} \right\}. \quad (18)$$

In fact, F is monotone decreasing in x , so to find its endpoints, one need only solve for the root of a smooth one-dimensional function. The monotonicity is a consequence of the fact that the truncated Gaussian distribution is a natural exponential family and hence has monotone likelihood ratio in μ [21]. Since the truncated Gaussian distribution is in the exponential family, we can also form uniformly most accurate unbiased confidence intervals by solving for a root of a system of two equations [21]. For simplicity, we will use the interval in (18), but we expect that the intervals can be shortened by using the uniformly most accurate unbiased intervals.

We now formalize the above observations in the following result, an immediate consequence of Theorem 5.3.

Corollary 6.2. Let η_j be defined as in (16), and let $L_\alpha = L_\alpha(\eta_j, (\hat{S}, \hat{s}))$ and $U_\alpha = U_\alpha(\eta_j, (\hat{S}, \hat{s}))$ be the (unique) values satisfying

$$F_{L_\alpha, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) = 1 - \frac{\alpha}{2} \quad F_{U_\alpha, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y) = \frac{\alpha}{2} \quad (19)$$

Then $[L_\alpha, U_\alpha]$ is a $(1 - \alpha)$ confidence interval for $\beta_{j \in \hat{S}}^*$, conditional on \hat{E} :

$$\mathbb{P} \left(\beta_{j \in \hat{S}}^* \in [L_\alpha, U_\alpha] \mid \{\hat{E} = E\} \right) = 1 - \alpha. \quad (20)$$

Proof. The confidence region of $\beta_{j \in \hat{S}}^*$ is the set of β_j such that the test of $H_0 : \beta_{j \in \hat{S}}^*$ accepts at the $1 - \alpha$ level. The function $F_{x, \sigma^2 \|\eta_j\|^2}^{[\mathcal{V}^-, \mathcal{V}^+]}(\eta_j^T y)$ is monotone in x , so solving for L_α and U_α identify the most extreme values where H_0 is still accepted. This gives a $1 - \alpha$ confidence interval. \square

In relation to the literature on False Coverage Rate (FCR) [1], our procedure also controls the FCR.

Lemma 6.3. For each $j \in \hat{S}$,

$$\Pr \left(\beta_{j \in \hat{S}}^* \in [L_\alpha^j, U_\alpha^j] \right) = 1 - \alpha. \quad (21)$$

Furthermore, the FCR of the intervals $\{[L_\alpha^j, U_\alpha^j]\}_{j \in \hat{E}}$ is α .

Proof. By (20), the conditional coverage of the confidence intervals are $1 - \alpha$. The coverage holds for every element of the partition $\{\hat{E}(y) = E\}$, so

$$\begin{aligned} & \Pr \left(\beta_{j \in \hat{S}}^* \in [L_\alpha^j, U_\alpha^j] \right) \\ &= \sum_E \Pr \left(\beta_{j \in \hat{S}}^* \in [L_\alpha, U_\alpha] \mid \{\hat{E} = E\} \right) \Pr(\hat{E} = E) \\ &= \sum_E (1 - \alpha) \Pr(\hat{E} = E) \\ &= (1 - \alpha) \sum_E \Pr(\hat{E} = E) \\ &= 1 - \alpha. \end{aligned}$$

\square

We summarize the algorithm for selecting and constructing confidence intervals below.




Figure 5: Comparison of adjusted and unadjusted 90% confidence interval for $TN(\mu, \sigma^2, a, b)$. The upper and lower bounds of 90% confidence intervals are based on $[a, b] = [-3\sigma, 3\sigma]$, and the x-axis plots the observation on the scale $\frac{x}{\sigma}$. We see that as long as the observation $\frac{x}{\sigma}$ is roughly 0.5σ away from either boundary, the size of the intervals is comparable to an unadjusted confidence interval. However, the adjusted intervals are guaranteed to have the correct coverage, whereas it is unknown when the unadjusted intervals have the correct coverage.

