

THE CCP SELECTOR: SCALABLE ALGORITHMS FOR SPARSE RIDGE REGRESSION FROM CHANCE-CONSTRAINED PROGRAMMING

BY WEIJUN XIE^{*} AND XINWEI DENG[†]

Department of Industrial and Systems Engineering^{} and Department of
Statistics[†], Virginia Tech, Blacksburg VA 24061*

Sparse regression and variable selection for large-scale data have been rapidly developed in the past decades. This work focuses on sparse ridge regression, which considers the exact L_0 norm to pursue the sparsity. We pave out a theoretical foundation to understand why many existing approaches may not work well for this problem, in particular on large-scale datasets. Inspired by reformulating the problem as a chance-constrained program, we derive a novel mixed integer second order conic (MISOC) reformulation and prove that its continuous relaxation is equivalent to that of the convex integer formulation proposed in a recent work. Based upon these two formulations, we develop two new scalable algorithms, the greedy and randomized algorithms, for sparse ridge regression with desirable theoretical properties. The proposed algorithms are proved to yield near-optimal solutions under mild conditions. In the case of much larger dimensions, we propose to integrate the greedy algorithm with the randomized algorithm, which can greedily search the features from the nonzero subset identified by the continuous relaxation of the MISOC formulation. The merits of the proposed methods are elaborated through a set of numerical examples in comparison with several existing ones.

MSC 2010 subject classifications: 62J07, 90C10, 90C15

Keywords and phrases: Ridge Regression, Chance Constraint, Mixed Integer, Conic Program, Approximation Algorithm

1. Introduction. As technology rapidly advances, modern statistical analysis often encounters regressions with a large number of explanatory variables (also known as features). Hence, sparse regression and variable selection have been studied intensively in the past decades. As an alternative to regular linear regression, ridge regression, first proposed by [22], has several desirable advantages including stable solution, estimator variance reduction, and efficient computation. However, its solution is usually neither sparse [18], nor applicable for variable selection.

To overcome these issues, we consider the sparse ridge regression problem as below:

$$(F0) \quad v^* = \min_{\boldsymbol{\beta}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 : \|\boldsymbol{\beta}\|_0 \leq k \right\},$$

where $\mathbf{y} \in \mathbb{R}^n$ denotes the response vector, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ represents the model matrix, $\boldsymbol{\beta} \in \mathbb{R}^p$ is the vector of regression coefficients (i.e., estimand), and $\lambda > 0$ is a positive tuning parameter for the ridge penalty (i.e., L_2 penalty). Here, $\|\boldsymbol{\beta}\|_0$ is the L_0 norm, which counts the number of nonzero entries of vector $\boldsymbol{\beta}$. The value of k represents the number of features to be chosen. In (F0), it aims to find a best k -sparse estimator, which minimizes the least square error with a squared L_2 penalty. It is easy to see that when $\lambda \rightarrow 0$, (F0) reduces to a special case, which is known as sparse regression. Note that in the signal process literature (cf. [40]), the formulation (F0) can also coincide with sparse signal recovery. Without loss of generality, let us assume that $k \leq n$ and $k \leq p$ and the data are normalized such that $\|\mathbf{x}_i\|_2 = 1$ for all $i \in [p] := \{1, 2, \dots, p\}$.

It is noted that the sparse ridge regression (F0) can be reformulated as a chance constrained program (CCP) with finite support [1, 26]. That is, we consider p scenarios with equal probability $\frac{1}{p}$, where the i th scenario set is $S^i := \{\boldsymbol{\beta} : \beta_i = 0\}$ for $i \in [p]$. The constraint $\|\boldsymbol{\beta}\|_0 \leq k$ means that at most $\frac{k}{p}$ portion of scenarios can be violated. Hence, we can reformulate (F0) as a CCP below

$$(F0\text{-CCP}) \quad v^* = \min_{\boldsymbol{\beta}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 : \frac{1}{p} \sum_{i \in [p]} \mathbb{I}(|\beta_i| \leq 0) \geq 1 - \frac{k}{p} \right\},$$

where $\mathbb{I}(\cdot)$ denotes the indicator function. In Section 2, we will investigate how applicable the recent progress on CCP (e.g., [1, 26, 31]) can be to solve (F0-CCP). It appears that many existing approaches may not work well due to the scalability issue or resulting in trivial solutions. In Section 4,

we propose two novel scalable algorithms and their integration to solve the sparse ridge regression with theoretical guarantees.

Relevant Literature. The ridge regression has been extensively studied in statistics [13, 27, 41]. However, although many desirable properties, the ridge estimator is often not sparse as it is computed by using the smoothed squared L_2 penalty. Enabling the sparsity in regression has also attracted a significant amount of work including the LASSO using L_1 penalty [38], the Bridge estimator using L_q ($q > 0$) penalty [23], the SCAD using non-convex penalty [43], the MCP using minimax concave penalty [44] among many others. Several excellent and comprehensive reviews of sparse regression can be found in [5], [19], and [14]. In particular, it is worthy of mentioning that in [46], Zou and Hastie proposed a well-known “elastic net” approach, which integrates the ridge penalty (i.e., squared L_2 penalty) and L_1 penalty into the ordinary least-square objective to obtain a sparse estimator. However, similar to the LASSO method, the elastic net might not consistently find the true sparse solution. On the contrary, instead, we introduce a constraint $\|\beta\|_0 \leq k$ in (F0), which strictly enforces the sparsity on β , and therefore, can obtain a best k -sparse estimator.

Note that it has been proven that the exact sparse linear regression (i.e., using L_0 norm) is NP-hard (cf., [30]), so is the sparse ridge regression (F0). There has been various effective approximation algorithms or heuristics introduced to solve sparse regression [11, 12, 16, 24, 25, 29]. For example, in [9], Das and Kempe studied greedy approach (or forward stepwise selection method) and proved its approximation guarantee when the covariance matrix is nearly identity and has constant bandwidth. However, the greedy approach has been found prohibitively expensive when the number of features (i.e., p) becomes large [20]. Recently, Hazimeh and Mazumder in [21] integrated coordinate descent method with local combinatorial search, and reported that the proposed method outperforms the existing ones. However, this method does not provide any provable guarantee on the solution quality. Many researchers have also attempted to solve sparse regression by developing exact algorithms (e.g., branch and cut), or using mixed integer program (MIP) solvers. It has been shown that for certain large-sized instances with large signal-to-noise ratios, the proposed MIP approaches with warm start (a good initial solution) work quite well and can yield very high-quality solutions [3, 4, 28]. In particular, in [4], Bertsimas and Van Parys also studied sparse ridge regression and developed a branch and cut algorithm. However, through our numerical study, these exact approaches can only solve medium-sized instances to the near optimality, and their performances highly rely on the speed of commercial solvers and can vary significantly from one dataset

to another. In this work, our emphasis is to develop fast approximation algorithms with attractive scalability property and theoretical performance guarantees.

Our Approaches and Contributions. In this work, we will focus on studying the sparse ridge regression (F0) from the angle of chance-constrained program (F0-CCP). We will first investigate various existing approaches of CCP to solve (F0-CCP). One particular approach, which has been used to solve sparse regressions [3], is to introduce one binary variable for each indicator function in (F0-CCP) and linearize it with big-M coefficient. Oftentimes, such a method can be very slow in computation, in particular for large-scale datasets. To overcome the aforementioned challenge, we develop a *big-M free* mixed integer second order conic (MISOC) reformulation for (F0-CCP). We further show that its continuous relaxation is equivalent to that of a mixed integer convex (MIC) formulation in [4, 11]. Moreover, these two formulations motivate us to construct a greedy approach (i.e., forward selection) in a much more efficient way than those in the literature. The performance guarantee of our greedy approach is also established. A randomized algorithm is studied by investigating the continuous relaxations of the proposed MISOC formulation. Numerical study shows that the proposed methods work quite well, in particular, the greedy approach outperforms the other methods both in running time and accuracy of variable selection. The contributions are summarized below:

- (i) We investigate theoretical properties of three existing approaches of CCP to solve (F0-CCP), i.e., the big-M method, the conditional-value-at-risk (i.e., CVaR) approach, and the heuristic algorithm from [1], and shed some lights on why those methods are not applicable to solve the sparse ridge regression (F0).
- (ii) We establish a mixed integer second order conic (MISOC) reformulation for (F0-CCP) from perspective formulation [17] and prove its continuous relaxation is equivalent to that of a mixed integer convex formulation in the work [4, 11]. We also show that the proposed MISOC formulation can be stronger than the naive big-M formulation.
- (iii) Based on the reformulations, we develop an efficient greedy approach for solving (F0-CCP), and prove its performance guarantee under a mild condition. From our analysis, the proposed greedy approach is theoretically sound and computationally efficient.
- (iv) Through establishing a relationship between the continuous relaxation value of the MISOC formulation and the optimal value of (F0-CCP) (i.e., v^*), we develop a randomized algorithm based on the optimal continuous relaxation solution of the MISOC formulation, and derive

its theoretical properties. Such a continuous relaxation solution can help reduce the number of potential features and thus can be integrated with greedy approach.

The remainder of the paper is organized as follows. Section 2 investigates the applicability of several existing approaches of CCP to the sparse ridge regression (F0). Section 3 develops two big-M free mixed integer convex program formulations and proves their equivalence. Section 4 proposes and analyzes two scalable algorithms and proves their performance guarantees. The numerical studies of the proposed scalable algorithms are presented in Section 5. We conclude this work with some discussion in Section 6.

The following notation is used throughout the paper. We use bold-letters (e.g., \mathbf{x} , \mathbf{A}) to denote vectors or matrices, and use corresponding non-bold letters to denote their components. Given a positive integer t , we let $[t] = \{1, \dots, t\}$ and let \mathbf{I}_t denote $t \times t$ identity matrix. Given a subset $S \subseteq [p]$, we let $\boldsymbol{\beta}_S$ denote the subvector of $\boldsymbol{\beta}$ with entries from subset S , and \mathbf{X}_S be a matrix with a subset S columns from \mathbf{X} . For a matrix \mathbf{Y} , we let $\sigma_{\min}(\mathbf{Y}), \sigma_{\max}(\mathbf{Y})$ denote its smallest and largest singular values, respectively. Given a vector \mathbf{x} , we let $\text{diag}(\mathbf{x})$ be a diagonal matrix with diagonal entries from \mathbf{x} . For a matrix \mathbf{W} , we let $\mathbf{W}_{\bullet i}$ denotes its i th column. Given a set T , we let $\text{conv}(T)$ denote its convex hull. Given a finite set S , we let $|S|$ denote its cardinality. Given two sets S, T , we let $S \setminus T$ denote the set of elements in S but not in T , let $S \cup T$ denote the union of S and T and let $S \Delta T$ be their symmetric difference, i.e., $S \Delta T = (S \setminus T) \cup (T \setminus S)$.

2. Existing Solution Approaches. In this section, we investigate three commonly used solution approaches to solve (F0-CCP).

2.1. *Big-M Method.* One typical method for a CCP is to formulate it as a mixed integer program (MIP) by introducing a binary variable z_i for each scenario $i \in [p]$, i.e., $\mathbb{I}(\beta_i \neq 0) \leq z_i$, and then using big-M method to linearize it, i.e., suppose that $|\beta_i| \leq M_i$ with a large positive number M_i , then $z_i \geq \mathbb{I}(\beta_i \neq 0)$ is equivalent to $|\beta_i| \leq M_i z_i$. Therefore, (F0-CCP) can be reformulated as the following MIP:

$$\text{(F0-big-M)} \quad v^* = \min_{\boldsymbol{\beta}, \mathbf{z}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 : \sum_{i \in [p]} z_i \leq k, |\beta_i| \leq M_i z_i, \mathbf{z} \in \{0, 1\}^n \right\}.$$

The advantage of (F0-big-M) is that it can be directly solved by off-the-shelf solvers (e.g., CPLEX, Gurobi). However, one has to choose the vector

$\mathbf{M} = (M_1, \dots, M_p)^\top$ properly.

