

Sparse Regression at Scale: Branch-and-Bound rooted in First-Order Optimization

Hussein Hazimeh*, Rahul Mazumder† and Ali Saab‡

Massachusetts Institute of Technology

April, 2020

Abstract

We consider the least squares regression problem, penalized with a combination of the ℓ_0 and ℓ_2 norms (a.k.a. $\ell_0\ell_2$ regularization). Recent work presents strong evidence that the resulting ℓ_0 -based estimators can outperform popular sparse learning methods, under many important high-dimensional settings. However, exact computation of ℓ_0 -based estimators remains a major challenge. Indeed, state-of-the-art mixed integer programming (MIP) methods for $\ell_0\ell_2$ -regularized regression face difficulties in solving many statistically interesting instances when the number of features $p \sim 10^4$. In this work, we present a new exact MIP framework for $\ell_0\ell_2$ -regularized regression that can scale to $p \sim 10^7$, achieving over 3600x speed-ups compared to the fastest exact methods. Unlike recent work, which relies on modern MIP solvers, we design a specialized nonlinear BnB framework, by critically exploiting the problem structure. A key distinguishing component in our algorithm lies in efficiently solving the node relaxations using specialized first-order methods, based on coordinate descent (CD). Our CD-based method effectively leverages information across the BnB nodes, through using warm starts, active sets, and gradient screening. In addition, we design a novel method for obtaining dual bounds from primal solutions, which certifiably works in high dimensions. Experiments on synthetic and real high-dimensional datasets demonstrate that our method is not only significantly faster than the state of the art, but can also deliver certifiably optimal solutions to statistically challenging instances that cannot be handled with existing methods. We open source the implementation through our toolkit L0BnB.

1 Introduction

We consider the sparse linear regression problem with additional ℓ_2 regularization [12, 29]. A natural way to impose sparsity in this context is through controlling the ℓ_0 (pseudo) norm of the estimator, which counts the corresponding number of nonzero entries. More concretely, let $X \in \mathbb{R}^{n \times p}$ be the data matrix (with n samples and p features) and $y \in \mathbb{R}^n$ be the response vector. We focus on the

*MIT Operations Research Center. Email: hazimeh@mit.edu

†MIT Sloan School of Management, Operations Research Center and MIT Center for Statistics. Email: rahulmaz@mit.edu

‡AJO. Email: saab@mit.edu

least squares problem with a combination of ℓ_0 and ℓ_2 regularization:

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_0 \|\beta\|_0 + \lambda_2 \|\beta\|_2^2, \quad (1)$$

where $\|\beta\|_0$ is defined as the number of nonzero entries in the regression coefficients $\beta \in \mathbb{R}^p$, and $\|\beta\|_2^2$ is the squared ℓ_2 -norm of β (also referred to as ridge regularization). The regularization parameters λ_0 and λ_2 are assumed to be specified by the practitioner¹. We note that the presence of a ridge regularization (i.e., $\lambda_2 > 0$) is important from a statistical viewpoint—see for example [27, 34, 28] for further discussions on this matter. Statistical properties of ℓ_0 -based procedures are well known, e.g., see [23, 39, 45, 17]. Recently there has been exciting work in developing Mixed Integer Programming (MIP)-based approaches to solve (1), e.g., [9, 38, 19, 12, 28, 11, 44]. Current work shows that under many important high-dimensional regimes, ℓ_0 -based estimators² possess statistical properties (variable selection, prediction, and estimation) that are superior to computationally friendlier alternatives such as ℓ_1 regularization (Lasso) [42] and stepwise regression—see [28] for an in-depth discussion.

Despite its appeal, Problem (1) is NP-Hard [35] and poses computational challenges. [9] demonstrated that modern MIP solvers can handle problem instances for p up to a thousand. Larger instances can be handled when λ_2 is sufficiently large and the feature correlations are low, e.g., see [12]; and also [18] for the classification variant of Problem (1). While these state-of-the-art global optimization approaches show very promising results, they are still relatively slow for practical usage [27, 28], as they typically take in the order of minutes to days to obtain optimal solutions. For example, our experiments show that these methods cannot terminate in two hours for typical instances with $p \sim 10^4$. On the other hand, the fast Lasso solvers, e.g., `glmnet` [22], and approximate methods for (1), such as `LOLearn` [28], can handle much larger instances, and they typically terminate in the order of milliseconds to seconds.

Our goal in this paper is to speed up the global optimization of Problem (1). In particular, we aim to (i) reduce the run time for solving typical instances with $p \sim 10^4$ from hours to seconds, and (ii) scale to larger instances with $p \sim 10^7$ in reasonable times (order of seconds to hours). To this end, we propose a specialized nonlinear branch-and-bound (BnB) framework that does not rely on commercial MIP solvers. We employ a first-order method, which carefully exploits the problem structure, to solve the node relaxations. This makes our approach quite different from prior work on global optimization for Problem (1), which rely on commercial MIP solvers, e.g., Gurobi and CPLEX. These MIP solvers are also based on a BnB framework, but they are equipped

¹The choice of (λ_0, λ_2) depends upon the particular application and/or dataset. We aim to compute solutions to (1) for a family of (λ_0, λ_2) -values.

²For example, available from optimal or near-optimal solutions to (1).

with general-purpose relaxation solvers and heuristics that do not take into account the specific structure in Problem (1).

Our BnB solves a mixed integer second order cone program (MISOCP), which is obtained by combining big-M and perspective reformulations [21] of Problem (1). The algorithm exploits the sparsity structure in the problem during the different stages: when solving node relaxations, branching, and obtaining upper bounds. The continuous node relaxations that appear in our BnB have not been studied at depth in earlier work. A main contribution of our work is to show that these relaxations, which involve seemingly complicated linear and conic constraints, can be efficiently handled using a primal coordinate descent (CD)-based algorithm. Indeed, this represents a radical change from the primal-dual relaxations solvers commonly used in state-of-the-art MIP solvers [7]. Our choice of CD is motivated by its ability to effectively share information across the BnB nodes (e.g., warm starts and gradient information), and more generally by its high scalability in the context of sparse learning, e.g., [22, 33, 28]. Along with CD, we propose additional strategies, namely, active set updates and gradient screening, which reduce the coordinate update complexity by exploiting the information shared across the BnB tree.

Although our CD-based algorithm is highly scalable, it only generates primal solutions. However, dual bounds are required for search space pruning in BnB. Thus, we propose a novel method to efficiently generate dual bounds from the primal solutions. We analyze these dual bounds and prove that their tightness is not affected by the number of features p , but rather by the number of nonzeros in the primal solution. This result serves as a theoretical justification for why our CD-based algorithm can lead to tight dual bounds in high dimensions.

Contributions and Structure: We summarize our key contributions below.

- We propose a hybrid formulation, a MISOCP, which combines big-M and perspective formulations of Problem (1). We provide a new analysis of the relaxation tightness, which shows that the hybrid formulation can significantly outperform the big-M and perspective formulations. See Sec. 2.
- To solve the MISOCP, we design a specialized nonlinear BnB, with the following main contributions (see Section 3):
 - We prove that the node relaxations, which involve linear and conic constraints, can be reformulated as a least squares problem with a non-differentiable but separable penalty. To solve the latter reformulation, we develop a primal CD algorithm, along with active set updates and gradient screening, which exploits the information shared across the BnB tree to reduce the coordinate update complexity.

- We develop a new efficient method for obtaining dual bounds from the primal solutions. We analyze these dual bounds and show that their tightness depends on the sparsity level rather than p .
- We introduce efficient methods that exploit sparsity when selecting branching variables and obtaining incumbents³.
- We perform a series of experiments on high-dimensional synthetic and real datasets, with p up to 8.3×10^6 . We study the effect of the regularization parameters and dataset characteristics on the run time, and perform ablation studies. The results indicate that our approach is not only over 3600x faster than the state of the art, but also capable of handling difficult statistical instances which were virtually unsolvable before. We open source the implementation through our toolkit L0BnB⁴. See Section 4.

Related Work: An impressive line of recent work considers solving Problem (1), or its cardinality-constrained variant, to optimality. [9] used Gurobi on a big-M based formulation, which can handle $n \sim p \sim 10^3$ in the order of minutes to hours. [12] scale the problem even further by applying outer-approximation (using Gurobi) on a boolean reformulation [38] of the problem. Their approach can handle $p \sim 10^5$ in the order of minutes when n and λ_2 are sufficiently large, and the feature correlations are sufficiently small. This outer-approximation approach has also been generalized to sparse classification in [10]. [44] considered solving a perspective formulation [21] of the problem directly using Gurobi and reported competitive timings with [12]—the largest problem instances considered have $p \sim 10^3$. [18] show that a variant of Problem (1) for classification can be solved through a sequence of MIPs (using Gurobi), each having a small number of binary variables, as opposed to p binary variables in the common approaches. Their approach can handle $n = 10^3$ and $p = 50,000$ in minutes if the feature correlations are sufficiently small. Our specialized BnB, on the other hand, can solve all the instance sizes mentioned above with speed-ups that exceed 3600x, and can scale to problems with $p \sim 10^7$. Moreover, our empirical comparisons show that, unlike the previous work, our BnB can handle difficult problems with relatively small λ_2 or high feature correlations.

In addition to the global optimization approaches discussed above, there is an interesting body of work in the broader optimization community on improved relaxations for sparse ridge regression, e.g., [38, 19, 4, 44], and algorithms that locally optimize an ℓ_0 -based objective [13, 5, 28].

There is also a rich literature on solving mixed integer nonlinear programs (MINLPs) using BnB (see e.g., [31, 7]). Our approach is based on the nonlinear BnB framework [16], where a nonlinear

³An incumbent, in the context of BnB, refers to the best integral solution found so far.

⁴<https://github.com/alisaab/L0BnB>

subproblem is solved at every node of the search tree. Interior point methods are a popular choice for these nonlinear subproblems, especially for MISOCPs [31], e.g., they are used in MOSEK [2] and CPLEX [40]. Generally, nonlinear solvers are not as effective in exploiting warm starts and sparsity as LP solvers [7], which led to an alternative approach known as outer-approximation (OA) [20]. In OA, a sequence of relaxations, consisting of mixed integer linear programs, are solved until converging to a solution of the MINLP, e.g., this is implemented in Baron [41] and Gurobi [25]. There is also a line of work on specialized OA reformulations and algorithms for mixed integer conic programs (which include MISOCPs), e.g., see [43, 32] and the references therein. However, for our problem, making use of its specific structure, we show that the relaxation of the MISOCP can be effectively handled by our proposed CD-based algorithm.