6.3 Experiments on Diabetes dataset

In Figure 1, we have already seen that the confidence intervals constructed using Algorithm 3 have exactly $1 - \alpha$ coverage proportion. In this section, we perform an experiment on real data where the linear model does not hold, the noise is not Gaussian, and the noise variance is unknown. The diabetes dataset contains $n = 442$ diabetes patients measured on $p = 10$ baseline variables [6]. The baseline variables are age, sex, body mass index, average blood pressure, and six blood serum measurements, and the response y is a quantitative measure of disease progression measured one year after the baseline. The goal is to use the baseline variables to predict y , the measure of disease progression after one year, and determine which baseline variables are statistically significant for predicting y .

Since the noise variance σ^2 is unknown, we estimate it by $\sigma^2 = \frac{\|y - \hat{y}\|^2}{n - p}$, where $\hat{y} = X\hat{\beta}$ and $\hat{\beta} = (X^T X)^{-1} X^T y$. For each trial we generated new responses $\tilde{y}_i = X\hat{\beta} + \tilde{\epsilon}$, and $\tilde{\epsilon}$ is bootstrapped from the residuals $r_i = y_i -$

Algorithm 3 Confidence intervals for selected variables

- 1: **Input:** Design matrix X , response y , model size k .
 - 2: Use Algorithm 1 to select a subset of variables \hat{S} and signs $\hat{s} = \text{sign}(X_{\hat{S}}^T y)$.
 - 3: Let $A = A(\hat{S}, \hat{s})$ and $b = b(\hat{S}, \hat{s})$ using (7). Let $\eta_j = (X_{\hat{S}}^T)^\dagger e_j$.
 - 4: Solve for L_α^j and U_α^j using Equation (19) where \mathcal{V}^- and \mathcal{V}^+ are computed via (10) and (11) using the A , b , and η_j previously defined.
 - 5: **Output:** Return the intervals $[L_\alpha^j, U_\alpha^j]$ for $j \in \hat{S}$.
-




Figure 6: Plot of $1 - \alpha$ vs the coverage proportion for diabetes dataset. The nominal curve is the line $y = x$. The coverage proportion of the adjusted intervals agree with the nominal coverage level, but the z-test coverage proportion is strictly below the nominal level. The adjusted intervals perform well, despite the noise being non-Gaussian, and σ^2 unknown.

\hat{y}_i . This is known as the residual bootstrap, and is a standard method for assessing statistical procedures when the underlying model is unknown [7]. We used marginal screening to select $k = 2$ variables, and then fit linear regression on the selected variables. The adjusted confidence intervals were constructed using Algorithm 3 with the estimated σ^2 . The nominal coverage level is varied across $1 - \alpha \in \{.5, .6, .7, .8, .9, .95, .99\}$. From Figure 6, we observe that the adjusted intervals always cover at the nominal level, whereas the z-test is always below. The experiment was repeated 2000 times.

7 Extensions

The purpose of this section is to illustrate the broad applicability of the condition on selection framework. This framework was first proposed in [16] to form valid hypothesis tests and confidence intervals after model selection via the Lasso. However, the framework is not restricted to the Lasso, and we have shown how to apply it to marginal screening. For expository purposes, we focused the paper on marginal screening where the framework is particularly easy to understand. In the rest of this section, we show how to apply the framework to marginal screening+Lasso, orthogonal matching pursuit, and non-negative least squares. This is a non-exhaustive list of selection procedures where the condition on selection framework is applicable, but we hope this incomplete list emphasizes the ease of constructing tests and confidence intervals post-model selection via conditioning.