There are many ways to choose the big-M coefficients (i.e., $\{M_i\}_{i \in [p]}$). One typical way is that for each $i \in [p]$, one can let M_i be equal to the largest value of $|\beta_i|$ given that the optimal value v^* of (F0-CCP) is bounded by v^U , i.e., let M_i be equal to the larger optimal value of the following two convex quadratic programs

$$(1) \quad \max \left\{ \max_{\boldsymbol{\beta}} \left\{ \beta_i : \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U \right\}, \right. \\ \left. \max_{\boldsymbol{\beta}} \left\{ -\beta_i : \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U \right\} \right\}.$$

To solve (1), it needs to compute an upper bound v^U of v^* . Note that the objective value of any feasible solution to (F0-CCP) suffices. A naive upper bound is $v^U = \|\mathbf{y}\|_2^2$ since $\boldsymbol{\beta} = 0$ is feasible to (F0-CCP). On the other hand, to obtain vector \mathbf{M} , one has to solve two convex quadratic programs in (1) for each $i \in [p]$, which can be very time-consuming, in particular when p is large.

Here we derive a slightly weaker but a closed-form vector \mathbf{M} . Note that all the convex programs in (1) share the same constraint $\frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U$. Thus, the key proof idea is to relax the constraint in (1) into a weaker one, which is more amenable for a closed-form upper bound of vector $|\boldsymbol{\beta}|$. This result is summarized below.

PROPOSITION 1. *Suppose that v^U is an upper bound to v^* , then vector $\mathbf{M} = (M_1, \dots, M_p)^\top$ can be chosen as*

$$(2) \quad M_i = \min \left\{ \sqrt{\frac{\frac{1}{n} \|\mathbf{y}\|_2^2 + v^U}{\frac{1}{2n} \sigma_{\min}(\mathbf{X}^\top \mathbf{X}) + \lambda}}, \sqrt{\frac{v^U}{\lambda}} \right\}, \forall i \in [p].$$

PROOF. In (1), it is sufficient to find an upper bound of vector $|\boldsymbol{\beta}|$ for any feasible $\boldsymbol{\beta}$ satisfies

$$\frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U.$$

First, the above constraint implies that

$$\frac{1}{2n} \|\mathbf{X}\boldsymbol{\beta}\|_2^2 - \frac{1}{n} \|\mathbf{y}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U,$$

where the first inequality is due to the inequality $(a - b)^2 \geq \frac{1}{2}b^2 - a^2$. Therefore, we have

$$\left(\frac{1}{2n} \sigma_{\min}(\mathbf{X}^\top \mathbf{X}) + \lambda \right) \|\boldsymbol{\beta}\|_2^2 \leq \boldsymbol{\beta}^\top \left(\frac{1}{2n} \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}_p \right) \boldsymbol{\beta} \leq \frac{1}{n} \|\mathbf{y}\|_2^2 + v^U,$$

where the first inequality is due to $\mathbf{X}^\top \mathbf{X} \geq \sigma_{\min}(\mathbf{X}^\top \mathbf{X}) \mathbf{I}_p$. Since $|\beta_i|^2 \leq \|\boldsymbol{\beta}\|_2^2$, thus we have

$$|\beta_i| \leq \sqrt{\frac{\frac{1}{n} \|\mathbf{y}\|_2^2 + v^U}{\frac{1}{2n} \sigma_{\min}(\mathbf{X}^\top \mathbf{X}) + \lambda}},$$

for each $i \in [p]$.

On the other hand, we note that

$$v^U \geq \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 \geq \lambda \|\boldsymbol{\beta}\|_2^2.$$

Thus, another upper bound can be developed by letting $\lambda \beta_i^2 \leq \lambda \|\boldsymbol{\beta}\|_2^2 \leq v^U$, which implies that $|\beta_i| \leq \sqrt{\frac{v^U}{\lambda}}$ for each $i \in [p]$. \square

It is known that this MIP (**F0-big-M**) with big-M coefficients typically has a very weak continuous relaxation value. Consequently, there has been significant research on improving the big-M coefficients of (**F0-big-M**), for example, [1, 3, 32, 33, 37]. However, the tightening procedures tend to be time consuming in particular for large-scale datasets. In Section 3, we will derive two big-M free MIP formulations, whose continuous relaxation can be proven to be stronger than that of (**F0-big-M**).

2.2. CVaR Approximation. Another well-known approximation of CCP is the so-called conditional value at risk (**CVaR**) approximation (see [31] for details), which is to replace the nonconvex probabilistic constraint by a convex **CVaR** constraint. For the sparse ridge regression in (**F0-CCP**), the resulting formulation is

$$(3) \quad v^{\mathbf{CVaR}} = \min_{\boldsymbol{\beta}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 : \inf_t \left[-\frac{k}{p}t + \frac{1}{p} \sum_{i \in [p]} (|\beta_i| + t)_+ \right] \leq 0 \right\},$$

where $(w)_+ = \max(w, 0)$. It is seen that (3) is a convex optimization problem and provides a feasible solution to (**F0-CCP**). Thus $v^{\mathbf{CVaR}} \geq v^*$. However, we observe that the only feasible solution to (3) is $\boldsymbol{\beta} = 0$.

PROPOSITION 2. *The only feasible solution to (3) is $\boldsymbol{\beta} = 0$, i.e., $v^{\mathbf{CVaR}} = \frac{1}{n} \|\mathbf{y}\|_2^2$.*

PROOF. We first observe that the infimum in (3) must be achievable. Indeed, $h(t) := -\frac{k}{p}t + \frac{1}{p} \sum_{i \in [p]} (|\beta_i| + t)_+$ is continuous and convex in t , and

$\lim_{t \rightarrow \infty} h(t) = \infty$ and $\lim_{t \rightarrow -\infty} h(t) = \infty$. Therefore, the infimum in (3) must exist. Hence, in (3), we can replace the infimum by existing operator as below:

$$v^{\text{CVaR}} = \min_{\beta} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_2^2 : \exists t, -\frac{k}{p}t + \frac{1}{p} \sum_{i \in [p]} (|\beta_i| + t)_+ \leq 0 \right\}.$$

Since $\frac{1}{p} \sum_{i \in [p]} (|\beta_i| + t)_+ \geq 0$ and $\frac{k}{p} > 0$, therefore, $t \geq 0$, i.e.

$$v^{\text{CVaR}} = \min_{\beta} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_2^2 : \exists t \geq 0, \frac{p-k}{p}t + \frac{1}{p} \sum_{i \in [p]} |\beta_i| \leq 0 \right\},$$

which implies that $t = 0$ and $\beta_i = 0$ for each $i \in [p]$. \square

Therefore, the **CVaR** approach yields a trivial solution for (F0-CCP). Hence, it is not a desirable approach and other alternatives are more preferred.

2.3. Heuristic Algorithm in [1]. In the recent work of [1], Ahmed et al. proposed a heuristic algorithm for a CCP with discrete distribution. It was reported that such a method can solve most of their numerical instances to the near-optimality (i.e., within 4% optimality gap). The key idea of the heuristic algorithm in [1] is to minimize the sum of infeasibilities for all scenarios when the objective value is upper bounded by v^U . Specifically, they considered the following optimization problem

$$(4) \quad \min_{\beta} \left\{ \sum_{i \in [p]} |\beta_i| : \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\beta\|_2^2 \leq v^U \right\}.$$

Let β_U^* be an optimal solution to (4) given an upper bound v^U of v^* . The heuristic algorithm is to decrease the value of v^U if $\|\beta_U^*\|_0 \leq k$, and increase it, otherwise. This searching (i.e., bisection) procedure will terminate after a finite number of iterations. The detailed procedure is described in Algorithm 1. Let v^{huer} denote the output solution from Algorithm 1. Then clearly,

PROPOSITION 3. *For Algorithm 1, the following two properties hold:*

- (i) *It terminates with at most $\lfloor \log_2(\frac{\|\mathbf{y}\|_2^2}{n\delta}) \rfloor + 1$ iterations; and*
- (ii) *It generates a feasible solution to (F0-CCP), i.e., $v^* \leq v^{\text{huer}}$.*

PROOF. (i) To prove the first part, according to the description of Algorithm 1, it will terminate if and only if $U - L \leq \hat{\delta}$. And after one iteration, the difference between U and L is halved. Thus, suppose Algorithm 1 will terminate with at most T steps, then we must have

$$\frac{\|\mathbf{y}\|_2^2}{n2^{T-1}} > \hat{\delta},$$

$$\text{i.e., } T < 1 + \log_2 \left(\frac{\|\mathbf{y}\|_2^2}{n\hat{\delta}} \right).$$

(ii) For the second part, we start with a feasible solution $\beta = 0$ to (F0-CCP). Thus, clearly, in Algorithm 1, we keep track of the feasible solutions from iteration to iteration. Thus, the output of Algorithm 1 is feasible to (F0-CCP), i.e., $v^* \leq v^{\text{huer}}$. □

Algorithm 1 Heuristic Algorithm in [1]

```

1: Let  $L = 0$  and  $U = \frac{\|\mathbf{y}\|_2^2}{n}$  be known lower and upper bounds for (F0-CCP), let  $\hat{\delta} > 0$ 
   be the stopping tolerance parameter.
2: while  $U - L > \hat{\delta}$  do
3:    $q \leftarrow (L + U)/2$ .
4:   Let  $\hat{\beta}$  be an optimal solution of (4) and set  $\hat{z}_i = \mathbb{I}(\hat{\beta}_i = 0)$  for all  $i \in [p]$ .
5:   if  $\sum_{i \in [p]} \hat{z}_i \geq p - k$  then
6:      $U \leftarrow q$ .
7:   else
8:      $L \leftarrow q$ .
9:   end if
10: end while
11: Output  $v^{\text{huer}} \leftarrow U$ .

```

It is worth remarking that for any given upper bound v^U , the formulation (4) is similar to the Dantzig selector proposed by [6]. The difference between Algorithm 1 and LASSO is that this iterative procedure simultaneously guarantees the sparsity and reduces the regression error while LASSO seeks a trade-off between the error and L_1 penalty of β . We also note that Algorithm 1 might not be computationally efficient since it requires to solve (4) multiple times. To the best of our knowledge, there is not a known performance guarantee of Algorithm 1.

3. Two Reformulations of Sparse Ridge Regression: Big-M Free.

Note that the Big-M formulation in (F0-big-M) is quite compact since it only involves $2p$ number of variables (i.e., β, \mathbf{z}). However, it is usually a weak formulation in the sense that the continuous relaxation value of (F0-big-M)

can be quite faraway from the optimal value v^* . In this section, we propose two big-M free reformulations of (F0-CCP) from the distinct perspectives and prove their equivalence.

3.1. Mixed Integer Second Order Conic (MISOC) Formulation. In this subsection, we will derive an MISOC formulation and its analytical properties. To begin with, we first make an observation from the perspective formulation in [17]. Let us consider a nonconvex set

$$(5) \quad W_i := \{(\beta_i, \mu_i, z_i) : \beta_i^2 \leq \mu_i, z_i \geq \mathbb{I}(\beta_i \neq 0), z_i \in \{0, 1\}\},$$

for each $i \in [p]$. The results in [17] shows that the convex hull of W_i , denoted as $\text{conv}(W_i)$, can be characterized as below.