Notation: We denote the set $\{1, 2, \dots, p\}$ by $[p]$. For any set A , the complement is denoted by A^c . $\|\cdot\|_q$ denotes the standard L_q norm with $q \in \{0, 1, 2, \infty\}$. For any vector $v \in \mathbb{R}^k$, $\text{sign}(v) \in \mathbb{R}^k$ refers to the vector whose i th component is given by $\text{sign}(v_i) = v_i/|v_i|$ if $v_i \neq 0$ and $\text{sign}(v_i) \in [-1, 1]$ if $v_i = 0$. We denote the support of $\beta \in \mathbb{R}^p$ by $\text{Supp}(\beta) = \{i : \beta_i \neq 0, i \in [p]\}$. For a set $S \subseteq [p]$, use $\beta_S \in \mathbb{R}^{|S|}$ to denote the subvector of β with indices in S . Similarly, X_S refers to the submatrix of X whose columns correspond to S .

2 MIP Formulations and Relaxations

In this section, we first discuss two popular MIP formulations for Problem (1), namely, the big-M and perspective formulations, and study the quality of their corresponding relaxations. We then propose a hybrid formulation, and identify regimes where it outperforms the big-M and perspective formulations, in terms of the relaxation quality.

2.1 Existing Formulations

The Big-M Formulation: Problem (1) can be modeled as a mixed integer quadratic program (MIQP) using a big-M formulation. We assume that there is a finite scalar M (a-priori specified) such that an optimal solution of (1), say β^* , satisfies: $\|\beta^*\|_\infty \leq M$. This leads to the following formulation:

$$\begin{aligned}
 \min_{\beta, z} \quad & \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_0 \sum_{i \in [p]} z_i + \lambda_2 \|\beta\|_2^2 \\
 \text{s.t.} \quad & -Mz_i \leq \beta_i \leq Mz_i, \quad i \in [p] \\
 & z_i \in \{0, 1\}, \quad i \in [p],
 \end{aligned} \tag{2}$$

where each binary variable z_i controls whether β_i is zero or not via the first constraint in (2)—i.e., if $z_i = 0$ then $\beta_i = 0$. Such big-M formulations are widely used in mixed integer programming and have been recently explored in multiple works on ℓ_0 regularization, e.g., [9, 34, 28, 44]. See [9, 44] for a discussion on how to estimate M in practice.

The Perspective Formulation: Problem (1) can be alternatively modeled using a perspective formulation [21, 24], which does not require big-M constraints. As before, we associate a binary variable z_i with every β_i . We consider a lifted formulation by using the auxiliary continuous variables: $s_i \geq 0, i \in [p]$, which leads to the following MISOCP:

$$\begin{aligned} \min_{\beta, z, s} \quad & \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_0 \sum_{i \in [p]} z_i + \lambda_2 \sum_{i \in [p]} s_i \\ \text{s.t.} \quad & \beta_i^2 \leq s_i z_i, \quad i \in [p] \\ & z_i \in \{0, 1\}, s_i \geq 0, \quad i \in [p], \end{aligned} \tag{3}$$

where each s_i takes the place of the β_i^2 term in the objective. The rotated second order cone constraint $\beta_i^2 \leq s_i z_i$ enforces: $z_i = 0 \implies \beta_i = 0$. Formulation (3) was used in [19], and by [44] for the cardinality constrained variant of Problem (1).

The Boxed Perspective Formulation: Later in this section, we study tightness of the relaxations for the big-M and perspective formulations. However, it is not fair to directly compare these two formulations, because the big-M formulation uses additional information, i.e., some optimal solution β^* to (1) satisfies $\|\beta^*\|_\infty \leq M$. To ensure a fair comparison, we augment the perspective formulation with box constraints, leading to the following MISOCP:

$$\begin{aligned} \min_{\beta, z, s} \quad & \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_0 \sum_{i \in [p]} z_i + \lambda_2 \sum_{i \in [p]} s_i \\ \text{s.t.} \quad & \beta_i^2 \leq s_i z_i, \quad i \in [p] \\ & -M \leq \beta_i \leq M, \quad i \in [p] \\ & z_i \in \{0, 1\}, s_i \geq 0, \quad i \in [p], \end{aligned} \tag{4}$$

where M has a similar interpretation as in (2). While formulation (4) does not appear explicitly in the literature, it is a natural candidate to consider if a good M is known.

In what follows, we compare the three formulations considered above, in terms of both the tightness of the relaxations and the integrality of the corresponding solutions.

Relaxation Quality: Recent works are showing promising results from perspective formulations similar to that in (3), e.g., [38, 19, 12, 10, 44]. We complement this line of work by showing that neither the big-M nor the perspective formulation dominates in terms of relaxation quality. More specifically, in Proposition 1, we establish a new result which characterizes the parameter ranges over which each formulation has a tighter relaxation. Our focus here is on the *interval relaxation* of each formulation, which is obtained by relaxing all the z_i 's in the formulation to the interval $[0, 1]$.

Before we present Proposition 1, we introduce some notation. Let $\mathcal{S}(\lambda_2)$ be the set of optimal solutions (in β) of the interval relaxation of (3); and define

$$\lambda_2^* := \inf\{\lambda_2 \geq 0 \mid \exists \beta \in \mathcal{S}(\lambda_2) \text{ s.t. } \|\beta\|_\infty \leq M\}. \quad (5)$$

Proposition 1. *Let M be the constant used in formulation (2) and λ_2^* be as defined in (5). Moreover, let v_1, v_2, v_3 be the optimal objectives of the interval relaxations of (2), (3), (4), respectively. Then, the following holds:*

$$v_1 \geq v_2 \geq v_3, \quad \text{if } M \leq \frac{1}{2}\sqrt{\lambda_0/\lambda_2} \quad (6)$$

$$v_1 \leq v_3, \quad \text{if } M \geq \sqrt{\lambda_0/\lambda_2}, \quad (7)$$

$$v_1 \leq v_2, \quad \text{if } M \geq \sqrt{\lambda_0/\lambda_2} \text{ and } \lambda_2 \geq \lambda_2^*. \quad (8)$$

Proposition 1 implies that if λ_2 is relatively small, then the big-M (with objective v_1) will have the tightest relaxation. On the other hand, if λ_2 is sufficiently large, then the two perspective relaxations (with objectives v_2 and v_3) will be tighter than the big-M. Moreover, the condition on λ_2 for which $v_1 \leq v_2$ holds, is generally more restrictive than that required for $v_1 \leq v_3$ to hold. We note that the result of Proposition 1 applies for any $M \geq 0$, even if it is mis-specified. We also point out that [44] studied the relaxations of several formulations for the cardinality-constrained variant of Problem (1), including big-M and perspective formulations. However, [44] does not compare the tightness of the big-M to the perspective relaxation.

In practice, we solve Problem (1) for a regularization path, i.e., for a two-dimensional grid of values in (λ_0, λ_2) . This grid can be wide, making both the big-M and perspective formulations useful—each for certain choices of the parameters (as detailed in Proposition 1).

Solution Integrality: The interval relaxation of the perspective formulation encourages integral solutions, which is an attractive property. In particular, if \tilde{z} is a solution to the interval relaxation of (3); then not only will many coordinates in \tilde{z} be zero (i.e., \tilde{z} is sparse), but many coordinates can be exactly one. This interesting observation was made by [38], who derived conditions under which

the relaxation of (3) leads to integral solutions⁵. Empirically, we see that for a sufficiently large value of λ_2 , the relaxation of (3) can set most of the entries in z to either 0 or 1. This observation is also supported by Proposition 1 which shows that the perspective formulation can outperform the big-M for sufficiently large λ_2 . The integrality of z in the relaxation can enhance the performance of BnB as there will be less variables to branch on. In contrast, the interval relaxation of the big-M formulation in (2) leads to an $\ell_1\ell_2$ -regularized least squares problem with a box constraint on β —this is a close cousin of the Elastic Net [47]. While the $\ell_1\ell_2$ relaxation may lead to a sparse solution \tilde{z} , it rarely results in \tilde{z}_i 's taking the value one, in practice.

2.2 A Hybrid Formulation

The discussions above suggest that neither of the interval relaxations for the big-M and perspective formulations dominates the other. Hence, it is natural to merge the big-M and perspective formulations into one hybrid formulation, so as to combine their benefits. To this end, we define the *hybrid formulation*:

$$\begin{aligned}
\min_{\beta, z, s} \quad & \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_0 \sum_{i \in [p]} z_i + \lambda_2 \sum_{i \in [p]} s_i \\
\text{s.t.} \quad & \beta_i^2 \leq s_i z_i, \quad i \in [p] \\
& -Mz_i \leq \beta_i \leq Mz_i, \quad i \in [p] \\
& z_i \in \{0, 1\}, s_i \geq 0, \quad i \in [p]
\end{aligned} \tag{9}$$

The MISOCP in (9) is equivalent to (1) (as long as M is suitably chosen). In Section 2.3, we study the interval relaxation of (9) and identify parameter ranges for which it can be significantly tighter than that of formulations (2),(3), and (4). We note that [44] mentions that combining perspective and big-M formulations can potentially lead to tighter relaxations for the cardinality constrained variant of (1). In our case, we provide new bounds that quantify the strengths of these relaxations (see Proposition 2). Moreover, we are the first to present an in-depth computational study of formulation (9).