7.1 Marginal screening + Lasso

The marginal screening+Lasso procedure was introduced in [9] as a variable selection method for the ultra-high dimensional setting of $p = O(e^{n^k})$. Fan et al. [9] recommend applying the marginal screening algorithm with $k = n - 1$, followed by the Lasso on the selected variables. This is a two-stage procedure, so to properly account for the selection we must encode the selection event of marginal screening followed by Lasso. This can be done by representing the two stage selection as a single event. Let (\hat{S}_m, \hat{s}_m) be the variables and signs selected by marginal screening, and the (\hat{S}_L, \hat{z}_L) be the variables and signs selected by Lasso [16]. In Proposition 2.2 of [16], it is shown how to encode the Lasso selection event (\hat{S}_L, \hat{z}_L) as a set of constraints $\{A_L y \leq b_L\}$ ⁴, and in Section 4 we showed how to encode the marginal screening selection event (\hat{S}_m, \hat{s}_m) as a set of constraints $\{A_m y \leq b_m\}$. Thus the selection event of marginal screening+Lasso can be encoded as $\{A_L y \leq b_L, A_m y \leq b_m\}$. Using these constraints, the hypothesis test and confidence intervals described in Algorithms 2 and 3 are valid for marginal screening+Lasso.

7.2 Orthogonal Matching Pursuit

Orthogonal matching pursuit (OMP) is a commonly used variable selection method. At each iteration, OMP selects the variable most correlated with

⁴The Lasso selection event is with respect to the Lasso optimization problem after marginal screening.

the residual r , and then recomputes the residual using the residual of least squares using the selected variables. The description of the OMP algorithm is given in Algorithm 4.

Algorithm 4 Orthogonal matching pursuit (OMP)

- 1: **Input:** Design matrix X , response y , and model size k .
 - 2: **for:** $i = 1$ to k
 - 3: $p_i = \arg \max_{j=1, \dots, p} |r_i^T x_j|$.
 - 4: $\hat{S}_i = \cup_{j=1}^i \{p_j\}$.
 - 5: $r_{i+1} = (I - X_{\hat{S}_i} X_{\hat{S}_i}^\dagger) y$.
 - 6: **end for**
 - 7: **Output:** $\hat{S} := \{p_1, \dots, p_k\}$, and $\hat{\beta}_{\hat{S}} = (X_{\hat{S}}^T X_{\hat{S}})^{-1} X_{\hat{S}}^T y$
-

Similar to Section 4, we can represent the OMP selection event as a set of linear constraints on y .

$$\begin{aligned}
\hat{E}(y) &= \{y : \text{sign}(x_{p_i}^T r_i) x_{p_i}^T r_i > \pm x_j^T r_i, \text{ for all } j \neq p_i \text{ and all } i \in [k]\} \\
&= \{y : \hat{s}_i x_{p_i}^T (I - X_{\hat{S}_{i-1}} X_{\hat{S}_{i-1}}^\dagger) y > \pm x_j^T (I - X_{\hat{S}_{i-1}} X_{\hat{S}_{i-1}}^\dagger) y \text{ and} \\
&\quad \hat{s}_i x_{p_i}^T (I - X_{\hat{S}_{i-1}} X_{\hat{S}_{i-1}}^\dagger) y > 0, \text{ for all } j \neq p_i, \text{ and all } i \in [k]\} \\
&= \left\{ y : A(\hat{S}_1, \dots, \hat{S}_k, \hat{s}_1, \dots, \hat{s}_k) \leq b(\hat{S}_1, \dots, \hat{S}_k, \hat{s}_1, \dots, \hat{s}_k) \right\}.
\end{aligned}$$

The selection event encodes that OMP selected a certain variable and the sign of the correlation of that variable with the residual, at steps 1 to k . The primary difference between the OMP selection event and the marginal screening selection event is that the OMP event also describes the order at which the variables were chosen. The marginal screening event only describes that the variable was among the top k most correlated, and not whether a variable was the most correlated or k th most correlated.

Since the selection event can be represented as constraints on y , the hypothesis test and confidence intervals described in Algorithms 2 and 3 are valid for OMP selected $\hat{\beta}_{\hat{S}}$.