LEMMA 1. (*Lemma 3.1. in [17]*) For each $i \in [p]$, the convex hull of set W_i is

$$(6) \quad \text{conv}(W_i) = \{(\beta_i, \mu_i, z_i) : \beta_i^2 \leq \mu_i z_i, z_i \in [0, 1]\}.$$

Lemma 1 suggests an extended formulation for (F0-CCP) without big-M coefficients. To achieve this goal, we first introduce a variable μ_i to be the upper bound of β_i^2 for each $i \in [p]$ and binary variable $z_i \geq \mathbb{I}(\beta_i \neq 0)$, thus, (F0-CCP) is equal to

$$v^* = \min_{\beta, \mu, z} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : \sum_{i \in [p]} z_i \leq k, (\beta_i, \mu_i, z_i) \in W_i, \forall i \in [p] \right\},$$

which can be equivalently reformulated as

$$(7) \quad \min_{\beta, \mu, z} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : (\beta_i, \mu_i, z_i) \in \text{conv}(W_i), z_i \in \{0, 1\}, \forall i \in [p], \sum_{i \in [p]} z_i \leq k \right\}.$$

Note that (i) in (7), we replace W_i by $\text{conv}(W_i)$ and enforce z_i to be binary for each $i \in [p]$; and (ii) from Lemma 1, $\text{conv}(W_i)$ can be described by (6).

The above result is summarized in the following theorem.

THEOREM 1. *The formulation (F0-CCP) is equivalent to*

$$(F0-MISOC) \quad v^* = \min_{\beta, \mu, z} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : \sum_{i \in [p]} z_i \leq k, \beta_i^2 \leq \mu_i z_i, z_i \in \{0, 1\}, \forall i \in [p] \right\}.$$

This formulation (**F0-MISOC**) introduces p more variables $\{\mu_i\}_{i \in [p]}$ than (**F0-big-M**), but it does not require any big-M coefficients.

Next, we show that the convex hull of the feasible region of (**F0-MISOC**) is equal to that of its continuous relaxation. Therefore, it suggests that we might not be able to improve the formulation by simply exploring the constraint system of (**F0-MISOC**). For notational convenience, let us denote T as the feasible region of (**F0-MISOC**), i.e.,

$$(8) \quad T = \left\{ (\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}) : \sum_{i \in [p]} z_i \leq k, \beta_i^2 \leq \mu_i z_i, z_i \in \{0, 1\}, \forall i \in [p] \right\}.$$

We show that the continuous relaxation of the set T is equivalent to $\text{conv}(T)$, i.e.,

PROPOSITION 4. *Let T denote as the feasible region of (**F0-MISOC**). Then*

$$\text{conv}(T) = \left\{ (\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}) : \sum_{i \in [p]} z_i \leq k, \beta_i^2 \leq \mu_i z_i, z_i \in [0, 1], \forall i \in [p] \right\}.$$

PROOF. Let \widehat{T} be the continuous relaxation set of T , i.e.,

$$\widehat{T} = \left\{ (\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}) : \sum_{i \in [p]} z_i \leq k, \beta_i^2 \leq \mu_i z_i, z_i \in [0, 1], \forall i \in [p] \right\}.$$

We would like to show that $\text{conv}(T) = \widehat{T}$.

- (i) It is clear that $\text{conv}(T) \subseteq \widehat{T}$.
- (ii) To prove $\widehat{T} \subseteq \text{conv}(T)$, we only need to show that for any given point $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\mu}}, \widehat{\mathbf{z}}) \in \widehat{T}$, we have $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\mu}}, \widehat{\mathbf{z}}) \in \text{conv}(T)$. Since $\widehat{\mathbf{z}} \in \{\mathbf{z} : \sum_{i \in [p]} z_i \leq k, \mathbf{z} \in [0, 1]^p\}$, which is an integral polytope, there exists K integral extreme points $\{\widehat{\mathbf{z}}^t\}_{t \in [K]} \subseteq \mathbb{Z}_+^p$ such that $\widehat{\mathbf{z}} = \sum_{t \in [K]} \lambda_t \widehat{\mathbf{z}}^t$ with $\lambda_t \in (0, 1)$ for all t and $\sum_{t \in [K]} \lambda_t = 1$. Now we construct $(\widehat{\boldsymbol{\beta}}^t, \widehat{\boldsymbol{\mu}}^t)$ for each $t \in [K]$ as follows:

$$\widehat{\boldsymbol{\mu}}^t = \begin{cases} \frac{\widehat{\mu}_i}{\widehat{z}_i} & \text{if } \widehat{z}_i^t = 1 \\ 0 & \text{otherwise} \end{cases}, \quad \widehat{\boldsymbol{\beta}}^t = \begin{cases} \frac{\widehat{\beta}_i}{\widehat{z}_i} & \text{if } \widehat{z}_i^t = 1 \\ 0 & \text{otherwise} \end{cases}, \forall i \in [p].$$

First of all, we claim that $(\widehat{\boldsymbol{\beta}}^t, \widehat{\boldsymbol{\mu}}^t, \widehat{\mathbf{z}}^t) \in T$ for all $t \in [K]$. Indeed, for any $t \in [K]$,

$$(\widehat{\beta}_i^t)^2 = \begin{cases} \frac{(\widehat{\beta}_i)^2}{\widehat{z}_i^2} & \text{if } \widehat{z}_i^t = 1 \\ 0 & \text{otherwise} \end{cases} \leq \widehat{\mu}_i^t \widehat{z}_i^t = \begin{cases} \frac{\widehat{\mu}_i}{\widehat{z}_i} & \text{if } \widehat{z}_i^t = 1 \\ 0 & \text{otherwise} \end{cases}, \forall i \in [p]$$

$$\begin{aligned} \sum_{i \in [p]} \bar{z}_i^t &\leq k \\ \bar{\mathbf{z}}^t &\in \{0, 1\}^p. \end{aligned}$$

As $\hat{\mathbf{z}} = \sum_{t \in [K]} \lambda_t \bar{\mathbf{z}}^t$, thus, for each $i \in [p]$, we have

$$\begin{aligned} \sum_{t \in [K]} \lambda_t \bar{\mu}_i^t &= \sum_{t \in [K]} \lambda_t \frac{\hat{\mu}_i}{\bar{z}_i^t} \bar{z}_i^t = \hat{\mu}_i \\ \sum_{t \in [K]} \lambda_t \bar{\beta}_i^k &= \sum_{t \in [K]} \lambda_t \frac{\hat{\beta}_i}{\bar{z}_i^t} \bar{z}_i^t = \hat{\beta}_i. \end{aligned}$$

Thus, $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\mu}}, \hat{\mathbf{z}}) \in \text{conv}(T)$.

□

Finally, we remark that if an upper bound \mathbf{M} of $\boldsymbol{\beta}$ is known, then (F0-MISOC) can be further strengthened by adding the constraints $\mu_i \leq M_i^2 z_i$ for each $i \in [p]$. This result is summarized in the following corollary.

PROPOSITION 5. *The formulation (F0-CCP) is equivalent to (F0-MISOC-M)*

$$v^* = \min_{(\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}) \in T} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : \mu_i \leq M_i^2 z_i, \forall i \in [p] \right\}$$

where vector $\mathbf{M} = (M_1, \dots, M_p)^\top$ can be chosen according to Proposition 1 and the set T is defined in (8).

3.2. Mixed Integer Convex (MIC) Formulation. In this subsection, we will derive an equivalent MIC formulation to (F0-CCP). The main idea is to separate the optimization in (F0-CCP) into two steps: (i) first, we optimize over $\boldsymbol{\beta}$ by fixing its nonzero entries with at most k , and (ii) then we select the best subset of nonzero entries with size at most k . After the first step, it turns out that we can arrive at a convex integer program, which is big-M free. We would like to acknowledge that this result has been independently observed by recent work in [4] and [11]. For the completeness of this paper, we also present a different way of proof here.

PROPOSITION 6. *The formulation (F0-CCP) is equivalent to (F0-MIC)*

$$v^* = \min_{\mathbf{z}} \left\{ f(\mathbf{z}) := \lambda \mathbf{y}^\top \left[n\lambda \mathbf{I}_n + \sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} : \sum_{i \in [p]} z_i \leq k, \mathbf{z} \in \{0, 1\}^p \right\}.$$

PROOF. We first reformulate (F0-CCP) as a combinatorial optimization problem. Let $S = \{i \in [p] : \beta_i \neq 0\}$ and we reformulate (F0-CCP) as

$$v^* = \min_{|S| \leq k} \min_{\beta_S} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}_S \beta_S\|_2^2 + \lambda \|\beta_S\|_2^2 \right\},$$

where \mathbf{X}_S is a submatrix of \mathbf{X} with columns from subset S . Note that the inner minimization has closed-form solution $\beta_S^* = \frac{1}{n} \left(\frac{1}{n} \mathbf{X}_S^\top \mathbf{X}_S + \lambda \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \mathbf{y}$. Hence, (F0-CCP) is equivalent to

$$(9) \quad v^* = \min_{|S| \leq k} \frac{1}{n} \mathbf{y}^\top \left[\mathbf{I}_n - \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \right] \mathbf{y}.$$

For any given $S \subseteq [p]$ with $|S| \leq k$, it remains to show that

$$\begin{aligned} & \mathbf{I}_n - \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top = \left[\mathbf{I}_n + \frac{1}{n\lambda} \mathbf{X}_S \mathbf{X}_S^\top \right]^{-1} \\ (\Leftrightarrow) & \left[\mathbf{I}_n + \frac{1}{n\lambda} \mathbf{X}_S \mathbf{X}_S^\top \right] \left[\mathbf{I}_n - \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \right] = \mathbf{I}_n \\ (\Leftrightarrow) & \left[\mathbf{I}_n + \frac{1}{n\lambda} \mathbf{X}_S \mathbf{X}_S^\top \right] \left[\mathbf{I}_n - \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \right] = \mathbf{I}_n \\ (\Leftrightarrow) & \frac{1}{n\lambda} \mathbf{X}_S \mathbf{X}_S^\top = \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \\ & + \frac{1}{n\lambda} \mathbf{X}_S \mathbf{X}_S^\top \mathbf{X}_S \left(\mathbf{X}_S^\top \mathbf{X}_S + \lambda n \mathbf{I}_{|S|} \right)^{-1} \mathbf{X}_S^\top \end{aligned}$$

By letting binary variable $z_i = 1$ if $i \in S$ and 0, otherwise, then the formulation in (9) is equivalent to the following mixed integer convex program

$$\begin{aligned} v^* &= \min_{\mathbf{z} \in \{0,1\}^p} \left\{ \frac{1}{n} \mathbf{y}^\top \left[\mathbf{I}_n + \frac{1}{n\lambda} \sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} : \sum_{i \in [p]} z_i \leq k \right\} \\ &= \min_{\mathbf{z} \in \{0,1\}^p} \left\{ \lambda \mathbf{y}^\top \left[n\lambda \mathbf{I}_n + \sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} : \sum_{i \in [p]} z_i \leq k \right\}. \end{aligned}$$

□

Note that in [4], Bertsimas and Van Parys proposed a branch and cut algorithm to solve (F0-MIC), which was shown to be effective in solving some large-sized instances. In the next subsection, we will show that the continuous relaxation of (F0-MIC) is equivalent to that of (F0-MISOC).

Therefore, it can be more appealing to solve (F0-MISOC) directly by MISOC solvers (e.g., CPLEX, Gurobi). Indeed, we numerically compares the branch and cut algorithm with directly solving (F0-MISOC) in Section 5.

Finally, we remark that given the set of selected features $S \subseteq [p]$, its corresponding estimator $\widehat{\boldsymbol{\beta}}$ can be computed by the following formula:

$$(10) \quad \begin{cases} \widehat{\boldsymbol{\beta}}_S = (\mathbf{X}_S^\top \mathbf{X}_S + n\lambda \mathbf{I}_{|S|})^{-1} \mathbf{X}_S^\top \mathbf{y} \\ \widehat{\beta}_i = 0 \quad \text{if } i \in [p] \setminus S \end{cases},$$

where $\widehat{\boldsymbol{\beta}}_S$ denotes a sub-vector of $\widehat{\boldsymbol{\beta}}$ with entries from subset S .