2.3 Relaxation of the Hybrid Formulation

In this section, we study the interval relaxation of (9) (i.e., by relaxing all z_i 's to $[0, 1]$) and discuss how this compares with the other formulations in (2), (3), and (4). In Theorem 1, we show that the interval relaxation of (9) can be reformulated in the β space. This leads to a regularized least squares criterion, where the regularizer involves the reverse Huber [37] penalty—a hybrid between

⁵The relaxation considered in [38] is derived using a duality-based argument, without explicitly using a perspective reformulation. However, as pointed out in [19] the relaxation of (3) is the same as that derived in [38].

the ℓ_1 and ℓ_2 (squared) penalties. The reverse Huber penalty $\mathcal{B} : \mathbb{R} \rightarrow \mathbb{R}$, is given by:

$$\mathcal{B}(t) = \begin{cases} |t| & |t| \leq 1 \\ (t^2 + 1)/2 & |t| \geq 1. \end{cases} \quad (10)$$

Theorem 1. (*Reduced Relaxation*) Let us define the function $\psi(\cdot)$ as

$$\psi(\beta_i; \lambda_0, \lambda_2, M) = \begin{cases} \psi_1(\beta_i; \lambda_0, \lambda_2) := 2\lambda_0\mathcal{B}(\beta_i\sqrt{\lambda_2/\lambda_0}) & \text{if } \sqrt{\lambda_0/\lambda_2} \leq M \\ \psi_2(\beta_i; \lambda_0, \lambda_2, M) := (\lambda_0/M + \lambda_2M)|\beta_i| & \text{if } \sqrt{\lambda_0/\lambda_2} \geq M \end{cases}$$

The interval relaxation of (9) is equivalent to:

$$\min_{\beta \in \mathbb{R}^p} F(\beta) := \frac{1}{2}\|y - X\beta\|_2^2 + \sum_{i \in [p]} \psi(\beta_i; \lambda_0, \lambda_2, M) \quad \text{s.t.} \quad \|\beta\|_\infty \leq M, \quad (11)$$

and we let v_4 denote the optimal objective value of (11).

Compared to the extended formulation involving (β, s, z) , the reduced formulation (11) plays an important role for both the subsequent analysis and the algorithmic development in Section 3. Theorem 1 shows that the conic and big-M constraints in the relaxation of (9) can be completely eliminated, at the expense of introducing (in the objective) the non-differentiable penalty function $\sum_i \psi(\beta_i; \lambda_0, \lambda_2, M)$ which is separable across the p coordinates $\beta_i, i \in [p]$. Depending on how small $\sqrt{\lambda_0/\lambda_2}$ is compared to M , the penalty $\psi(\beta_i; \lambda_0, \lambda_2, M)$ is either the reverse Huber or the ℓ_1 , both of which are sparsity-inducing. The appearance of a pure ℓ_1 penalty is interesting in this case since the original formulation in (9) has a ridge ($\sum_i s_i$) term in the objective. It turns out that when $\sqrt{\lambda_0/\lambda_2} \geq M$, the constraint $|\beta_i| \leq Mz_i$ becomes active at any optimal solution, which turns the ridge term into an ℓ_1 penalty—see the proof of Theorem 1 for more insight. Next, we analyze the tightness of the hybrid relaxation in (11).

Hybrid Relaxation Tightness: [19] has shown that the interval relaxation of the perspective formulation (3) simplifies to:

$$v_2 = \min_{\beta \in \mathbb{R}^p} G(\beta) := \frac{1}{2}\|y - X\beta\|_2^2 + \sum_{i \in [p]} \psi_1(\beta_i; \lambda_0, \lambda_2), \quad (12)$$

where $\psi_1(\beta_i; \lambda_0, \lambda_2)$ is defined in Theorem 1. Moreover, the interval relaxation of the boxed per-

spective formulation in (4) can be similarly simplified to:

$$v_3 := \min_{\beta \in \mathbb{R}^p} G(\beta) \quad \text{s.t.} \quad \|\beta\|_\infty \leq M. \quad (13)$$

For $\sqrt{\lambda_0/\lambda_2} \leq M$, the hybrid relaxation in Theorem 1 exactly matches that in (13). Thus, for $\sqrt{\lambda_0/\lambda_2} \leq M$, the big-M constraints in the hybrid formulation (9) can be replaced with the box constraint $\|\beta\|_\infty \leq M$, without affecting the interval relaxation. Moreover, in this case, the hybrid relaxation is generally tighter than (12) (because of the additional box constraint). However, when $\sqrt{\lambda_0/\lambda_2} \geq M$, the hybrid formulation can be (strictly) tighter than both the big-M and perspective relaxations, as we show in Proposition 2. In Table 1, we summarize the formulations considered.

Table 1: Formulations considered in Proposition 2.

Formulation	Relaxation Objective
Big-M (2)	v_1
Perspective (3)	v_2
Boxed Perspective (4)	v_3
Hybrid (9)	v_4

Proposition 2. *Suppose the same constant M is used in all the formulations considered, i.e., (2), (4), and (9). Let β^* be an optimal solution to (11), and define the function $h(\lambda_0, \lambda_2, M) = \lambda_0/M + \lambda_2 M - 2\sqrt{\lambda_0\lambda_2}$. Then, the following holds for $\sqrt{\lambda_0/\lambda_2} \geq M$:*

$$v_4 \geq v_1 + \lambda_2(M\|\beta^*\|_1 - \|\beta^*\|_2^2) \quad (14)$$

$$v_4 \geq v_3 + h(\lambda_0, \lambda_2, M)\|\beta^*\|_1 \geq v_2 + h(\lambda_0, \lambda_2, M)\|\beta^*\|_1 \quad (15)$$

We make a couple of remarks. For bound (14), we always have $(M\|\beta^*\|_1 - \|\beta^*\|_2^2) \geq 0$ since $\|\beta^*\|_\infty \leq M$. If there is an $i \in [p]$ with $0 < |\beta_i^*| < M$, then $(M\|\beta^*\|_1 - \|\beta^*\|_2^2) > 0$, and consequently $v_4 > v_1$ (as long as $\lambda_2 > 0$).

For $\sqrt{\lambda_0/\lambda_2} \geq M$, $h(\lambda_0, \lambda_2, M)$ is non-negative and monotonically decreasing in M . In particular, for $\sqrt{\lambda_0/\lambda_2} > M$, we have $h(\lambda_0, \lambda_2, M) > 0$, and consequently $v_4 > v_3 \geq v_2$ (as long as $\beta^* \neq 0$). Moreover, it is easy to see that if the parameters are chosen so that $\beta^* \neq 0$ and M is chosen to be sufficiently small, then $\|\beta^*\|_\infty = M$, so $\|\beta^*\|_1 \geq M$, and the bound in (15) can be relaxed to: $v_4 \geq v_3 + h(\lambda_0, \lambda_2, M)M$ (which does not depend on β^*). In Section 4.2, our experiments empirically validate Proposition 2: using a hybrid formulation with a sufficiently tight big-M can speed up the same BnB solver by more than 10x compared to a pure perspective formulation.

3 A Specialized Branch-and-bound (BnB) Framework

In this section, we develop a specialized nonlinear BnB framework for solving the hybrid formulation in (9). First, we briefly recall the high-level mechanism behind nonlinear BnB, in the context of our problem.

Nonlinear BnB at a Glance: The algorithm starts by solving a nonlinear relaxation for (9), i.e., the root node. Then, it selects a branching variable, say variable $j \in [p]$, and creates two new nodes (optimization subproblems): one node with $z_j = 0$ and another with $z_j = 1$, where all the other z_i 's are relaxed to the interval $[0, 1]$. For every unvisited node, the algorithm proceeds recursively, i.e., by solving an optimization subproblem at the current node and then branching on a new variable to create two new nodes. This leads to a search tree with nodes corresponding to optimization subproblems and edges representing branching decisions.

To reduce the size of the search tree, our BnB prunes a node (i.e., does not branch on it) in either one of the following situations: (i) the relaxation at the current node has an integral z or (ii) the objective of the current relaxation exceeds the best available upper bound on (9). The relaxation need not be solved exactly; lower bounds (a.k.a. dual bounds) can be used instead.

Note that at any level in the search tree, the minimum objective value across the nodes of that level is a valid lower bound on (9). The lower bound increases with every new level in the search tree. Thus, we can terminate the algorithm early on if the gap between the best lower and upper bounds is below a pre-specified threshold.

Overview of Our Strategies: Our discussion above outlines how nonlinear BnB operates in general. Of course, the specific strategies used, e.g., solving the relaxations, passing information across the nodes, and selecting branching variables, can have a key impact on scalability. In the rest of this section, we will give a detailed account of the strategies we use in our BnB. We first provide an overview of these strategies:

- **A Primal Relaxation Solver:** Unlike the state-of-the-art approaches for nonlinear BnB, which employ primal-dual relaxation solvers [7], we rely solely on a primal method. We design a highly scalable CD-based algorithm for solving the continuous node subproblems corresponding to the hybrid formulation in (9). In particular, our CD solves the subproblems in the β space as opposed to the extended (β, s, z) space—these subproblems are variants of the reduced relaxation we introduced in (11). The algorithm heavily shares and exploits warm starts, active sets, and information on the gradients, across the BnB tree. This will be developed in Section 3.1.

- **Dual Bounds:** Dual bounds on the relaxation are required by the BnB for search space pruning, yet our relaxation solver works in the primal space for scalability considerations. We develop a new efficient method for obtaining dual bounds from the primal solutions. We provide an analysis of this method and show that the tightness of the dual bounds depends on the sparsity level and *not* on the number of features p . See Section 3.2.
- **Branching and Incumbents:** We present an efficient variant of strong branching, which leverages the solutions and active sets of previous node relaxations to make optimization tractable. Moreover, we employ several efficient heuristics to obtain incumbents. See Section 3.3.

Without loss of generality, in the remainder of Section 3, we assume that the columns of X and y have unit ℓ_2 norm.

3.1 Primal Relaxation Solver: Active-set Coordinate Descent

To simplify the presentation, we will focus on solving the root relaxation, through using the reduced formulation in (11). Note that the reduced formulation can improve the efficiency of the algorithm as compared to solving interval relaxation of (9) in the extended (β, s, z) space. The rest of the nodes in BnB involve fixing some binary variables in (9) to 0 or 1, so their relaxations can be obtained by minor modifications to the root relaxation.