7.3 Nonnegative Least Squares

Non-negative least squares (NNLS) is a simple modification of the linear regression estimator with non-negative constraints on β :

$$\arg \min_{\beta: \beta \geq 0} \frac{1}{2} \|y - X\beta\|^2. \tag{22}$$

Under a positive eigenvalue conditions on X , several authors [26, 22] have shown that NNLS is comparable to the Lasso in terms of prediction and estimation errors. The NNLS estimator also does not have any tuning parameters, since the sign constraint provides a natural form of regularization. NNLS has found applications when modeling non-negative data such as prices, incomes, count data. Non-negativity constraints arise naturally in non-negative matrix factorization, signal deconvolution, spectral analysis, and network tomography; we refer to [5] for a comprehensive survey of the applications of NNLS.

We show how our framework can be used to form exact hypothesis tests and confidence intervals for NNLS estimated coefficients. The primal dual solution pair $(\hat{\beta}, \hat{\lambda})$ is a solution iff the KKT conditions are satisfied,

$$\begin{aligned}\hat{\lambda}_i &:= -x_i^T(y - X\hat{\beta}) \geq 0 \text{ for all } i \\ \hat{\beta} &\geq 0.\end{aligned}$$

Let $\hat{S} = \{i : -x_i^T(y - X\hat{\beta}) = 0\}$. By complementary slackness $\hat{\beta}_{-\hat{S}} = 0$, where $-\hat{S}$ is the complement to the “active” variables \hat{S} chosen by NNLS. Given the active set we can solve the KKT equation for the value of $\hat{\beta}_{\hat{S}}$,

$$\begin{aligned}-X_{\hat{S}}^T(y - X\hat{\beta}) &= 0 \\ -X_{\hat{S}}^T(y - X_{\hat{S}}\hat{\beta}_{\hat{S}}) &= 0 \\ \hat{\beta}_{\hat{S}} &= X_{\hat{S}}^\dagger y,\end{aligned}$$

which is a linear contrast of y . The NNLS selection event is

$$\begin{aligned}\hat{E}(y) &= \{y : X_{\hat{S}}^T(y - X\hat{\beta}) = 0, X_{-\hat{S}}^T(y - X\hat{\beta}) > 0\} \\ &= \{y : X_{\hat{S}}^T(y - X\hat{\beta}) \geq 0, -X_{-\hat{S}}^T(y - X\hat{\beta}) \geq 0, X_{-\hat{S}}^T(y - X\hat{\beta}) > 0\} \\ &= \{y : X_{\hat{S}}^T(I - X_{\hat{S}}X_{\hat{S}}^\dagger)y \geq 0, -X_{-\hat{S}}^T(I - X_{\hat{S}}X_{\hat{S}}^\dagger)y \geq 0, X_{-\hat{S}}^T(I - X_{\hat{S}}X_{\hat{S}}^\dagger)y > 0\} \\ &= \{y : A(\hat{S})y \leq 0\}.\end{aligned}$$

The selection event encodes that for a given y the NNLS optimization program will select a subset of variables $\hat{S}(y)$. Similar to the case in OMP and marginal screening, we can use Algorithms 2 and 3, since the selection event is represented by a set of linear constraints $\{y : A(\hat{S})y \leq 0\}$.

8 Conclusion

Due to the increasing size of datasets, marginal screening has become an important method for fast variable selection. However, the standard hy-

pothesis tests and confidence intervals used in linear regression are invalid after using marginal screening to select important variables. We have described a method to perform hypothesis and form confidence intervals after marginal screening. The conditional on selection framework is not restricted to marginal screening, and also applies to OMP, marginal screening + Lasso, and NNLS.