3.3. Formulation Comparisons. In this subsection, we will focus on comparing (F0-big-M), (F0-MISOC), (F0-MISOC-M) and (F0-MIC) according to their continuous relaxation bounds. First, let v_1, v_2, v_3, v_4 denote the continuous relaxation of (F0-big-M), (F0-MISOC), (F0-MISOC-M) and (F0-MIC), respectively, i.e.,

(11a)

$$v_1 = \min_{\boldsymbol{\beta}, \mathbf{z}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_2^2 : \sum_{i \in [p]} z_i \leq k, |\beta_i| \leq M_i z_i, \mathbf{z} \in [0, 1]^p \right\},$$

(11b)

$$v_2 = \min_{\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : \beta_i^2 \leq \mu_i z_i, \forall i \in [p], \sum_{i \in [p]} z_i \leq k, \mathbf{z} \in [0, 1]^p \right\},$$

(11c)

$$v_3 = \min_{\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{z}} \left\{ \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\mu}\|_1 : \beta_i^2 \leq \mu_i z_i, \mu_i \leq M_i^2 z_i, \forall i \in [p], \sum_{i \in [p]} z_i \leq k, \mathbf{z} \in [0, 1]^p \right\},$$

(11d)

$$v_4 = \min_{\mathbf{z}} \left\{ f(\mathbf{z}) = \lambda \mathbf{y}^\top \left[n\lambda \mathbf{I}_n + \sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} : \sum_{i \in [p]} z_i \leq k, \mathbf{z} \in [0, 1]^p \right\}.$$

Next, in the following theorem, we will show a comparison of proposed formulations, i.e., (F0-big-M), (F0-MISOC), (F0-MISOC-M) and (F0-MIC). In particular, we prove that $v_2 = v_4$, i.e., the continuous relaxation bounds of (F0-MISOC) and (F0-MIC) coincide. In addition, we show that by adding

big-M constraints $\mu_i \leq M_i^2 z_i$ for each $i \in [p]$ into (F0-MISOC), we arrive at a tighter relaxation bound than that of (F0-big-M), i.e., $v_3 \geq v_1$.

THEOREM 2. *Let v_1, v_2, v_3, v_4 denote optimal values of (11a), (11b), (11c) and (11d), respectively, then*

- (i) $v_2 = v_4 \leq v_3$; and
- (ii) $v_1 \leq v_3$.

PROOF. $v_2 = v_4$. By Lemma A.1. [34], we note that (11c) is equivalent to

$$\begin{aligned} v_4 &= \min_{\gamma_0, \gamma, \mathbf{z}} \lambda \left(\|\gamma_0\|_2^2 + \sum_{i \in [p]} \frac{\gamma_i^2}{z_i} \right), \\ \text{s.t. } & \sqrt{\lambda n} \gamma_0 + \sum_{i \in [p]} \mathbf{x}_i \gamma_i = \mathbf{y}, \\ & \sum_{i \in [p]} z_i \leq k, \\ & \mathbf{z} \in [0, 1]^p, \gamma_0 \in \mathbb{R}^n, \gamma_i \in \mathbb{R}, \forall i \in [p], \end{aligned}$$

where by default, we let $\frac{0}{0} = 0$. Now let $\beta_i = \gamma_i$ and introduce a new variable μ_i to denote $\mu_i \geq \frac{\beta_i^2}{z_i}$ for each $i \in [p]$. Then the above formulation is equivalent to

$$\begin{aligned} v_4 &= \min_{\gamma_0, \beta, \mu, \mathbf{z}} \lambda (\|\gamma_0\|_2^2 + \|\boldsymbol{\mu}\|_1), \\ \text{s.t. } & \sqrt{\lambda n} \gamma_0 + \sum_{i \in [p]} \mathbf{x}_i \beta_i = \mathbf{y}, \\ & \beta_i^2 \leq \mu_i z_i, \forall i \in [p], \\ & \sum_{i \in [p]} z_i \leq k, \\ & \mathbf{z} \in [0, 1]^p, \gamma_0 \in \mathbb{R}^n, \mu_i \in \mathbb{R}_+, \forall i \in [p]. \end{aligned}$$

Finally, in the above formulation, replace $\gamma_0 = \frac{1}{\sqrt{\lambda n}} \left(\mathbf{y} - \sum_{i \in [p]} \mathbf{x}_i \beta_i \right) = \frac{1}{\sqrt{\lambda n}} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})$. Then we arrive at (11b).

$v_2 \leq v_3$. Note that the set of the constraints in (11b) is a subset of those in (11c). Thus, $v_2 \leq v_3$.

$v_1 \leq v_3$. Let $(\boldsymbol{\beta}^*, \boldsymbol{\mu}^*, \mathbf{z}^*)$ be optimal solution of (11c). Clearly, we must have $\mu_i^* = \frac{(\beta_i^*)^2}{z_i^*}$ for each $i \in [p]$, otherwise, suppose that there exists $i_0 \in [p]$

such that $\mu_{i_0}^* > \frac{(\beta_{i_0}^*)^2}{z_{i_0}^*}$, then the objective value of (11c) can be strictly less than v_2 by letting $\boldsymbol{\beta} = \boldsymbol{\beta}^*$ and $\mu_i = \mu_i^*$ for each $i \neq i_0$ and $\mu_{i_0} = \frac{(\beta_{i_0}^*)^2}{z_{i_0}^*}$, a contradiction of the optimality of $(\boldsymbol{\beta}^*, \boldsymbol{\mu}^*, \mathbf{z}^*)$.

Therefore, we have

$$\begin{aligned} \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 + \lambda \|\boldsymbol{\beta}^*\|_2^2 &\leq \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 + \lambda \sum_{i \in [p]} \frac{(\beta_i^*)^2}{z_i^*} \\ &= \frac{1}{n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^*\|_2^2 + \lambda \|\boldsymbol{\mu}^*\|_1 = v_3, \end{aligned}$$

where the inequality is because $z_i^* \in [0, 1]$ for each $i \in [p]$.

It remains to show that $|\beta_i^*| \leq M_i z_i^*$ for each $i \in [p]$. This indeed holds because of $(\beta_i^*)^2 \leq \mu_i^* z_i^*$ and $\mu_i^* \leq M_i^2 z_i^*$. Hence, $(\boldsymbol{\beta}^*, \mathbf{z}^*)$ is feasible to (11a) with a smaller objective value. Thus, $v_1 \leq v_3$. \square

Based on the results established in Theorem 2, we could directly solve the second order conic program (11b) to obtain the continuous relaxation of MIC (**F0-MIC**), which can be solved quite efficiently by existing solvers (e.g., CPLEX, Gurobi). In addition, adding big-M constraints $\mu_i \leq M_i^2 z_i$ for each $i \in [p]$ into (11b), the relaxation bound can be further improved.

Finally, we would like to elaborate that by choosing vector \mathbf{M} differently, continuous relaxation bound v_2 of (**F0-MISOC**) can dominate v_1 , the relaxation bound of (**F0-big-M**) and vice versa.

EXAMPLE 1. Consider the following instance of (**F0-CCP**) with $n = 2, p = 2, k = 1$ and $\mathbf{y} = (1, 1)^\top, \mathbf{X} = \mathbf{I}_2$. Thus, in this case, we have $v^* = \frac{\lambda}{1+2\lambda} + \frac{1}{2}, v_2 = \frac{4\lambda}{1+4\lambda}$. There are two different choices about $\mathbf{M} = (M_1, M_2)^\top$:

- (i) If we choose \mathbf{M} loosely, i.e., $M_1 = M_2 = \sqrt{\frac{\|\mathbf{y}\|_2^2}{n\lambda}} = \sqrt{\frac{1}{\lambda}}$, then

$$v_1 = \frac{2\lambda}{1+2\lambda} < v_2 < v^*,$$

given that $\lambda > 0$.

- (ii) If we choose \mathbf{M} to be the tightest bound of the optimal solutions of (**F0-CCP**), i.e., $M_1 = M_2 = \frac{1}{1+2\lambda}$, then

$$v_2 < v_1 = \frac{8\lambda + 1}{8\lambda + 4} < v^*,$$

given that $\lambda \in (0, 1/4)$.

4. Two Scalable Algorithms and their Performance Guarantees.

In this section, we will study two scalable algorithms based upon two equivalent formulations (F0-MISOC) and (F0-MIC). That is, the greedy approach based on (F0-MIC), and the randomized algorithm based on (F0-MISOC).

4.1. *The Greedy Approach based on MIC Formulation.* The Greedy approach (i.e., forward selection) has been commonly used for the best subset selection [10, 36, 45]. The idea of the greedy approach is to select a feature which minimizes the marginal decrement of objective value in (F0-MIC) at each iteration until the number of selected features reaches k . Note that given a selected subset $S \subseteq [p]$ and an index $j \notin S$, the marginal objective value difference by adding j to S can be computed explicitly via Sherman-Morrison formula [35] as below:

$$\lambda \mathbf{y}^\top \left[\mathbf{A}_S + \mathbf{x}_j \mathbf{x}_j^\top \right]^{-1} \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} = - \frac{\lambda (\mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{x}_j)^2}{1 + \mathbf{x}_j^\top \mathbf{A}_S^{-1} \mathbf{x}_j},$$

$$\mathbf{A}_S^{-1} = \mathbf{A}_S^{-1} - \frac{\mathbf{A}_S^{-1} \mathbf{x}_j \mathbf{x}_j^\top \mathbf{A}_S^{-1}}{1 + \mathbf{x}_j^\top \mathbf{A}_S^{-1} \mathbf{x}_j},$$

where $\mathbf{A}_S = n\lambda \mathbf{I}_n + \sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^\top$.

This motivates us an efficient implementation of the greedy approach, which is described in Algorithm 2. Note that in Algorithm 2, at each iteration, we only need to keep track of $\{\mathbf{A}_S^{-1} \mathbf{x}_j\}_{j \in [p]}$, $\{\mathbf{x}_j \mathbf{A}_S^{-1} \mathbf{x}_j\}_{j \in [p]}$ and $\{\mathbf{y} \mathbf{A}_S^{-1} \mathbf{x}_j\}_{j \in [p]}$, which has space complexity $O(np)$ and update them from one iteration to another iteration, which costs $O(np)$ operations per iteration. Therefore, the space and time complexity of Algorithm 2 are $O(np)$ and $O(npk)$, respectively.

Algorithm 2 Greedy Approach

- 1: Initialize $S = \emptyset$ and $\mathbf{A}_S = n\lambda \mathbf{I}_n$
 - 2: **for** $i = 1, \dots, k$ **do**
 - 3: Let $j^* \in \arg \min_{j \in [p] \setminus S} \left\{ - \frac{\lambda (\mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{x}_j)^2}{1 + \mathbf{x}_j^\top \mathbf{A}_S^{-1} \mathbf{x}_j} \right\}$
 - 4: Let $S = S \cup \{j^*\}$ and $\mathbf{A}_S = \mathbf{A}_S + \mathbf{x}_{j^*} \mathbf{x}_{j^*}^\top$, $\mathbf{A}_S^{-1} = \mathbf{A}_S^{-1} - \frac{\mathbf{A}_S^{-1} \mathbf{x}_{j^*} \mathbf{x}_{j^*}^\top \mathbf{A}_S^{-1}}{1 + \mathbf{x}_{j^*}^\top \mathbf{A}_S^{-1} \mathbf{x}_{j^*}}$
 - 5: **end for**
 - 6: Output $v^G \leftarrow \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y}$.
-

From our empirical study, the greedy approach works very well. Indeed, next we are going to investigate the greedy solution and to prove that it can

be very close to the true optimal, in particular when $n\lambda$ is not too small. To begin with, let us define θ_s to be the largest singular value of all the matrices $\mathbf{X}_S \mathbf{X}_S^\top$ with $|S| = s$, i.e.,

$$(12) \quad \theta_s := \max_{|S|=s} \sigma_{\max}^2(\mathbf{X}_S) = \max_{|S|=s} \sigma_{\max}(\mathbf{X}_S \mathbf{X}_S^\top),$$

for each $s \in [p]$. By definition (12), we have $1 = \theta_1 \leq \theta_2 \leq \dots \leq \theta_p$, and by default, we let $\theta_0 = 0$.