In relaxation (11), the objective has a composite form [36]: a smooth loss function along with a non-smooth but separable penalty. The feasible set, consisting of the constraints $|\beta_i| \leq M$, $i \in [p]$ is also separable across the coordinates. This makes Problem (11) amenable to cyclic coordinate descent (CD). To our knowledge, cyclic CD has not been applied to the reverse Huber penalty which appears as part of relaxation (11). We also emphasize that a direct application of cyclic CD to Problem (11) will face scalability issues, and more importantly, it cannot lead to dual bounds. The additional strategies we develop later in this section are essential for both scaling the algorithm and obtaining (provably) high quality dual bounds.

Cyclic CD visits the coordinates according to a fixed ordering, updating one coordinate at a time, as detailed in Algorithm 1.

Algorithm 1: Cyclic CD for Relaxation (11)

- **Input:** Initialization $\hat{\beta}$
- **While** not converged:
 - **For** $i \in [p]$:

$$\hat{\beta}_i \leftarrow \arg \min_{\beta_i \in \mathbb{R}} F(\hat{\beta}_1, \dots, \beta_i, \dots, \hat{\beta}_p) \quad \text{s.t. } |\beta_i| \leq M. \tag{16}$$

Every limit point of Algorithm 1 is a global minimizer of (11) [8], and the objective values converge at a (non-asymptotic worst-case) rate of $\mathcal{O}(1/k)$, where k is the iteration counter [6]. We will show that a solution of (16) can be computed in closed-form. To this end, since the columns of X have unit ℓ_2 -norm, we note that (16) is equivalent to:

$$\min_{\beta_i \in \mathbb{R}} \frac{1}{2}(\beta_i - \tilde{\beta}_i)^2 + \psi(\beta_i; \lambda_0, \lambda_2, M) \quad \text{s.t. } |\beta_i| \leq M, \quad (17)$$

where $\tilde{\beta}_i := \langle y - \sum_{j \neq i} X_j \hat{\beta}_j, X_i \rangle$. Given non-negative scalar parameters a and m , we define the *boxed soft-thresholding operator* $T : \mathbb{R} \rightarrow \mathbb{R}$ as

$$T(t; a, m) := \begin{cases} 0 & \text{if } |t| \leq a \\ (|t| - a) \text{sign}(t) & \text{if } a \leq |t| \leq a + m \\ m \text{sign}(t) & \text{o.w.} \end{cases}$$

For $\sqrt{\lambda_0/\lambda_2} \leq M$, the solution of (17) is given by:

$$\hat{\beta}_i = \begin{cases} T(\tilde{\beta}_i; 2\sqrt{\lambda_0\lambda_2}, M) & \text{if } |\tilde{\beta}_i| \leq 2\sqrt{\lambda_0\lambda_2} + \sqrt{\lambda_0/\lambda_2} \\ T(\tilde{\beta}_i(1 + \lambda_2)^{-1}; 0, M) & \text{o.w.} \end{cases} \quad (18)$$

and for $\sqrt{\lambda_0/\lambda_2} \geq M$, the solution of (17) is:

$$\hat{\beta}_i = T(\tilde{\beta}_i; \lambda_0/M + \lambda_2 M, M). \quad (19)$$

Thus, update (16) in Algorithm 1 can be computed in closed-form using expressions (18) and (19).

Warm Starts: The interior-point methods used in the state-of-the-art nonlinear BnB solvers cannot easily utilize warm starts. On the other hand, cyclic CD has proven to be very effective in exploiting warm starts, e.g., see [22, 28]. For every node in BnB, we use its parent's primal solution as a warm start. Recall that the parent and children nodes have the same relaxations, except that a single binary variable is fixed to zero or one in the children. Intuitively, this can make the warm start close to the optimal solution. Moreover, the supports of the optimal solutions of the parent and children nodes are typically very close. We exploit this observation by sharing information about the supports across the BnB tree, as we describe next in Section 3.1.1.

3.1.1 Active Sets

Update (16) in Algorithm 1 requires $\mathcal{O}(n)$ operations, so every full cycle (i.e., across all p coordinates) of the algorithm has an $\mathcal{O}(np)$ time complexity. This becomes a major bottleneck for large n or p , since CD can require many cycles before convergence. In practice, the majority of the variables stay zero during the course of Algorithm 1 (assuming that the regularization parameters are chosen so that the optimal solution is sparse and good warm starts are used). Motivated by this observation, we run Algorithm 1 restricted to a small subset of the variables $\mathcal{A} \subseteq [p]$, which we refer to as the *active set*, i.e., we solve the following restricted problem:

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} F(\beta) \quad \text{s.t.} \quad \|\beta\|_\infty \leq M, \beta_{\mathcal{A}^c} = 0. \quad (20)$$

After solving (20), we augment the active set with any variable $i \in \text{Supp}(\hat{\beta})^c$ that violates the following optimality condition:

$$0 = \arg \min_{\beta_i \in \mathbb{R}} F(\hat{\beta}_1, \dots, \beta_i, \dots, \hat{\beta}_p) \quad \text{s.t.} \quad |\beta_i| \leq M.$$

We repeat this procedure of solving the restricted problem (20) and then augmenting \mathcal{A} with the violating variables, until there are no more violations. The algorithm is summarized below.

Algorithm 2: The Active-set Algorithm⁶

- **Input:** Initial solution $\hat{\beta}$ and initial active set \mathcal{A} .
- **Repeat:**
 - Step 1: Solve (20) using Algorithm 1 to get a solution $\hat{\beta}$.
 - Step 2: $\mathcal{V} \leftarrow \{i \in \text{Supp}(\hat{\beta})^c \mid 0 \neq \arg \min_{|\beta_i| \leq M} F(\hat{\beta}_1, \dots, \beta_i, \dots, \hat{\beta}_p)\}$.
 - Step 3: If \mathcal{V} is empty **terminate**, otherwise, $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{V}$.

The quality of the initial active set affects the number of iterations in Algorithm 2. For example, if \mathcal{A} is a super set of the support of an optimal solution, then \mathcal{V} in the first iteration of Algorithm 2 will be empty and the algorithm will terminate in a single iteration. For every node in BnB, we choose the initial active set to be the same as the final active set of the parent node. This works well in practice because parent and children nodes typically have very similar supports. For the root relaxation, we obtain the initial active set by choosing a small subset of features which have the highest (absolute) correlation with y .

In Section 3.2, we present a novel method that makes use of the updates in Algorithm 2 to obtain provably high-quality dual bounds. We note that our approach goes beyond the standard

⁶Algorithm 2 solves the root relaxation. For other nodes, the objective function $F(\beta)$ will need to be modified to account for the fixed variables, and all the variables fixed to zero at the current node should be excluded from the set \mathcal{V} in Step 2.

active-set methods [22]. In particular, (i) we use active sets in the context of a BnB tree, percolating information from parents to children (as opposed to warm-start continuation across a grid of regularization parameters); and (ii) we exploit the active sets to deliver dual bounds—connections to dual bounds are not discussed in standard active-set methods that deliver approximate primal solutions. In the next remark, we discuss how Step 1 in Algorithm 2 can be performed inexactly.

Remark 1. *In practice, we solve the restricted optimization problem in Step 1 of Algorithm 2 inexactly, e.g., a heuristic termination criterion can be used in Algorithm 1. However, we perform Step 3 exactly (which ensures there are no optimality violations outside the active set)—this step is essential to obtain the tight dual bounds discussed in Section 3.2.*

3.1.2 Gradient Screening

In practice, Algorithm 2 effectively reduces the number of coordinate updates, through restricting optimization to the active set. However, checking the optimality conditions outside the active set, i.e., performing Step 2, requires $\mathcal{O}(np)$ operations. This is a bottleneck when p is large, even if a small number of such checks (or passes) is performed. To mitigate this, we present a *gradient screening* method which reduces the time complexity of these optimality checks. Our method is inspired by the gradient screening proposed in [29] for a different problem: learning sparse hierarchical interactions. In the current paper, the optimality checks in Step 2 of Algorithm 2 essentially require computing a gradient of the least squares loss, in order to construct \mathcal{V} —gradient screening will avoid computing the “non-essential” parts of this gradient by using previously computed quantities in the BnB tree.

In the following proposition, we give an explicit way to construct the set \mathcal{V} in Step 2 of Algorithm 2.

Proposition 3. *Let $\hat{\beta}$ and \mathcal{V} be as defined in Algorithm 2, and define $\hat{r} = y - X\hat{\beta}$. Then, the set \mathcal{V} can be equivalently written as follows:*

$$\mathcal{V} = \{i \in \text{Supp}(\hat{\beta})^c \mid |\langle \hat{r}, X_i \rangle| > c(\lambda_0, \lambda_2, M)\}, \quad (21)$$

where $c(\lambda_0, \lambda_2, M) = 2\sqrt{\lambda_0\lambda_2}$ if $\sqrt{\lambda_0/\lambda_2} \leq M$, and $c(\lambda_0, \lambda_2, M) = (\lambda_0/M + \lambda_2M)$ if $\sqrt{\lambda_0/\lambda_2} \geq M$.

By Proposition 3, constructing \mathcal{V} directly costs $\mathcal{O}(n(p - \|\hat{\beta}\|_0))$. Next, we will discuss how to compute \mathcal{V} with a lower cost, by making use of previously computed quantities. Suppose we have access to a set $\hat{\mathcal{V}} \subset [p]$ such that $\mathcal{V} \subset \hat{\mathcal{V}}$. Then, we can construct \mathcal{V} by restricting the checks in (21) to $\hat{\mathcal{V}}$ instead of $\text{Supp}(\hat{\beta})^c$, i.e., the following holds:

$$\mathcal{V} = \{i \in \hat{\mathcal{V}} \mid |\langle \hat{r}, X_i \rangle| > c(\lambda_0, \lambda_2, M)\}. \quad (22)$$

Assuming $\hat{\mathcal{V}}$ is available, the complexity of computing \mathcal{V} in (22) is $\mathcal{O}(n|\hat{\mathcal{V}}|)$, which can be significantly better than that of (21), if $|\hat{\mathcal{V}}|$ is sufficiently small. Next, we present a method that can obtain a relatively small $\hat{\mathcal{V}}$ in practice, thus speeding up the computation of \mathcal{V} .