References

- [1] Yoav Benjamini and Daniel Yekutieli. False discovery rate-adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469):71–81, 2005.
- [2] Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *Annals of Statistics*, 41(2):802–837, 2013.
- [3] Peter Lukas Bühlmann and Sara A van de Geer. *Statistics for High-dimensional Data*. Springer, 2011.
- [4] George Casella and Roger L Berger. *Statistical inference*, volume 70. Duxbury Press Belmont, CA, 1990.
- [5] Donghui Chen and Robert J Plemmons. Nonnegativity constraints in numerical analysis. In *Symposium on the Birth of Numerical Analysis*, pages 109–140, 2009.
- [6] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- [7] Bradley Efron and Robert Tibshirani. *An introduction to the bootstrap*, volume 57. CRC press, 1993.
- [8] Jianqing Fan, Shaojun Guo, and Ning Hao. Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society. Series B (Methodological)*, 74(1):37–65, 2012.
- [9] Jianqing Fan and Jinchi Lv. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5):849–911, 2008.
- [10] Jianqing Fan, Richard Samworth, and Yichao Wu. Ultrahigh dimensional feature selection: beyond the linear model. *The Journal of Machine Learning Research*, 10:2013–2038, 2009.
- [11] Jianqing Fan and Rui Song. Sure independence screening in generalized linear models with np-dimensionality. *The Annals of Statistics*, 38(6):3567–3604, 2010.

- [12] Christopher R Genovese, Jiashun Jin, Larry Wasserman, and Zhigang Yao. A comparison of the lasso and marginal regression. *The Journal of Machine Learning Research*, 98888:2107–2143, 2012.
- [13] Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *The Journal of Machine Learning Research*, 3:1157–1182, 2003.
- [14] Adel Javanmard and Andrea Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *arXiv preprint arXiv:1306.3171*, 2013.
- [15] Jason Lee, Yuekai Sun, and Jonathan E Taylor. On model selection consistency of penalized m-estimators: a geometric theory. In *Advances in Neural Information Processing Systems*, pages 342–350, 2013.
- [16] Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact inference after model selection via the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- [17] Hannes Leeb and Benedikt M Pötscher. The finite-sample distribution of post-model-selection estimators and uniform versus nonuniform approximations. *Econometric Theory*, 19(1):100–142, 2003.
- [18] Hannes Leeb and Benedikt M Pötscher. Model selection and inference: Facts and fiction. *Econometric Theory*, 21(1):21–59, 2005.
- [19] Hannes Leeb and Benedikt M Pötscher. Can one estimate the conditional distribution of post-model-selection estimators? *The Annals of Statistics*, pages 2554–2591, 2006.
- [20] Jeff Leek. Prediction: the lasso vs just using the top 10 predictors. <http://simplystatistics.tumblr.com/post/18132467723/prediction-the-lasso-vs-just-using-the-top-10>.
- [21] Erich L. Lehmann and Joseph P. Romano. *Testing Statistical Hypotheses*. Springer, 3 edition, 2005.
- [22] Nicolai Meinshausen et al. Sign-constrained least squares estimation for high-dimensional regression. *Electronic Journal of Statistics*, 7:1607–1631, 2013.
- [23] Nicolai Meinshausen, Lukas Meier, and Peter Bühlmann. P-values for high-dimensional regression. *Journal of the American Statistical Association*, 104(488), 2009.

- [24] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [25] Stephen Reid, Robert Tibshirani, and Jerome Friedman. A study of error variance estimation in lasso regression. *Preprint*, 2013.
- [26] Martin Slawski, Matthias Hein, et al. Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization. *Electronic Journal of Statistics*, 7:3004–3056, 2013.
- [27] Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [28] Virginia Goss Tusher, Robert Tibshirani, and Gilbert Chu. Significance analysis of microarrays applied to the ionizing radiation response. *Proceedings of the National Academy of Sciences*, 98(9):5116–5121, 2001.
- [29] Sara van de Geer, Peter Bühlmann, and Ya’acov Ritov. On asymptotically optimal confidence regions and tests for high-dimensional models. *arXiv preprint arXiv:1303.0518*, 2013.
- [30] M.J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). 55(5):2183–2202, 2009.
- [31] Larry Wasserman and Kathryn Roeder. High dimensional variable selection. *Annals of statistics*, 37(5A):2178, 2009.
- [32] Cun-Hui Zhang and S Zhang. Confidence intervals for low-dimensional parameters with high-dimensional data. *arXiv preprint arXiv:1110.2563*, 2011.
- [33] P. Zhao and B. Yu. On model selection consistency of lasso. 7:2541–2563, 2006.