Note that the well-known restricted isometry property (RIP) in the sparse regression literature [6, 7] states as below:

$$(1 - \delta_s) \|\boldsymbol{\beta}\|_2^2 \leq \|\mathbf{X}\boldsymbol{\beta}\|_2^2 \leq (1 + \delta_s) \|\boldsymbol{\beta}\|_2^2, \forall s \in [p], \boldsymbol{\beta} : \|\boldsymbol{\beta}\|_0 = s,$$

where $\boldsymbol{\delta} \in (0, 1)^p$ is a constant. Then under RIP condition, clearly, vector $\boldsymbol{\theta}$ is upper bounded by $\mathbf{e} + \boldsymbol{\delta}$, i.e., we have $\theta_s \leq (1 + \delta_s)$ for all $s \in [p]$. Therefore, vector $\boldsymbol{\theta}$ usually is quite smaller in particular if the data points $\{\mathbf{x}_i\}_{i \in [p]}$ are nearly uncorrelated.

Our main results of near optimality of the greedy approach are stated as below. That is, if $n\lambda > \theta_k$, then the solution of greedy approach will be quite close to any optimal estimator from (F0-CCP).

THEOREM 3. *Suppose $n\lambda > \theta_k$ with θ defined in (12). Then (i) the output (i.e., v^G) of the greedy approach (i.e., Algorithm 2) is bounded by*

$$(13) \quad v^* \leq v^G \leq v^* + \frac{3\theta_k^2}{n^2\lambda(n\lambda + \theta_k)} \|\mathbf{y}\|_2^2,$$

and (ii) the greedy solution is asymptotically optimal when $\lim_{n \rightarrow \infty} \frac{\|\mathbf{y}\|_2^2}{n} < \infty$ and $\lim_{n \rightarrow \infty} \frac{\theta_k}{n\lambda} = 0$.

PROOF. First of all, for any \mathbf{z} satisfying $\sum_{i \in [p]} z_i \leq k$, we have $\|\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top\|_2 \leq \theta_k < n\lambda$, by Taylor expansion of inverse matrix function, the objective function $f(\mathbf{z})$ is equal to

$$\begin{aligned} f(\mathbf{z}) &= \lambda \mathbf{y}^\top \left[n\lambda \mathbf{I}_n + \sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} \\ &= \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2\lambda} \sum_{i \in [p]} z_i (\mathbf{x}_i^\top \mathbf{y})^2 + \sum_{\tau=0}^{\infty} \frac{(-1)^\tau}{n^{\tau+3} \lambda^{\tau+2}} \mathbf{y}^\top \left(\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right)^{\tau+2} \mathbf{y} \end{aligned}$$

$$\begin{aligned}
&\geq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2 \lambda} \sum_{i \in [p]} z_i (\mathbf{x}_i^\top \mathbf{y})^2 - \|\mathbf{y}\|_2^2 \sum_{\tau=0}^{\infty} \frac{\theta_k^{2\tau+3}}{n^{2\tau+4} \lambda^{2\tau+3}} \\
(14) \quad &= \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2 \lambda} \sum_{i \in [p]} z_i (\mathbf{x}_i^\top \mathbf{y})^2 - \frac{\theta_k^3}{n^2 \lambda (n^2 \lambda^2 - \theta_k^2)} \|\mathbf{y}\|_2^2,
\end{aligned}$$

where the first inequality is because we ignore the negative terms when τ is odd and $\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \preceq \theta_k \mathbf{I}_n$ since $\sigma_{\max} \left(\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right) \leq \theta_k$.

Therefore, if we optimize both sides in (14) over \mathbf{z} subject to $\sum_{i \in [p]} z_i \leq k$, we have

$$(15) \quad v^* \geq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2 \lambda} \sum_{i \in [k]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2 - \frac{\theta_k^3}{n^2 \lambda (n^2 \lambda^2 - \theta_k^2)} \|\mathbf{y}\|_2^2,$$

where \mathcal{P} is a permutation of $[p]$ such that $[p] = \{\mathcal{P}(1), \dots, \mathcal{P}(p)\}$ and $(\mathbf{x}_{\mathcal{P}(1)}^\top \mathbf{y})^2 \geq (\mathbf{x}_{\mathcal{P}(2)}^\top \mathbf{y})^2 \geq \dots \geq (\mathbf{x}_{\mathcal{P}(p)}^\top \mathbf{y})^2$.

On the other hand, according to Step 3 of Algorithm 2, for any given $S, |S| = s < k$, the minimization procedure can be upper bounded as below. Given $\mathbf{A}_S = n\lambda \mathbf{I}_n + \sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^\top$ and $j \in [p] \setminus S$, then we have

$$\begin{aligned}
&\lambda \mathbf{y}^\top \left[\mathbf{A}_S + \mathbf{x}_j \mathbf{x}_j^\top \right]^{-1} \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} = -\frac{\lambda (\mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{x}_j)^2}{1 + \mathbf{x}_j^\top \mathbf{A}_S^{-1} \mathbf{x}_j} \\
&\leq -\left(\mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{x}_j \right)^2 \left(\lambda - \frac{1}{n} \right) \\
&= -\left(\frac{1}{n\lambda} \mathbf{y}^\top \mathbf{x}_j - \frac{1}{n\lambda} \mathbf{y}^\top (I - n\lambda \mathbf{A}_S^{-1}) \mathbf{x}_j \right)^2 \left(\lambda - \frac{1}{n} \right) \\
&\leq -\left(\frac{1}{n\lambda} \left| \mathbf{y}^\top \mathbf{x}_j \right| - \frac{1}{n\lambda} \frac{\theta_s}{n\lambda + \theta_s} \left| \mathbf{y}^\top \mathbf{x}_j \right| \right)^2 \left(\lambda - \frac{1}{n} \right) \\
(16) \quad &= -\frac{1}{(n\lambda + \theta_s)^2} (\mathbf{x}_j^\top \mathbf{y})^2 \left(\lambda - \frac{1}{n} \right),
\end{aligned}$$

where the first inequality is because $\mathbf{A}_S \succeq n\lambda \mathbf{I}_n$, $\|\mathbf{x}_j\|_2 = 1$ and $\frac{1}{1+a} \geq 1-a$ for any $0 \leq a \leq 1$, the second and third inequalities is due to $\sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^\top \preceq \theta_s \mathbf{I}_n$, $n\lambda \geq \theta_k \geq 1$ and $\mathbf{A}_S \leq (n\lambda + \theta_s) \mathbf{I}_n$.

Thus, from (16), we can prove by induction that the greedy value is upper bounded by

$$(17) \quad v^G \leq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{(n\lambda + \theta_k)^2} \left(\lambda - \frac{1}{n} \right) \sum_{i \in [k]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2.$$

Indeed, if $k = 0$, (17) holds. Suppose that $k = t \geq 0$, (17) holds. Now let $k = t + 1$ and let S be the selected subset at iteration t . By induction, we have

$$\lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} \leq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{(n\lambda + \theta_t)^2} \left(\lambda - \frac{1}{n} \right) \sum_{i \in [t]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2.$$

And by the greedy selection procedure, we further have

$$\begin{aligned} v^G &= \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} + \min_{j \in [p] \setminus S} \lambda \mathbf{y}^\top \left[\mathbf{A}_S + \mathbf{x}_j \mathbf{x}_j^\top \right]^{-1} \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} \\ &\leq \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} + \min_{j \in [p] \setminus S} -\frac{1}{(n\lambda + \theta_{t+1})^2} (\mathbf{x}_j^\top \mathbf{y})^2 \left(\lambda - \frac{1}{n} \right) \\ &\leq \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y} - \frac{1}{(n\lambda + \theta_{t+1})^2} (\mathbf{x}_{\mathcal{P}(t+1)}^\top \mathbf{y})^2 \left(\lambda - \frac{1}{n} \right) \\ &\leq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{(n\lambda + \theta_{t+1})^2} \left(\lambda - \frac{1}{n} \right) \sum_{i \in [t+1]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2, \end{aligned}$$

where the first inequality is due to (16), the second inequality is because $n\lambda > 1$, $([p] \setminus S) \cap \{\mathcal{P}(i)\}_{i \in [t+1]} \neq \emptyset$ and $(\mathbf{x}_{\mathcal{P}(1)}^\top \mathbf{y})^2 \geq \dots \geq (\mathbf{x}_{\mathcal{P}(t)}^\top \mathbf{y})^2 \geq (\mathbf{x}_{\mathcal{P}(t+1)}^\top \mathbf{y})^2$, and the third inequality is due to the induction and $\theta_t \leq \theta_{t+1}$.

According to $2n\lambda\theta_k + \theta_k^2 + \lambda n \leq (2\theta_k + 1)(n\lambda + \theta_k)$, (17) is further upper bounded by

$$\begin{aligned} v^G &\leq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2\lambda} \sum_{i \in [k]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2 + \frac{2\theta_k + 1}{n^2\lambda(n\lambda + \theta_k)} \mathbf{y}^\top \left(\sum_{i \in [k]} \mathbf{x}_{\mathcal{P}(i)} \mathbf{x}_{\mathcal{P}(i)}^\top \right) \mathbf{y} \\ &\leq \frac{1}{n} \|\mathbf{y}\|_2^2 - \frac{1}{n^2\lambda} \sum_{i \in [k]} (\mathbf{x}_{\mathcal{P}(i)}^\top \mathbf{y})^2 + \frac{(2\theta_k + 1)\theta_k}{n^2\lambda(n\lambda + \theta_k)} \|\mathbf{y}\|_2^2 \\ &\leq v^* + \left[\frac{(2\theta_k + 1)\theta_k}{n^2\lambda(n\lambda + \theta_k)} + \frac{\theta_k^3}{n^2\lambda(n^2\lambda^2 - \theta_k^2)} \right] \|\mathbf{y}\|_2^2 \\ (18) \quad &\leq v^* + \frac{3\theta_k^2}{n^2\lambda(n\lambda + \theta_k)} \|\mathbf{y}\|_2^2, \end{aligned}$$

where the second inequality is because $\sum_{i \in [k]} \mathbf{x}_{\mathcal{P}(i)} \mathbf{x}_{\mathcal{P}(i)}^\top \leq \theta_k \mathbf{I}_n$, the third inequality is due to (15), and the fourth inequality is because $n\lambda(2\theta_k + 1)\theta_k - \theta_k^3 - \theta_k^2 \leq 3\theta_k^2(n\lambda - \theta_k)$ and $n\lambda > \theta_k \geq 1$.

Thus, if $\lim_{n \rightarrow \infty} \frac{\|\mathbf{y}\|_2^2}{n} < \infty$ and $\lim_{n \rightarrow \infty} \frac{\theta_k}{n\lambda} = 0$, then we have $v^G \xrightarrow{n \rightarrow \infty} v^*$, i.e., the greedy solution is asymptotically optimal. \square

Note that the condition of the greedy approach to be of good quality requires $\lambda > \frac{\theta_k}{n}$. Therefore, a relatively small value of θ_k suffices. For example, in the RIP condition [6], we have $\theta_k \leq 1 + \delta_k \leq 2$, which is indeed a small constant.

In the next subsection, we will derive a randomized algorithm and prove its approximation guarantee under a weaker condition of λ .

In addition, we remark that the estimator β^G of the greedy approach can be computed by (10), where S denotes the set of selected features by greedy approach. In the next theorem, we will show that the derived estimator from greedy approach (i.e., β^G) can be also quite close to an optimal solution β^* of (F0-CCP).

THEOREM 4. *Let β^* be an optimal solution to (F0-CCP) with set of selected features S^* and β^G be the estimator from the greedy approach with set of selected features S^G . If $n\lambda > \theta_k$, then we have*

$$\|\beta^G - \beta^*\|_2 \leq \frac{\sqrt{4n\theta_{|S^G \setminus S^*|} v^*}}{n\lambda + \sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U})} + \sqrt{\frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n\lambda(n\lambda + \theta_k)(n\lambda + \sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U}))}},$$

where $S^U = S^G \cup S^*$, i.e., the union of set S^G and set S^* .