Computation of $\hat{\mathcal{V}}$: Proposition 4 presents a method to construct a set $\hat{\mathcal{V}}$ which satisfies $\mathcal{V} \subset \hat{\mathcal{V}}$ (as discussed above), using information from a warm start β^0 (e.g., the solution of the relaxation from the parent node in BnB).

Proposition 4. *Let $\hat{\beta}$ and \mathcal{V} be as defined in Algorithm 2. Let β^0 be an arbitrary vector in \mathbb{R}^p . Define $\hat{r} = y - X\hat{\beta}$, $r^0 = y - X\beta^0$, and $\epsilon = \|X\beta^0 - X\hat{\beta}\|_2$. Then, the following holds:*

$$\mathcal{V} \subseteq \hat{\mathcal{V}} \stackrel{\text{def}}{=} \left\{ i \in \text{Supp}(\hat{\beta})^c \mid |\langle r^0, X_i \rangle| > c(\lambda_0, \lambda_2, M) - \epsilon \right\}. \quad (23)$$

Note that the set $\hat{\mathcal{V}}$ defined in (23) depends solely on the quantities $|\langle r^0, X_i \rangle|$, $i \in \text{Supp}(\hat{\beta})^c$, which we will assume to be sorted and available along with the warm start β^0 . This is in contrast to a direct evaluation of \mathcal{V} in (21), which requires the costly computation of $|\langle \hat{r}, X_i \rangle|$, $i \in \text{Supp}(\hat{\beta})^c$. Given the sorted values $|\langle r^0, X_i \rangle|$, $i \in \text{Supp}(\hat{\beta})^c$, we can identify $\hat{\mathcal{V}}$ in (23) with $\mathcal{O}(\log p)$ operations, using a variant of binary search. Thus, the overall cost of computing \mathcal{V} using (22) is $\mathcal{O}(n|\hat{\mathcal{V}}| + \log p)$.

From Proposition 4, the closer $X\beta^0$ is to $X\hat{\beta}$, the smaller $\hat{\mathcal{V}}$ is, i.e., the lower is the time complexity of computing \mathcal{V} in (22). Thus, for the method to be successful, we need a good warm start β^0 . In practice, we obtain β^0 for the current node in BnB from its parent, and we update β^0 as necessary during the course of Algorithm 2 (as detailed below). Next, we present the gradient screening procedure formally: we replace Step 2 of Algorithm 2 by the following:

Gradient Screening

1. If this is the root node and the first iteration of Algorithm 2, then: $\beta^0 \leftarrow \hat{\beta}$, $r^0 = y - X\beta^0$, compute and sort $|\langle r^0, X_i \rangle|$, $i \in [p]$, compute \mathcal{V} using (21), and skip all the steps below.
2. If this is the first iteration of Algorithm 2, get β^0 from the parent node in the BnB tree.
3. Compute $\hat{\mathcal{V}}$ using (23) and then \mathcal{V} using (22)⁷.
4. If $|\hat{\mathcal{V}}| > 0.05p$, $\beta^0 \leftarrow \hat{\beta}$, $r^0 = y - X\beta^0$, and re-compute and sort $|\langle r^0, X_i \rangle|$, $i \in [p]$.

If $X\beta^0$ is not a good estimate of $X\hat{\beta}$, then $\hat{\mathcal{V}}$ might become large. To avoid this issue, we update β^0 in Step 4 above, every time the set $\hat{\mathcal{V}}$ becomes relatively large ($|\hat{\mathcal{V}}| > 0.05p$). This updated β^0 will be passed to the children in the BnB tree. While Step 4 above can be costly, it is not performed often in practice as the solutions of the relaxations in the parent and children nodes are typically close.

⁷Each variable i fixed to zero by BnB at the current node should be excluded when computing \mathcal{V} and $\hat{\mathcal{V}}$.

3.2 Dual Bounds

In practice, we use Algorithm 2 to obtain inexact primal solutions for relaxation (11), as discussed in Remark 1. However, dual bounds are needed to perform search space pruning in BnB. Here, we present a new efficient method to obtain dual bounds from the primal solutions. Moreover, we prove that our method can obtain dual bounds whose tightness depends on the sparsity level rather than p . We start by introducing a Lagrangian dual of relaxation (11) in Theorem 2.

Theorem 2. For $\sqrt{\lambda_0/\lambda_2} \leq M$, a dual of Problem (11) is given by:

$$\max_{\alpha \in \mathbb{R}^n, \gamma \in \mathbb{R}^p} h_1(\alpha, \gamma) := -\frac{1}{2}\|\alpha\|_2^2 - \alpha^T y - \sum_{i \in [p]} v(\alpha, \gamma_i), \quad (24)$$

where $v : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ is defined as follows:

$$v(\alpha, \gamma_i) := \left[\frac{(\alpha^T X_i - \gamma_i)^2}{4\lambda_2} - \lambda_0 \right]_+ + M|\gamma_i|. \quad (25)$$

Otherwise, if $\sqrt{\lambda_0/\lambda_2} \geq M$, a dual of Problem (11) is given by

$$\begin{aligned} \max_{\rho \in \mathbb{R}^n, \mu \in \mathbb{R}^p} h_2(\rho, \mu) &:= -\frac{1}{2}\|\rho\|_2^2 - \rho^T y - M\|\mu\|_1 \\ \text{s.t.} \quad |\rho^T X_i| - \mu_i &\leq \lambda_0/M + \lambda_2 M, \quad i \in [p]. \end{aligned} \quad (26)$$

Let β^* be an optimal solution to Problem (11) and define $r^* = y - X\beta^*$. The optimal dual variables for (24) are given by: $\alpha^* = -r^*$ and $\gamma_i^* = \mathbb{1}_{[|\beta_i^*|=M]}(\alpha^{*T} X_i - 2M\lambda_2 \text{sign}(\alpha^{*T} X_i))$, $i \in [p]$. Moreover, the optimal dual variables for (26) are: $\rho^* = -r^*$ and $\mu_i^* = \mathbb{1}_{[|\beta_i^*|=M]}(|\rho^{*T} X_i| - \lambda_0/M - \lambda_2 M)$, $i \in [p]$.

Note that strong duality holds for relaxation (11) since it satisfies Slater's condition [8], so the optimal objective of the dual in Theorem 2 matches that of (11).

Dual Feasible Solutions: Let $\hat{\beta}$ be an inexact primal solution obtained using Algorithm 2 and define $\hat{r} = y - X\hat{\beta}$. We discuss next how to construct a dual feasible solution using $\hat{\beta}$, i.e., without solving the dual in Theorem 2.

We first consider the case of $\sqrt{\lambda_0/\lambda_2} \leq M$. Here, we construct a dual solution $\hat{\alpha}, \hat{\gamma}$ for Problem (24) as follows:

$$\hat{\alpha} := -\hat{r} \quad \text{and} \quad \hat{\gamma} \in \arg \max_{\gamma \in \mathbb{R}^p} h_1(\hat{\alpha}, \gamma). \quad (27)$$

The choice $\hat{\alpha} = -\hat{r}$ is motivated by the optimality conditions in Theorem 2, while $\hat{\gamma}$ maximizes the dual objective (with α fixed to $\hat{\alpha}$). Note that the constructed solution is (trivially) feasible

since the dual in (24) is unconstrained. Since (24) is separable across the γ_i 's, $\hat{\gamma}_i$ is equivalently the solution of $\min_{\gamma_i \in \mathbb{R}} v(\hat{\alpha}, \gamma_i)$, which is an ℓ_1 -penalized convex optimization problem. It is easy to check that the corresponding solution is given by:

$$\hat{\gamma}_i = T(\hat{\alpha}^T X_i; 2M\lambda_2, \infty). \quad (28)$$

Thus, $\hat{\gamma}$ can be obtained in closed-form using (28). Since $\hat{\beta}$ is the output of Algorithm 2, we know that the corresponding set \mathcal{V} in Algorithm 2 is empty. Thus, by Proposition 3, we have $|\hat{r}^T X_i| \leq 2\sqrt{\lambda_0\lambda_2} \leq 2M\lambda_2$ for any $i \in \text{Supp}(\hat{\beta})^c$. Using the latter bound in (28), we get that $\hat{\gamma}_i = 0$ for any $i \in \text{Supp}(\hat{\beta})^c$. However, for $i \in \text{Supp}(\hat{\beta})$, $\hat{\gamma}_i$ can be potentially nonzero.

We now consider the second case: $\sqrt{\lambda_0/\lambda_2} \geq M$, where the dual is given by (26). Here, we construct a dual feasible solution $(\hat{\rho}, \hat{\mu})$ for (26) as follows:

$$\hat{\rho} := -\hat{r} \quad \text{and} \quad \hat{\mu} = \arg \max_{\mu \in \mathbb{R}^p} h_2(\hat{\rho}, \mu) \text{ s.t. } (\hat{\rho}, \mu) \text{ is feasible for (26)}. \quad (29)$$

Similar to (29), the choice $\hat{\rho} := -\hat{r}$ is motivated by the optimality conditions in Theorem 2, whereas the choice $\hat{\mu}$ maximizes the dual objective under the condition that $\rho = -\hat{r}$ (while ensuring feasibility). It can be readily seen that $\hat{\mu}_i$ is given in closed form by:

$$\hat{\mu}_i = \left[|\hat{\rho}^T X_i| - \lambda_0/M - \lambda_2 M \right]_+.$$

Note that for any $i \in \text{Supp}(\hat{\beta})^c$, we have $|\hat{\rho}^T X_i| \leq \lambda_0/M + \lambda_2 M$ (by Proposition 3), which implies that $\hat{\mu}_i = 0$. However, for $i \in \text{Supp}(\hat{\beta})$, $\hat{\mu}_i$ can be potentially nonzero.

Quality of the Dual Bounds: In Theorem 3, we quantify the tightness of the dual bounds obtained from the dual feasible solutions (27) and (29).

Theorem 3. *Let α^* , γ^* , ρ^* , and μ^* be the optimal dual variables defined in Theorem 2. Let β^* be an optimal solution to (11), and $\hat{\beta}$ be an inexact solution obtained using Algorithm 2. Define the primal gap $\epsilon = \|X(\beta^* - \hat{\beta})\|_2$. Moreover, let $k = \|\hat{\beta}\|_0$. Assuming M and λ_2 are fixed, the following holds for the dual solution $\hat{\alpha}, \hat{\gamma}$ defined in (27):*

$$h_1(\hat{\alpha}, \hat{\gamma}) \geq h_1(\alpha^*, \gamma^*) - k\mathcal{O}(\epsilon) - k\mathcal{O}(\epsilon^2), \quad (30)$$

and for the dual solution $\hat{\rho}, \hat{\mu}$ defined in (29):

$$h_2(\hat{\rho}, \hat{\mu}) \geq h_2(\rho^*, \mu^*) - k\mathcal{O}(\epsilon) - \mathcal{O}(\epsilon^2). \quad (31)$$

Interestingly, the bounds established in Theorem 3 do not depend on p , but rather on the support size k . Specifically, the constants in $\mathcal{O}(\epsilon)$ and $\mathcal{O}(\epsilon^2)$ only involve M and λ_2 (these constants are made explicit in the proof). In practice, we seek highly sparse solutions, i.e., $k \ll p$ —suggesting that the quality of the dual bounds deteriorates with k and not p . The main driver behind these tight bounds is Algorithm 2, which performs optimality checks on the coordinates outside the support. If vanilla CD was used instead of Algorithm 2, then the term k appearing in the bounds (30) and (31) will be replaced by p , making the bounds loose⁸.

Efficient Computation of the Dual Bounds: A direct computation of the dual bound $h_1(\hat{\alpha}, \hat{\gamma})$ or $h_2(\hat{\rho}, \hat{\mu})$ costs $\mathcal{O}(np)$ operations. Interestingly, we show that this cost can be reduced to $\mathcal{O}(n\|\hat{\beta}\|_0)$ (where we recall that $\hat{\beta}$ is a solution from Algorithm 2). First, we consider the case of $\sqrt{\lambda_0/\lambda_2} \leq M$, where the goal is to compute $h_1(\hat{\alpha}, \hat{\gamma})$. By Lemma 2 (in the appendix), we have $v(\hat{\alpha}, \hat{\gamma}_i) = 0$ for every $i \in \text{Supp}(\hat{\beta})^c$. Thus, $h_1(\hat{\alpha}, \hat{\gamma})$ can be simplified to:

$$h_1(\hat{\alpha}, \hat{\gamma}) = -\frac{1}{2}\|\hat{\alpha}\|_2^2 - \hat{\alpha}^T y - \sum_{i \in \text{Supp}(\hat{\beta})} v(\hat{\alpha}, \hat{\gamma}_i). \quad (32)$$

Now, we consider the case of $\sqrt{\lambda_0/\lambda_2} \geq M$, where the goal is to evaluate $h_2(\hat{\rho}, \hat{\mu})$. By construction, the solution (29) is dual feasible, which means that the constraints in (26) need not be checked when computing the bound. Moreover, $\hat{\mu}_i = 0$ for every $i \in \text{Supp}(\hat{\beta})^c$ (see the discussion after (29)). Thus, the dual bound can be expressed as follows:

$$h_2(\hat{\rho}, \hat{\mu}) = -\frac{1}{2}\|\hat{\rho}\|_2^2 - \hat{\rho}^T y - M \sum_{i \in \text{Supp}(\hat{\beta})} |\hat{\mu}_i|. \quad (33)$$

The expressions in (32) and (33) can be computed in $\mathcal{O}(n\|\hat{\beta}\|_0)$ operations.

3.3 Branching and Incumbents

Many branching strategies for BnB have been explored in the literature, e.g., random branching, strong branching, and pseudo-cost branching [7, 14, 3, 1]. Among these strategies, strong branching has proven to be very effective in minimizing the size of the search tree [3, 14, 7]. Strong branching selects the variable which leads to the maximum increase in the lower bounds of the children nodes. To select such a variable, two temporary node subproblems should be solved for every non-integral variable in the current relaxation. This can become a computational bottleneck, as each temporary subproblem involves solving a nonlinear optimization problem similar to (11). To address this challenge, we use a fast (approximate) version of strong branching, in which we restrict

⁸This holds by using the argument in the proof of Theorem 3 for Algorithm 1 instead of Algorithm 2.

the optimization in these temporary subproblems to the active set of the current node (instead of optimizing over all p variables). In practice, this leads to very similar search trees compared to exact strong branching, since the active set of the parent is typically close to the support of the children.

We obtain the initial upper bound using `L0Learn` [28], which uses coordinate descent and efficient local search algorithms to obtain feasible solutions for Problem (9). Moreover, at every node in the BnB tree, we attempt to improve the incumbent by exploiting the support of the node’s relaxation. Specifically, we solve the following ℓ_2 regularized least squares problem restricted to the relaxation’s support S : $\min_{\beta_S \in \mathbb{R}^{|S|}} \frac{1}{2} \|y - X_S \beta_S\|_2^2 + \lambda_2 \|\beta_S\|_2^2$. Since S is typically small, the problem above can be solved quickly using QR factorization-based solvers. Moreover, if S is similar to the support of the parent node, then the QR factorization for X_S can be efficiently computed by exploiting the QR factorization in the parent, instead of factorizing from scratch, e.g., see [26].

4 Experiments

We perform a series of high-dimensional experiments to study the run time of our BnB and compare to state-of-the-art approaches. While our dataset and parameter choices are well-grounded from a statistical perspective, we note that our goal here is not to study the statistical properties of ℓ_0 estimators. We refer the reader to [9, 12, 28] for empirical studies of the statistical properties.

4.1 Experimental Setup

Synthetic Data Generation: We generate a multivariate Gaussian data matrix with samples drawn from $MVN(0, \Sigma_{p \times p})$, a sparse coefficient vector $\beta^\dagger \in \mathbb{R}^p$ with k^\dagger equi-spaced nonzero entries all set to 1, and a noise vector $\epsilon_i \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$. The response is then obtained from the linear model $y = X\beta^\dagger + \epsilon$. We define the *signal-to-noise ratio (SNR)* as follows: $\text{SNR} = \text{Var}(X\beta^\dagger) / \sigma^2$. In all the experiments, we set σ^2 to achieve $\text{SNR} = 5$ —this is a relatively difficult setting which still allows for full support recovery, under suitable choices of n , p , and Σ (see [28] for a discussion on appropriate levels of SNR).

Warm Starts, λ_0 , λ_2 , and M : We obtain the warm start from `L0Learn`⁹ [28] and use it for all the MIP solvers considered. Unless otherwise specified, we fix λ_0 to a value λ_0^* which leads to a support size of k^\dagger (i.e., that of the true model). The parameters λ_2 and M can affect the run time significantly, so we study the sensitivity to these choices in our experiments. We consider choices of λ_2 that are relevant from a statistical perspective. We define λ_2^* as the λ_2 which minimizes the

⁹We use the default CD-based algorithm in `L0Learn`. We remind the reader that `L0Learn` is a local optimization method that does not provide certificates of global optimality.

ℓ_2 estimation error, among all k^\dagger -sparse solutions. More formally, for every λ_0, λ_2 , let $\beta(\lambda_0, \lambda_2)$ denote an optimal solution to Problem (1). We define λ_0^* and λ_2^* as a solution of:

$$\lambda_0^*, \lambda_2^* \in \arg \min_{\lambda_0, \lambda_2 \geq 0} \|\beta^\dagger - \beta(\lambda_0, \lambda_2)\|_2 \quad \text{s.t.} \quad \|\beta(\lambda_0, \lambda_2)\|_0 = k^\dagger. \quad (34)$$

We estimate λ_0^* and λ_2^* using `L0Learn` by doing a two-dimensional grid search over λ_0 and λ_2 ; in particular, $\lambda_2 \in [10^{-4}, 10]$ and the range for λ_0 is selected automatically by the toolkit. In the experiments, we report our λ_2 choices as a fraction or multiple of λ_2^* (e.g., $\lambda_2 = 0.1\lambda_2^*$). Moreover, we define M^* as the minimum M for which the hybrid formulation (9) is equivalent to (1). We estimate M^* as the ℓ_∞ norm of the solution obtained by `L0Learn`, and we report our choices in terms of M^* . Note that in almost all cases considered, the support of β^\dagger is correctly recovered by `L0Learn`, which yields high quality estimates for M^* and λ_2^* .

Solvers and Settings: Our solver, `L0BnB`¹⁰, is written in Python with critical code sections optimized using Numba [30]. We compare `L0BnB` with Gurobi and MOSEK on formulation (9). We also compare against [12] who solve the cardinality-constrained variant of (1); here we set the number of nonzeros to k^\dagger . For all approaches, we set the threshold for the relative optimality gap¹¹ to 1.0%. See the appendix for details on the computing setup.

4.2 Comparison with the State of the Art

In this section, we study the scalability of the different solvers in terms of the number of features p . We generate synthetic datasets with $n = 10^3$, $p \in \{10^3, 10^4, 10^5, 10^6\}$, and $k^\dagger = 10$. We consider a constant correlation setting, where $\Sigma_{ij} = 0.1 \forall i \neq j$ and 1 otherwise. The parameters λ_2 and M can have a significant effect on the run time. Thus, we report the timings for different choices of these parameters. In particular, in Table 2 (top panel), we fix $M = 1.5M^*$ and report the timings for the choices $\lambda_2 \in \{0.1\lambda_2^*, \lambda_2^*, 10\lambda_2^*\}$ (where λ_2^* and M^* are defined in Section 4.1). In Table 2 (bottom panel), we fix $\lambda_2 = \lambda_2^*$ and report the timings for $M \in \{M^*, 2M^*, 4M^*, \infty\}$.

The results in the top panel of Table 2 indicate significant speed-ups, reaching over 200,000x compared to Gurobi, 5000x compared to MOSEK, and 3600x compared to [12]. At $\lambda_2 = \lambda_2^*$, `L0BnB` is the only solver that can handle $p \geq 10^5$, and can, in fact, handle $p = 10^6$ in less than an hour. Recall that λ_2^* minimizes the ℓ_2 estimation error and leads to an estimator that is the closest to the ground truth. For $\lambda_2 = 0.1\lambda_2^*$, `L0BnB` is again the only solver that can handle $p \geq 10^5$. For $\lambda_2 = 10\lambda_2^*$, the optimization problem seems to become easier: `L0BnB` is up to 600x faster compared to λ_2^* , and [12] can handle up to 10^6 . However, `L0BnB` in this case, is ~ 100 times faster than [12] at $p = 10^6$. The speed-ups for $\lambda_2 = 10\lambda_2^*$ can be attributed to the fact that a larger λ_2 adds a large

¹⁰<https://github.com/alisaab/l0bnb>

¹¹Given the upper bound UB and lower bound LB, the relative optimality gap is defined as $(UB - LB)/UB$.