PROOF. Note that the greedy estimator β^G can be computed through (10) by setting S to be S^G , the set of selected features by greedy approach. Moreover, we define $\tilde{\mathbf{X}}$ as follows:

$$\begin{cases} \tilde{\mathbf{X}}_{S^G \setminus S^*} = \mathbf{X}_{S^G \setminus S^*} \\ \tilde{\mathbf{X}}_{\bullet i} = 0 \quad \text{if } i \in [p] \setminus (S^G \setminus S^*) \end{cases}.$$

Then we have,

$$\begin{aligned} & \frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta^G\|_2^2 + \lambda \|\beta^G\|_2^2 - \left[\frac{1}{n} \|\mathbf{y} - \mathbf{X}\beta^*\|_2^2 + \lambda \|\beta^*\|_2^2 \right] \leq \frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n^2 \lambda (n\lambda + \theta_k)} \\ (\Leftrightarrow) & -2(\beta^* - \beta^G)^\top \left[-\frac{1}{n} \mathbf{X}^\top (\mathbf{y} - \mathbf{X}\beta^*) + \lambda \beta^* \right] \\ & + (\beta^* - \beta^G)^\top \left[\frac{1}{n} \mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}_p \right] (\beta^* - \beta^G) \leq \frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n^2 \lambda (n\lambda + \theta_k)} \\ (\Leftrightarrow) & -2(\beta^* - \beta^G)^\top \left[-\frac{1}{n} \tilde{\mathbf{X}}^\top (\mathbf{y} - \mathbf{X}\beta^*) \right] \\ & + (\beta_{S^U}^* - \beta_{S^U}^G)^\top \left[\frac{1}{n} \mathbf{X}_{S^U}^\top \mathbf{X}_{S^U} + \lambda \mathbf{I}_{|S^U|} \right] (\beta_{S^U}^* - \beta_{S^U}^G) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n^2 \lambda (n\lambda + \theta_k)} \\
(\Rightarrow) &-\frac{2}{n} \|\tilde{\mathbf{X}}\|_2 \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^*\|_2 \|\boldsymbol{\beta}_{S^G \setminus S^*}^* - \boldsymbol{\beta}_{S^G \setminus S^*}^G\|_2 \\
&+ \left(\lambda + \frac{\sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U})}{n} \right) \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^G\|_2^2 \leq \frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n^2 \lambda (n\lambda + \theta_k)} \\
(\Rightarrow) &-\sqrt{\frac{4\theta_{|S^G \setminus S^*|} v^*}{n}} \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^G\|_2 + \left(\lambda + \frac{\sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U})}{n} \right) \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^G\|_2^2 \\
&\leq \frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n^2 \lambda (n\lambda + \theta_k)} \\
(\Rightarrow) &\|\boldsymbol{\beta}^G - \boldsymbol{\beta}^*\|_2 \\
&\leq \frac{\sqrt{4n\theta_{|S^G \setminus S^*|} v^*}}{n\lambda + \sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U})} + \sqrt{\frac{3\theta_k^2 \|\mathbf{y}\|_2^2}{n\lambda (n\lambda + \theta_k) (n\lambda + \sigma_{\min}(\mathbf{X}_{S^U}^\top \mathbf{X}_{S^U}))}},
\end{aligned}$$

where the second equivalence is due to the optimality condition of $\boldsymbol{\beta}^*$, i.e., $-\frac{1}{n} \mathbf{X}_{S^*}^\top (\mathbf{y} - \mathbf{X}_{S^*} \boldsymbol{\beta}_{S^*}^*) + \lambda \boldsymbol{\beta}_{S^*}^* = 0$, and the nonzero entries of $\boldsymbol{\beta}^* - \boldsymbol{\beta}^G$ are only from subset $S^U := S^G \cup S^*$. The first implication is due to submultiplicativity of matrix norm and $\|\mathbf{A}\|_2 \geq \sigma_{\min}(\mathbf{A})$, the second implication is because of $\|\tilde{\mathbf{X}}\|_2 \leq \sqrt{\theta_k}$, $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^*\|_2 \leq \sqrt{nv^*}$, and the last implication is because any solution of the following quadratic inequality $at^2 - bt - c \leq 0$ with $a, b, c > 0$ is upper bounded by $\frac{b}{a} + \sqrt{\frac{c}{a}}$. \square

Note that in Theorem 4, the first term of the error bound vanishes when $S^G = S^*$, i.e., when the greedy approach can exactly identify all the features.

4.2. The Randomized Algorithm based on MISOC Formulation. In this subsection, we develop a randomized algorithm based on the continuous relaxation solution of (F0-MISOC), i.e., the optimal solution to (11b), which can be efficiently solved via the interior point method or other convex optimization approaches [2].

Suppose that $\hat{\mathbf{z}}$ is the optimal solution of the continuous relaxation model (11b). For each $i \in [p]$, the column \mathbf{x}_i will be picked by probability \hat{z}_i . The detailed implementation is illustrated in Algorithm 3.

Next, we will show that if λ is not too small, then with high probability, the output S of Algorithm 3 yields its corresponding objective value close to the optimal value v^* . To begin with, we present the following matrix concentration bound.

Algorithm 3 Randomized Algorithm

```

1: Let  $\hat{z}$  be the optimal solution to (11b)
2: Initialize set  $S = \emptyset$  and vector  $\tilde{z} = 0 \in \mathbb{R}^p$ 
3: for  $i = 1, \dots, p$  do
4:   Sample a standard uniform random variable  $U$ 
5:   if  $U \leq \hat{z}_i$  then
6:     Let  $S = S \cup \{i\}$  and  $\tilde{z}_i = 1$ 
7:   end if
8: end for
9: Output  $S, \tilde{z}$ 

```

LEMMA 2. (Theorem 1.4., [39]) Consider a finite sequence $\{\mathbf{Y}_k\}$ of independent, random, symmetric matrices with dimension d . Assume that each random matrix satisfies $E[\mathbf{Y}_k] = 0$ and $\|\mathbf{Y}_k\|_2^2 \leq R^2$ almost surely. Then, for all $t \geq 0$, we have

$$(19) \quad \mathbb{P} \left\{ \left\| \sum_k \mathbf{Y}_k \right\|_2 \geq t \right\} \leq d \exp \left(-\frac{t^2}{2\nu^2 + 2/3Rt} \right),$$

where $\nu^2 := \|\sum_k \mathbb{E}[\mathbf{Y}_k^2]\|_2$.

Lemma 2 implies that if λ is not too small, then with high probability, $\lambda n \mathbf{I}_n + \sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^\top$ has the similar eigenvalues as $\lambda n \mathbf{I}_n + \sum_{i \in [p]} \hat{z}_i \mathbf{x}_i \mathbf{x}_i^\top$, where \hat{z} is the optimal solution to (11b) and S is the output of Algorithm 3.

LEMMA 3. Let \hat{z} be the optimal solution to (11b) and S be the output of Algorithm 3. Given that $\alpha \in (0, 1)$ and

$$\lambda \geq \frac{\log(2n/\alpha)}{3n\epsilon} + \frac{\sqrt{2\theta_k \log(2n/\alpha)}}{2n\epsilon},$$

then with probability at least $1 - \frac{\alpha}{2}$, we have

$$(1 - \epsilon) \mathbf{u}^\top \boldsymbol{\Sigma}_* \mathbf{u} \leq \mathbf{u}^\top \widehat{\boldsymbol{\Sigma}} \mathbf{u} \leq (1 + \epsilon) \mathbf{u}^\top \boldsymbol{\Sigma}_* \mathbf{u}, \forall \mathbf{u} \in \mathbb{R}^n,$$

where $\boldsymbol{\Sigma}_* = \lambda n \mathbf{I}_n + \sum_{i \in [p]} \hat{z}_i \mathbf{x}_i \mathbf{x}_i^\top$ and $\widehat{\boldsymbol{\Sigma}} = \lambda n \mathbf{I}_n + \sum_{i \in S} \mathbf{x}_i \mathbf{x}_i^\top$.

PROOF. Let \hat{z} be the optimal solution to (11b) and let $\{r_i\}_{i \in [p]}$ be independent Bernoulli random variables with $\mathbb{P}\{r_i = 1\} = \hat{z}_i$ for each $i \in [p]$. Consider the random matrix defined as for each $i \in [p]$,

$$\mathbf{A}_i = (r_i - \hat{z}_i) \mathbf{x}_i \mathbf{x}_i^\top$$

and $\mathbb{E}[\mathbf{A}_i] = 0$. On the other hand, by assumption we have $\|\mathbf{x}_i\|_2 = 1$ for each $i \in [p]$, thus

$$\|\mathbf{A}_i\|_2 = |r_i - \widehat{z}_i| \|\mathbf{x}_i\|_2^2 = |r_i - \widehat{z}_i|.$$

Also,

$$\begin{aligned} \left\| \sum_{i \in [p]} \mathbb{E}[\mathbf{A}_i^2] \right\|_2 &= \left\| \sum_{i \in [p]} \widehat{z}_i (1 - \widehat{z}_i) \|\mathbf{x}_i\|_2^2 \mathbf{x}_i \mathbf{x}_i^\top \right\|_2 = \left\| \sum_{i \in [p]} \widehat{z}_i (1 - \widehat{z}_i) \mathbf{x}_i \mathbf{x}_i^\top \right\|_2 \\ &\leq \left\| \sum_{i \in [p]} \widehat{z}_i \mathbf{x}_i \mathbf{x}_i^\top \right\|_2 \leq \theta_k, \end{aligned}$$

where the first inequality is due to triangle inequality and $\|\mathbf{x}_i\|_2^2 = 1$ for each $i \in [p]$, the second inequality is due to $1 - \widehat{z}_i \in [0, 1]$ for all $i \in [p]$ and the last one is due to

$$\begin{aligned} &\max_{\mathbf{z} \in [0,1]^p} \left\{ \sigma_{\max} \left(\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right) : \sum_{i \in [p]} z_i = k \right\} \\ &= \max_{\mathbf{z} \in \{0,1\}^p} \left\{ \sigma_{\max} \left(\sum_{i \in [p]} z_i \mathbf{x}_i \mathbf{x}_i^\top \right) : \sum_{i \in [p]} z_i = k \right\} := \theta_k. \end{aligned}$$

Now by Lemma 2 with $\sigma_{\min}(\boldsymbol{\Sigma}_*)$ denoting the smallest eigenvalue of $\boldsymbol{\Sigma}_*$ and $t = \epsilon \sigma_{\min}(\boldsymbol{\Sigma}_*)$, we have

$$\mathbb{P} \left\{ \left\| \sum_{i \in [p]} (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_*) \right\|_2 \geq \epsilon \sigma_{\min}(\boldsymbol{\Sigma}_*) \right\} \leq n \exp \left(-\frac{\epsilon^2 \sigma_{\min}^2(\boldsymbol{\Sigma}_*)}{2\theta_k + 2/3\epsilon \sigma_{\min}(\boldsymbol{\Sigma}_*)} \right).$$

We would like to ensure that the right-hand side of above inequality is at most $\frac{\alpha}{2}$.