Table 2: (Sensitivity to λ_2 and M) Running time in seconds for solving (1) (via formulation (9)) by our proposal (L0BnB), Gurobi (GRB) and MOSEK (MSK). The method [12] solves the cardinality-constrained variant of (1). The symbols “*” or “-” mean that the method does not terminate in 2 hours. “***” means that the method terminates before 2 hours due to insufficient memory (30GB). Optimality gap is shown in parentheses for “*” and “***”, and is 100% for “-”. Choices of λ_2 and M are discussed in the text.

p	$\lambda_2 = \lambda_2^*$				$\lambda_2 = 0.1\lambda_2^*$				$\lambda_2 = 10\lambda_2^*$			
	L0BnB	GRB	MSK	[12]	L0BnB	GRB	MSK	[12]	L0BnB	GRB	MSK	[12]
10^3	6	95	216	1079	2	81	273	-	0.01	2372	54	0.4
10^4	14	-	5354	*(45%)	33	-	6693	-	0.8	-	1511	0.6
10^5	255	-	*(35%)	*(70%)	544	-	*(41%)	-	10	-	*(3%)	13
10^6	3468	-	-	*(88%)	*(7%)	-	-	-	43	-	-	4123

p	$M = M^*$			$M = 2M^*$			$M = 4M^*$			$M = \infty$		
	L0BnB	GRB	MSK	L0BnB	GRB	MSK	L0BnB	GRB	MSK	L0BnB	GRB	MSK
10^3	0.7	35	106	1	199	279	1	1636	259	1	2307	336
10^4	2	-	1909	21	-	6646	23	-	*(7%)	23	-	-
10^5	25	-	*(9%)	543	-	*(59%)	588	-	-	628	-	-
10^6	309	-	-	7180	-	-	** (3%)	-	-	** (3%)	-	-

amount of strong convexity to the objective (via the perspective term)—improving the performance of the relaxation solvers¹². It is worth emphasizing that our L0BnB is prototyped in Python; as opposed to the highly efficient BnB routines available in commercial solvers such as Gurobi and MOSEK.

Ideally, we desire a solver that can address (1) over a range of λ_2 values, which includes values in the neighborhood of λ_2^* . However, the results in Table 2 suggest that the state-of-the-art methods (except L0BnB) seem to only work for quite large values of λ_2 (which in this case, do not correspond to solutions that are interesting from a statistical viewpoint). On the other hand, L0BnB seems to be the only method that can scale to $p \sim 10^6$ while being relatively robust to the choice of λ_2 .

In the bottom panel of Table 2, the results also indicate that L0BnB significantly outperforms Gurobi and MOSEK for different choices of M . For all the solvers, the run time increases with M , and the longest run times are for $M = \infty$, which corresponds to the (pure) perspective formulation (3). This empirically validates our result in Proposition 2, where we show that for a sufficiently small value of M , the hybrid formulation can be better than a pure perspective formulation, in terms of the relaxation quality. However, even with $M = \infty$, L0BnB can still achieve a decent gap (3%) for $p = 10^6$ in less than two hours.

¹²Gurobi is the only exception to this observation. We investigated this: Gurobi generates additional cuts only for the case of $10\lambda_2^*$, which seem to slow down the relaxation solver.

4.3 Sensitivity to Data Parameters

We study how L0BnB’s running time is affected by the following data specific parameters: number of samples (n), feature correlations (Σ), and number of nonzero coefficients (k^\dagger). In the experiments below, we fix $\lambda_2 = \lambda_2^*$ and $M = 1.5M^*$.

Number of Samples: We fix $k^\dagger = 10$, $p = 10^4$, and consider a constant correlation setting $\Sigma_{ij} = 0.1$ for $i \neq j$ and 1 otherwise¹³. The timings for $n \in \{10^2, 10^3, 10^4, 10^5\}$ are in Table 3. The results indicate that the problem can be solved in reasonable times (order of seconds to minutes) for all values of n . The problems can be solved the fastest when n is close to p , and the extreme cases $n = 10^2$ and $n = 10^5$ are the slowest. We contend that for large n , the CD updates (which cost $\mathcal{O}(n)$ each) become a bottleneck—this is verified by the (roughly) 10x increase in run time between $n = 10^4$ and $n = 10^5$. For $n = 10^2$, the CD updates are cheaper, however, the size of the search tree is significantly larger—suggesting that a small value of n can lead to loose relaxations and require more branching. Note also that the underlying statistical problem is the most difficult for $n = 10^2$ compared to other larger values of n .

Feature Correlations: We generate datasets with $n = 10^3$, $p = 10^5$, $k^\dagger = 10$. In this setup, our experiments indicate that instances with a constant correlation matrix with $\Sigma_{ij} = 0.4$ (for $i \neq j$) can be solved in ~ 18 seconds, however, the run time and memory requirements significantly increase after the 0.4 threshold (in fact, our machine ran out of memory after around an hour for a correlation of 0.45). Here we consider a more tractable setting where Σ has a block diagonal structure: $\Sigma = \text{Diag}(\Sigma^{(1)}, \Sigma^{(2)}, \dots, \Sigma^{(k^\dagger)})$ —such block structures are commonly used in the sparse learning literature as support recovery may not be possible with constant correlations, e.g., see [46]. Given a correlation parameter $\rho \in (0, 1)$, for each $l \in [k^\dagger]$, we assume that $\Sigma_{ij}^{(l)} = \rho^{|i-j|}$ for any $i, j \in [p/k^\dagger]$, i.e., the correlation is exponentially decaying in each block. We report the timings for different values of the parameter ρ in Table 3. The results indicate that higher correlations lead to an increase in the run time, roughly a 10x increase from $\rho = 0.1$ to $\rho = 0.8$. However, for $\rho = 0.9$, L0BnB runs out of memory (after around an hour) with a gap of 16%.

Sparsity of the Regression Coefficients: We consider datasets with $n = 10^3$, $p = 10^4$, and the same correlation setting as the experiment for the number of samples. We increase k^\dagger until L0BnB runs out of memory. The results in Table 3 show that problems with 20 nonzeros can be handled in reasonable time (360 seconds). After this threshold, more memory would be required for L0BnB to terminate. We also note that for larger λ_2 or tighter M choices, larger values of k^\dagger can be handled.

¹³We choose $p = 10^4$ to ensure that the matrix fits into memory when n is large.

Table 3: Run time in seconds (denoted by $t(s)$) for L0BnB. The symbol “*” means that L0BnB runs out of memory before a 2hr time limit (optimality gap is shown in parenthesis).

Varying n					Varying correlation coefficient ρ							Varying k^\dagger					
n	10^2	10^3	10^4	10^5	ρ	0.1	0.3	0.5	0.7	0.8	0.9	k^\dagger	5	10	15	20	25
$t(s)$	216	14	32	319	$t(s)$	110	115	145	361	1126	*(16%)	$t(s)$	1	14	139	360	*(16%)

4.4 Real Data and Ablation Studies

High-dimensional Real Data: Here we investigate the run time of L0BnB on the Riboflavin dataset [15]—a genetics dataset used for predicting Vitamin B2 production levels. The original dataset has $p = 4088$ and $n = 71$. We augment the dataset with pairwise feature interactions to get $p = 8,362,004$. On the augmented data, we run 5-fold cross-validation in L0Learn (with default parameters) to find the optimal regularization parameters $\hat{\lambda}_0$ and $\hat{\lambda}_2$. We set M to be 10 times the ℓ_∞ norm of the solution obtained from cross-validation. In L0BnB, we solve the problem for $\lambda_2 \in \{0.1\hat{\lambda}_2, \hat{\lambda}_2\}$ and vary λ_0 to generate solutions of different support sizes. The run time, for each λ_2 , as a function of the support size is reported in Figure 1. Interestingly, at $\hat{\lambda}_2$, all the support sizes obtained (up to 15 nonzeros) can be handled in less than a minute. Moreover, the increase in time is relatively slow as the number of nonzeros increases. When λ_2 becomes smaller (i.e., $\lambda_2 = 0.1\hat{\lambda}_2$), the run times increase (though they are reasonable) as the problem becomes more difficult. When an optimal solution has six nonzeros, L0BnB for $\lambda_2 = 0.1\lambda_2^*$ is approximately 20x slower than $\lambda_2 = \lambda_2^*$.

Ablation Study: We perform an ablation study to measure the effect of key choices in our BnB on the run time. Particularly, we consider the following changes: (i) replacing our relaxation solver with MOSEK, (ii) turning off warm starts in our solver, and (iii) turning off active sets in our solver. We measure the run time before and after these changes on synthetic data with $n = 10^3$, k^\dagger , $\Sigma_{ij} = 0.1$ for all $i \neq j$ and 1 otherwise, and with $\lambda_2 = \lambda_2^*$ and $M = 1.5M^*$. The run times for different choices of p are reported in Table 4. The results show that replacing our relaxation solver with MOSEK will slow down the BnB by more than 500x at $p = 10^5$. The results for MOSEK are likely to be even slower for $p = 10^6$ as it still had a 100% gap after 2 hours. We note that MOSEK employs a state-of-the-art conic optimization solver (based on an interior point method). The significant speed-ups here can be attributed to our CD-based algorithm, which is designed to effectively exploit the sparsity structure in the problem. The results also indicate that warm starts and active sets are important for run time, e.g., removing warm starts and active sets at $p = 10^5$ can slow down the algorithm by more than 5x and 23x, respectively.