Note that

$$\lambda n \leq \sigma_{\min}(\boldsymbol{\Sigma}_*) \leq \frac{\mathbf{Tr}(\boldsymbol{\Sigma}_*)}{n} = \lambda n + \frac{1}{n} \sum_{i \in [p]} \widehat{z}_i \mathbf{x}_i^\top \mathbf{x}_i = \lambda n + \frac{k}{n},$$

where the first inequality is due to $\boldsymbol{\Sigma}_* \succeq \lambda n \mathbf{I}_n$ and second inequality is because of the well known fact that the smallest eigenvalue of a symmetric matrix is no larger than the arithmetic mean of its trace. Thus, if

$$\mathbb{P} \left\{ \left\| \sum_{i \in [p]} (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_*) \right\|_2 \geq \epsilon \sigma_{\min}(\boldsymbol{\Sigma}_*) \right\} \leq \mathbb{P} \left\{ \left\| \sum_{i \in [p]} (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_*) \right\|_2 \geq \epsilon \lambda n \right\}$$

$$\leq n \exp\left(-\frac{\epsilon^2(\lambda n + \frac{k}{n})^2}{2\theta_k + 2/3\epsilon n\lambda}\right) \leq n \exp\left(-\frac{\epsilon^2\sigma_{\min}^2(\boldsymbol{\Sigma}_*)}{2\theta_k + 2/3\epsilon\sigma_{\min}(\boldsymbol{\Sigma}_*)}\right) \leq \frac{\alpha}{2},$$

i.e., if

$$\lambda \geq \frac{\log(2n/\alpha)}{3n\epsilon} + \frac{\sqrt{2\theta_k \log(2n/\alpha)}}{2n\epsilon},$$

then with probability at least $1 - \frac{\alpha}{2}$, we have

$$\left\| \sum_{i \in [p]} (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_*) \right\|_2 \leq \epsilon\sigma_{\min}(\boldsymbol{\Sigma}_*).$$

Then the conclusion follows directly by Weyl's theorem [15, 42]. \square

Based on Lemma 3, we can imply the following bi-criteria approximation of (F0).

THEOREM 5. *Let $(S, \tilde{\mathbf{z}})$ be the output of Algorithm 3. Given that $\alpha \in (0, 1)$ and*

$$\lambda \geq \frac{\log(2n/\alpha)}{3n\epsilon} + \frac{\sqrt{2\theta_k \log(2n/\alpha)}}{2n\epsilon},$$

then with probability at least $1 - \alpha$, we have

$$(20) \quad \lambda \mathbf{y}^\top \left[\lambda n \mathbf{I}_n + \sum_{i \in [p]} \tilde{z}_i \mathbf{x}_i \mathbf{x}_i^\top \right]^{-1} \mathbf{y} \leq (1 + \epsilon)v^*$$

and

$$(21) \quad \sum_{i \in [p]} \tilde{z}_i \leq \left(1 + \sqrt{\frac{3 \log(2/\alpha)}{k}} \right) k.$$

PROOF. Note that (20) follows from Lemma 3. The result in (21) holds due to the Chernoff bound [8], i.e.,

$$\mathbb{P} \left\{ \sum_{i \in [p]} \tilde{z}_i \leq \left(1 + \sqrt{\frac{3 \log(2/\alpha)}{k}} \right) k \right\} \geq 1 - e^{-\frac{\left(\sqrt{\frac{3 \log(2/\alpha)}{k}} \right)^2 k}{3}} \geq 1 - \frac{\alpha}{2}.$$

Therefore, by union bound or Boole's inequality, we then arrive at the conclusion. \square

Next, let β^R be the estimator from Algorithm 3, which can be computed according to (10) by letting S be the output from Algorithm 3. Then we can show that the distance between β^R and β^* (i.e., $\|\beta^R - \beta^*\|_2$) can be also quite small, where β^* is an optimal solution to (F0).

THEOREM 6. *Let β^* be an optimal solution to (F0) with set of selected features S^* and β^R be the estimator from Algorithm 3 with set of selected features S^R . Given $\alpha \in (0, 1)$, if $\lambda \geq \frac{\log(2n/\alpha)}{3n\epsilon} + \frac{\sqrt{2\theta_k \log(2n/\alpha)}}{2n\epsilon}$, then with probability at least $1 - \alpha$, we have*

$$\|\beta^R - \beta^*\|_2 \leq \frac{\sqrt{4n\theta_{|S^R \setminus S^*|} v^*}}{n\lambda + \sigma_{\min}(\mathbf{X}_{S^R \cup S^*}^\top \mathbf{X}_{S^R \cup S^*})} + \sqrt{\frac{n\epsilon v^*}{n\lambda + \sigma_{\min}(\mathbf{X}_{S^R \cup S^*}^\top \mathbf{X}_{S^R \cup S^*})}}.$$

PROOF. The proof is almost identical to that of Theorem 4, thus is omitted here. \square

Finally, we remark that we can apply the greedy approach based upon the support of continuous relaxation solution of (F0-MISOC). That is, given that $\hat{\mathbf{z}}$ is the optimal solution to (11b) and $\delta > 0$ is a positive constant, then we first let set $\mathcal{C} := \{i \in [p] : \hat{z}_i \geq \delta\}$ and apply greedy approach (Algorithm 2) to set \mathcal{C} rather than $[p]$, which could save a significant amount of computational time, in particular when continuous relaxation solution $\hat{\mathbf{z}}$ is very sparse. The detailed description can be found in Algorithm 4.

Algorithm 4 Restricted Greedy Approach

- 1: Let $\hat{\mathbf{z}}$ be the optimal solution to (11b)
 - 2: Initialize $\delta > 0$ (e.g., $\delta = 0.01$), $\mathcal{C} := \{i \in [p] : \hat{z}_i \geq \delta\}$
 - 3: Let $S = \emptyset$ and $\mathbf{A}_S = n\lambda \mathbf{I}_n$
 - 4: **for** $i = 1, \dots, k$ **do**
 - 5: Let $j^* \in \arg \min_{j \in \mathcal{C} \setminus S} \left\{ -\frac{\lambda(\mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{x}_j)^2}{1 + \mathbf{x}_j^\top \mathbf{A}_S^{-1} \mathbf{x}_j} \right\}$
 - 6: Let $S = S \cup \{j^*\}$ and $\mathbf{A}_S = \mathbf{A}_S + \mathbf{x}_{j^*} \mathbf{x}_{j^*}^\top$, $\mathbf{A}_S^{-1} = \mathbf{A}_S^{-1} - \frac{\mathbf{A}_S^{-1} \mathbf{x}_{j^*} \mathbf{x}_{j^*}^\top \mathbf{A}_S^{-1}}{1 + \mathbf{x}_{j^*}^\top \mathbf{A}_S^{-1} \mathbf{x}_{j^*}}$
 - 7: **end for**
 - 8: Output $v^{RG} \leftarrow \lambda \mathbf{y}^\top \mathbf{A}_S^{-1} \mathbf{y}$.
-

5. Numerical Illustration. In this section, we conduct numerical studies to evaluate the performance of the proposed method. The input and response of the data are generated from the underlying linear model

$$y = \mathbf{x}^\top \beta^0 + \tilde{\epsilon},$$

where $\tilde{\epsilon} \sim N(0, \sigma^2)$. The i.i.d. sample of \mathbf{x} are generated from a multivariate normal distribution with

$$\mathbf{x}_i \sim N(0, \Sigma), \quad i = 1, \dots, n,$$

where Σ is the covariance matrix with its entries to be $\sigma_{ij} = \rho^{|i-j|}$ for each $i, j \in [p]$. Here we choose $\rho = 0.5$. The values of nonzero coefficients in $\beta^0 = (\beta_1^0, \dots, \beta_p^0)^\top$ are drawn randomly from the uniform distribution $Unif(-3, 3)$. To control the signal-to-noise ratio (SNR), we choose the value of σ^2 such that $SNR = var(\mathbf{x}^\top \beta^0) / var(\tilde{\epsilon}) = 9$. By generating an i.i.d. sample of noise $\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n$ with $\tilde{\epsilon}_i \sim N(0, \sigma^2)$ for each $i \in [n]$, we simulate the response values, i.e., $y_i = \mathbf{x}_i^\top \beta^0 + \tilde{\epsilon}_i$ for each $i \in [n]$.

Recall that the goal is to find a best k -sparse estimator for a given k . The performance of the methods in comparison are evaluated by the selection accuracy and computational time. Here we consider different combinations of k, n, p to generate the simulation data, where $p \in \{1000, 5000\}$, $n \in \{500, 1000, 5000\}$ and $k \in \{10, 20, 30\}$. Each simulation setting is repeated by 10 times, i.e., for each tuple (k, n, p) , we generate 10 repetitions¹. For simplicity, for all the testing instances, we set tuning parameter $\lambda = 0.08$.

The methods in comparison include the branch-and-cut algorithm proposed by [4] based on (F0-MIC), directly solving (F0-MISOC), the heuristic Algorithm 1 in [1], the proposed greedy Algorithm 2, the proposed randomized Algorithm 3 and the proposed restricted greedy Algorithm 4). Note that the heuristic Algorithm 1 in [1] is similar to the LASSO on the use of L_1 norm to achieve the sparsity. The commercial solver Gurobi 7.5 with its default setting is used to solve (F0-MISOC) and its continuous relaxation. We set time limit to be an hour (3600 seconds). Due to out-of-memory and out-of-time-limit issues, in the case of $p = 5000$, we only compute two of the most effective algorithms: the proposed greedy Algorithm 2 and the proposed restricted greedy Algorithm 4. The comparison results are listed in Table 1 to Table 3, where the *Avg. Value*, *Avg. Gap*, *Avg. Time*, and *Avg. False Alarm Rate* denotes the average objective function value, average optimality gap (of exact methods), average computational time (in seconds), and average percent of falsely detected features, respectively. All the computations were executed on a MacBook Pro with a 2.80 GHz processor and 16GB RAM.

Table 1 reports the comparison results between directly solving (F0-MISOC) and the branch-and-cut algorithm based upon (F0-MIC). It is seen that that directly solving (F0-MISOC) outperforms the branch-and-cut algorithm for

¹We restrict the simulation to 10 repetitions because certain existing methods is very slow in computation.

TABLE 1
 Comparison of the Branch and Cut algorithm in [4] and directly solving (F0-MISOC)
 with $p = 1000$

p	k	n	Branch and Cut in [4]				Solving (F0-MISOC)			
			Avg. Value	Avg. Time	Avg. Gap	Avg. False Alarm Rate	Avg. Value	Avg. Time	Avg. Gap	Avg. False Alarm Rate
1000	10	500	9.71	3438.51	47.2%	26.0%	6.83	3505.82	7.1%	5.0%
		1000	7.11	2451.47	10.4%	5.0%	7.27	3562.61	9.7%	7.0%
		5000	NA*	NA	NA	NA	6.67	387.44	0.0%	0.0%
	20	500	23.02	3600.00	141.5%	45.0%	11.98	3600.00	21.4%	20.0%
		1000	31.52	3600.00	131.2%	50.5%	11.55	3600.00	11.7%	18.0%
		5000	NA	NA	NA	NA	11.30	2434.64	0.3%	0.5%
	30	500	39.62	3600.00	189.3%	51.3%	20.42	3600.00	31.4%	27.0%
		1000	50.63	3600.00	175.9%	55.0%	19.16	3600.00	18.1%	22.3%
		5000	NA	NA	NA	NA	17.79	3600.00	1.3%	5.0%

* The NA represents for out of memory instances.

most of instances, in particular when k becomes large. This is because that (i) we prove in Theorem 2 that continuous relaxations of (F0-MIC) and (F0-MISOC) are equivalent, thus directly solving (F0-MISOC) should perform at least as good as branch and cut algorithm; and (ii) the branch-and-cut algorithm needs to compute the gradient of the objective function in (F0-MIC), which involves very time-consuming $n \times n$ matrix inversion. However, for both approaches, they reach the time limit for most of the cases and the average false alarm rates are higher than the approximation algorithms in Table 2. Therefore, for large-scale instances, these approaches might not be very desirable.

From Table 2 and Table 1, the proposed greedy Algorithm 2 and restricted greedy Algorithm 4 apparently perform best among all comparison methods. We see that for the instances with $k = 10$, the heuristic Algorithm 1, greedy Algorithm 2 and restricted greedy Algorithm 4 find almost all the features, while the randomized Algorithm 3 performs slightly worse. When the number of active features, k , grows, all the methods in comparison have relatively larger false alarm rates. Their performance of identifying right features improves as the sample size n increases, i.e., providing more information. For the heuristic Algorithm 1 in [1], it is less accurate and takes much longer time. Thus, it might not be a good option for the large-scale instances either. In contrast, we note that the greedy Algorithm 2 is much accurate. It runs very fast with the computation time increasing proportionally to n, p, k . But the randomized Algorithm 3, which depends on solution time of solving continuous relaxation of (F0-MISOC), is quite insensitive to k in terms of computation time. Therefore, by integrating these two together, the restricted greedy Algorithm 4 can be advantageous for large

TABLE 2
 Comparison of Heuristic Algorithm 1 in [1], Greedy Algorithm 2, Randomized Algorithm 3 and Restricted Greedy Algorithm 4 with $p = 1000$

p	k	n	Heuristic Algorithm 1 in [1]			Proposed Greedy Algorithm 2		
			Avg. Value	Avg. Time	Avg. False Alarm	Avg. Value	Avg. Time	Avg. False Alarm Rate
1000	10	500	9.59	579.36	3.0%	6.60	0.47	0.0%
		1000	7.88	45.78	0.0%	6.54	0.59	0.0%
		5000	7.24	737.06	0.0%	6.67	1.41	0.0%
	20	500	15.87	589.66	14.5%	10.86	0.79	9.0%
		1000	13.42	47.92	11.5%	10.91	2.02	4.0%
		5000	12.66	738.55	4.5%	11.30	2.37	0.0%
	30	500	28.87	583.98	17.0%	16.88	1.13	10.7%
		1000	23.53	43.92	12.7%	17.19	1.43	6.7%
		5000	19.74	678.10	6.0%	17.74	3.28	2.0%
p	k	n	Proposed Randomized Algorithm 3			Proposed Restricted Greedy Algorithm 4		
			Avg. Value	Avg. Time	Avg. False Alarm Rate	Avg. Value	Avg. Time	Avg. False Alarm Rate
1000	10	500	7.79	4.06	14.0%	6.60	3.84	0.0%
		1000	6.86	11.21	6.0%	6.54	10.58	0.0%
		5000	6.67	181.77	0.0%	6.67	186.81	0.0%
	20	500	12.88	4.01	23.5%	10.86	3.80	9.0%
		1000	11.68	10.84	18.0%	10.91	13.81	4.0%
		5000	11.40	199.31	6.5%	11.30	202.66	0.0%
	30	500	20.89	4.21	26.3%	16.89	4.06	11.0%
		1000	19.89	10.58	24.0%	17.19	11.94	6.7%
		5000	18.11	167.95	10.0%	17.74	170.14	2.0%

k , providing accurate estimation with fast computation. For the numerical study with $p = 5000$ below, we choose these two most efficient algorithms for comparison.

In Table 3, we observe that the greedy Algorithm 2 and the restricted greedy Algorithm 4 have exactly the same false alarm rates. But the greedy Algorithm 2 is much faster than the restricted greedy Algorithm 4. This is mainly because it takes much longer time to solve the continuous relaxation to the optimality and for these instances, k is relatively small. In particular, for a large-scale datasets (e.g., $n = p = 5000$), the computation time of the restricted greedy Algorithm 4 is much longer time than those in the case with $p = 1000$. But, the greedy Algorithm 2 can still find very high-quality solutions within 30 seconds of computation time. On the other hand, we note that the accuracy of both approaches grows when the sample size increases. Thus, we would recommend to find a reasonable sample size that the greedy methods can work efficiently and identify the features accurately.

TABLE 3
 Comparison of Greedy Algorithm 2 and Restricted Greedy Algorithm 4 with $p = 5000$

p	k	n	Proposed Greedy Algorithm 2			Proposed Restricted Greedy Algorithm 4		
			Avg. Value	Avg. Time	Avg. False Alarm Rate	Avg. Value	Avg. Time	Avg. False Alarm Rate
5000	10	500	4.57	2.31	0.0%	4.57	15.81	0.0%
		1000	4.59	3.13	0.0%	4.59	39.06	0.0%
		5000	4.68	9.04	0.0%	4.68	1451.78	0.0%
	20	500	12.86	4.31	8.0%	12.86	15.69	8.0%
		1000	13.35	5.41	2.5%	13.35	38.14	2.5%
		5000	13.27	14.58	0.0%	13.27	1426.93	0.0%
	30	500	14.02	5.98	20.7%	14.02	16.24	20.7%
		1000	14.97	8.21	12.7%	14.97	39.41	12.7%
		5000	15.60	20.52	3.3%	15.60	1503.48	3.3%

6. Conclusion. This paper studies the sparse ridge regression with the use of exact L_0 norm for sparsity. We first show that many existing approaches cannot work well for this problem. Then we propose a mixed integer second order conic (MISOC) formulation, which is big-M free and is derived based on perspective formulation. We prove that the continuous relaxation of this MISOC reformulation is equivalent to the convex integer program (CIP) formulation studied by literature, and can be stronger than straightforward big-M formulation. Based on these two formulations, we propose two scalable algorithms, the greedy and randomized algorithms, for solving the sparse ridge regression. Under mild conditions, both algorithm can find near-optimal solutions with performance grantees. Finally, we conduct a series of numerical illustrations and show that greedy algorithm works the best among the other algorithms in comparison.

References.

- [1] Ahmed, S., Luedtke, J., Song, Y., and Xie, W. (2016). Nonanticipative duality and mixed-integer programming formulations for chance-constrained stochastic programs. To appear in *Mathematical Programming*.
- [2] Ben-Tal, A. and Nemirovski, A. (2001). *Lectures on Modern Convex Optimization*. Society for Industrial and Applied Mathematics.
- [3] Bertsimas, D., King, A., Mazumder, R., et al. (2016). Best subset selection via a modern optimization lens. *The Annals of Statistics*, 44(2):813–852.
- [4] Bertsimas, D. and Van Parys, B. (2017). Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *arXiv preprint arXiv:1709.10029*.
- [5] Bühlmann, P. and Van De Geer, S. (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media.
- [6] Candes, E. and Tao, T. (2007). The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351.
- [7] Candes, E. J. (2008). The restricted isometry property and its implications for compressed sensing. *Comptes rendus mathématique*, 346(9-10):589–592.

- [8] Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507.
- [9] Das, A. and Kempe, D. (2008). Algorithms for subset selection in linear regression. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 45–54. ACM.
- [10] Das, A. and Kempe, D. (2011). Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *arXiv preprint arXiv:1102.3975*.
- [11] Dong, H. (2016). On the exact recovery of sparse signals via conic relaxations. *arXiv preprint arXiv:1603.04572*.
- [12] Dong, H., Chen, K., and Linderoth, J. (2015). Regularization vs. relaxation: A conic optimization perspective of statistical variable selection. *arXiv preprint arXiv:1510.06083*.
- [13] Draper, N. R. and Van Nostrand, R. C. (1979). Ridge regression and james-stein estimation: review and comments. *Technometrics*, 21(4):451–466.
- [14] Fan, J. and Lv, J. (2011). Nonconcave penalized likelihood with np-dimensionality. *IEEE Transactions on Information Theory*, 57(8):5467–5484.
- [15] Franklin, J. N. (1968). *Matrix theory*. Courier Corporation.
- [16] Friedman, J. H. (2012). Fast sparse regression and classification. *International Journal of Forecasting*, 28(3):722–738.
- [17] Günlük, O. and Linderoth, J. (2012). Perspective reformulation and applications. *Mixed Integer Nonlinear Programming*, pages 61–89.
- [18] Gusnanto, A. and Pawitan, Y. (2015). Sparse alternatives to ridge regression: a random effects approach. *Journal of Applied Statistics*, 42(1):12–26.
- [19] Hastie, T., Tibshirani, R., and Friedman, J. (2009). Unsupervised learning. In *The elements of statistical learning*, pages 485–585. Springer.
- [20] Hastie, T., Tibshirani, R., and Tibshirani, R. J. (2017). Extended comparisons of best subset selection, forward stepwise selection, and the lasso. *arXiv preprint arXiv:1707.08692*.
- [21] Hazimeh, H. and Mazumder, R. (2018). Fast best subset selection: Coordinate descent and local combinatorial optimization algorithms. *arXiv preprint arXiv:1803.01454*.
- [22] Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.
- [23] Huang, J., Horowitz, J. L., and Ma, S. (2008). Asymptotic properties of bridge estimators in sparse high-dimensional regression models. *The Annals of Statistics*, 36(2):587–613.
- [24] Kowalski, M. (2009). Sparse regression using mixed norms. *Applied and Computational Harmonic Analysis*, 27(3):303–324.
- [25] Kuo, L. and Mallick, B. (1998). Variable selection for regression models. *Sankhyā: The Indian Journal of Statistics, Series B*, pages 65–81.
- [26] Luedtke, J. (2014). A branch-and-cut decomposition algorithm for solving chance-constrained mathematical programs with finite support. *Mathematical Programming*, 146(1-2):219–244.
- [27] Marquardt, D. W. and Snee, R. D. (1975). Ridge regression in practice. *The American Statistician*, 29(1):3–20.
- [28] Mazumder, R. and Radchenko, P. (2017). The discrete dantzig selector: Estimating sparse linear models via mixed integer linear optimization. *IEEE Transactions on Information Theory*, 63(5):3053–3075.
- [29] Miller, A. (2002). *Subset selection in regression*. CRC Press.
- [30] Natarajan, B. K. (1995). Sparse approximate solutions to linear systems. *SIAM*

- journal on computing*, 24(2):227–234.
- [31] Nemirovski, A. and Shapiro, A. (2006). Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17(4):969–996.
 - [32] Pavlikov, K., Veremyev, A., and Pasiliao, E. L. (2018). Optimization of value-at-risk: computational aspects of mip formulations. *Journal of the Operational Research Society*, 69:127–141.
 - [33] Qiu, F., Ahmed, S., Dey, S. S., and Wolsey, L. A. (2014). Covering linear programming with violations. *INFORMS Journal on Computing*, 26(3):531–546.
 - [34] Sagnol, G., Harman, R., et al. (2015). Computing exact d -optimal designs by mixed integer second-order cone programming. *The Annals of Statistics*, 43(5):2198–2224.
 - [35] Sherman, J. and Morrison, W. J. (1950). Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127.
 - [36] Smola, A. J. and Bartlett, P. L. (2001). Sparse greedy gaussian process regression. In *Advances in neural information processing systems*, pages 619–625.
 - [37] Song, Y., Luedtke, J. R., and Küçükyavuz, S. (2014). Chance-constrained binary packing problems. *INFORMS Journal on Computing*.
 - [38] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288.
 - [39] Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434.
 - [40] Tropp, J. A. and Gilbert, A. C. (2007). Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on information theory*, 53(12):4655–4666.
 - [41] Vinod, H. D. (1978). A survey of ridge regression and related techniques for improvements over ordinary least squares. *The Review of Economics and Statistics*, pages 121–131.
 - [42] Weyl, H. (1912). The asymptotic distribution law of the eigenvalues of linear partial differential equations (with an application to the theory of cavity radiation). *Mathematical Annals*, 71(4):441–479.
 - [43] Xie, H. and Huang, J. (2009). Scad-penalized regression in high-dimensional partially linear models. *The Annals of Statistics*, 37(2):673–696.
 - [44] Zhang, C.-H. et al. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics*, 38(2):894–942.
 - [45] Zhang, T. (2011). Adaptive forward-backward greedy algorithm for learning sparse representations. *IEEE transactions on information theory*, 57(7):4689–4708.
 - [46] Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

DEPARTMENT OF INDUSTRIAL AND SYSTEMS ENGINEERING
 VIRGINIA TECH, BLACKSBURG VA 24061
 E-MAIL: wxie@vt.edu

DEPARTMENT OF STATISTICS
 VIRGINIA TECH, BLACKSBURG VA 24061
 E-MAIL: xdeng@vt.edu