Table 4: Time (seconds) after the following changes to our relaxation solver: (i) replacing it with MOSEK, (ii) removing warm starts, and (iii) removing active sets. “*”: Does not terminate in 2 hours (opt. gap in parenthesis).

p	10^4	10^5	10^6
L0BnB	14	255	3468
(i)	*(18%)	*(24%)	*(100%)
(ii)	57	1389	*(10%)
(iii)	174	5964	*(15%)

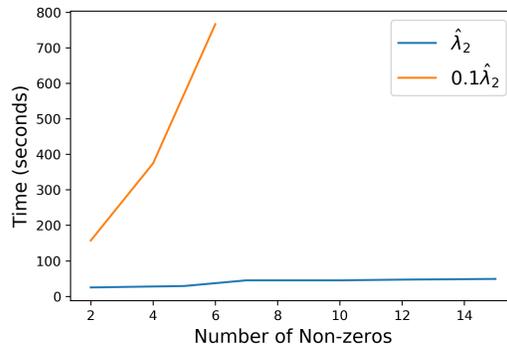


Figure 1: L0BnB run times on a real genomics dataset (Riboflavin) $n = 71, p \approx 8.3 \times 10^6$.

5 Conclusion

We considered the exact computation of estimators from the $\ell_0\ell_2$ -regularized least squares problem. We developed a highly specialized nonlinear BnB framework for solving the problem. Unlike the state-of-the-art MIP solvers, which rely on primal-dual relaxation solvers, we designed a primal first-order method. The method consists of coordinate descent along with active set updates and gradient screening, which exploit the information shared across the search tree to reduce the coordinate update complexity. Moreover, we proposed a new method for obtaining dual bounds from the primal coordinate descent solutions and showed that the quality of these bounds depends on the sparsity level, rather than the number of features p . Our experiments on both real and synthetic data indicate that our method is over 3600x compared to the fastest solvers, handling high-dimensional instances with $p = 8.3 \times 10^6$ in the order of seconds to few minutes. Moreover, the results show that our method is more robust to the choices of the regularization parameters and can handle difficult statistical problems (e.g., relatively high correlations or small number of samples). Our work demonstrates that carefully designed first-order methods can be highly effective within a BnB framework; and can perhaps, be applied to more general mixed integer programs involving sparsity.

Acknowledgements

Hussein acknowledges research funding from the Office of Naval Research ONR-N000141812298. Rahul acknowledges research funding from the Office of Naval Research ONR-N000141812298 (Young Investigator Award), the National Science Foundation (NSF-IIS-1718258), and IBM.

References

- [1] Achterberg, T., Koch, T., Martin, A.: Branching rules revisited. *Operations Research Letters* **33**(1), 42–54 (2005)
- [2] Andersen, E.D., Roos, C., Terlaky, T.: On implementing a primal-dual interior-point method for conic quadratic optimization. *Mathematical Programming* **95**(2), 249–277 (2003)
- [3] Applegate, D., Bixby, R., Cook, W., Chvátal, V.: On the solution of traveling salesman problems (1998)
- [4] Atamturk, A., Gomez, A.: Rank-one convexification for sparse regression. arXiv preprint arXiv:1901.10334 (2019)
- [5] Beck, A., Eldar, Y.C.: Sparsity constrained nonlinear optimization: Optimality conditions and algorithms. *SIAM Journal on Optimization* **23**(3), 1480–1509 (2013). DOI 10.1137/120869778. URL <https://doi.org/10.1137/120869778>
- [6] Beck, A., Tretuashvili, L.: On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization* **23**(4), 2037–2060 (2013). DOI 10.1137/120887679. URL <http://dx.doi.org/10.1137/120887679>
- [7] Belotti, P., Kirches, C., Leyffer, S., Linderoth, J., Luedtke, J., Mahajan, A.: Mixed-integer nonlinear optimization. *Acta Numerica* **22**, 1–131 (2013)
- [8] Bertsekas, D.: *Nonlinear Programming*. Athena scientific optimization and computation series. Athena Scientific (2016). URL <https://books.google.com/books?id=Tw0ujgEACAAJ>
- [9] Bertsimas, D., King, A., Mazumder, R., et al.: Best subset selection via a modern optimization lens. *The Annals of Statistics* **44**(2), 813–852 (2016)
- [10] Bertsimas, D., Pauphilet, J., Van Parys, B.: Sparse classification: a scalable discrete optimization perspective. arXiv preprint arXiv:1710.01352 (2017)
- [11] Bertsimas, D., Pauphilet, J., Van Parys, B.: Sparse regression: Scalable algorithms and empirical performance. arXiv preprint arXiv:1902.06547 (2019)
- [12] Bertsimas, D., Van Parys, B.: Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. arXiv preprint arXiv:1709.10029 (2017)
- [13] Blumensath, T., Davies, M.E.: Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis* **27**(3), 265–274 (2009)

- [14] Bonami, P., Lee, J., Leyffer, S., Wächter, A.: More branch-and-bound experiments in convex nonlinear integer programming. Preprint ANL/MCS-P1949-0911, Argonne National Laboratory, Mathematics and Computer Science Division (2011)
- [15] Bühlmann, P., Kalisch, M., Meier, L.: High-dimensional statistics with a view toward applications in biology (2014)
- [16] Dakin, R.J.: A tree-search algorithm for mixed integer programming problems. *The computer journal* **8**(3), 250–255 (1965)
- [17] David, G., Ilias, Z.: High dimensional regression with binary coefficients. estimating squared error and a phase transtition. In: *Conference on Learning Theory*, pp. 948–953 (2017)
- [18] Dedieu, A., Hazimeh, H., Mazumder, R.: Learning sparse classifiers: Continuous and mixed integer optimization perspectives. *arXiv preprint arXiv:2001.06471* (2020)
- [19] Dong, H., Chen, K., Linderoth, J.: Regularization vs. Relaxation: A conic optimization perspective of statistical variable selection. *ArXiv e-prints* (2015)
- [20] Duran, M.A., Grossmann, I.E.: An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical programming* **36**(3), 307–339 (1986)
- [21] Frangioni, A., Gentile, C.: Perspective cuts for a class of convex 0–1 mixed integer programs. *Mathematical Programming* **106**(2), 225–236 (2006)
- [22] Friedman, J., Hastie, T., Tibshirani, R.: Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software* **33**(1), 1–22 (2010). URL <http://www.jstatsoft.org/v33/i01/>
- [23] Greenshtein, E., et al.: Best subset selection, persistence in high-dimensional statistical learning and optimization under l1 constraint. *The Annals of Statistics* **34**(5), 2367–2386 (2006)
- [24] Günlük, O., Linderoth, J.: Perspective reformulations of mixed integer nonlinear programs with indicator variables. *Mathematical programming* **124**(1-2), 183–205 (2010)
- [25] Gurobi Optimization, I.: Gurobi optimizer reference manual. URL <http://www.gurobi.com> (2020)
- [26] Hammarling, S., Lucas, C.: Updating the qr factorization and the least squares problem. Manchester Institute for Mathematical Sciences, University of Manchester (2008)
- [27] Hastie, T., Tibshirani, R., Tibshirani, R.J.: Extended comparisons of best subset selection, forward stepwise selection, and the lasso. *arXiv preprint arXiv:1707.08692* (2017)

- [28] Hazimeh, H., Mazumder, R.: Fast Best Subset Selection: Coordinate Descent and Local Combinatorial Optimization Algorithms. ArXiv e-prints (2018)
- [29] Hazimeh, H., Mazumder, R.: Learning hierarchical interactions at scale: A convex optimization approach. arXiv preprint arXiv:1902.01542 (2019)
- [30] Lam, S.K., Pitrou, A., Seibert, S.: Numba: A llvm-based python jit compiler. In: Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, pp. 1–6 (2015)
- [31] Lee, J., Leyffer, S.: Mixed integer nonlinear programming, vol. 154. Springer Science & Business Media (2011)
- [32] Lubin, M., Yamangil, E., Bent, R., Vielma, J.P.: Extended formulations in mixed-integer convex programming. In: International Conference on Integer Programming and Combinatorial Optimization, pp. 102–113. Springer (2016)
- [33] Mazumder, R., Friedman, J.H., Hastie, T.: Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association* **106**(495), 1125–1138 (2011). DOI 10.1198/jasa.2011.tm09738. URL <https://doi.org/10.1198/jasa.2011.tm09738>. PMID: 25580042
- [34] Mazumder, R., Radchenko, P., Dedieu, A.: Subset selection with shrinkage: Sparse linear modeling when the snr is low. arXiv preprint arXiv:1708.03288 (2017)
- [35] Natarajan, B.K.: Sparse approximate solutions to linear systems. *SIAM journal on computing* **24**(2), 227–234 (1995)
- [36] Nesterov, Y.: Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization* **22**(2), 341–362 (2012). DOI 10.1137/100802001. URL <https://doi.org/10.1137/100802001>
- [37] Owen, A.B.: A robust hybrid of lasso and ridge regression. *Contemporary Mathematics* **443**(7), 59–72 (2007)
- [38] Pilanci, M., Wainwright, M.J., El Ghaoui, L.: Sparse learning via boolean relaxations. *Mathematical Programming* **151**(1), 63–87 (2015)
- [39] Raskutti, G., Wainwright, M.J., Yu, B.: Minimax rates of estimation for high-dimensional linear regression over l_q -balls. *IEEE transactions on information theory* **57**(10), 6976–6994 (2011)

- [40] Studio-CPLEX, I.I.C.O.: Users manual-version 12 release 6. IBM ILOG CPLEX Division: Incline Village, NV, USA (2013)
- [41] Tawarmalani, M., Sahinidis, N.V.: A polyhedral branch-and-cut approach to global optimization. *Mathematical programming* **103**(2), 225–249 (2005)
- [42] Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **58**(1), 267–288 (1996)
- [43] Vielma, J.P., Ahmed, S., Nemhauser, G.L.: A lifted linear programming branch-and-bound algorithm for mixed-integer conic quadratic programs. *INFORMS Journal on Computing* **20**(3), 438–450 (2008)
- [44] Xie, W., Deng, X.: Scalable algorithms for the sparse ridge regression (2020)
- [45] Zhang, Y., Wainwright, M.J., Jordan, M.I.: Lower bounds on the performance of polynomial-time algorithms for sparse linear regression. In: *Conference on Learning Theory*, pp. 921–948 (2014)
- [46] Zhao, P., Yu, B.: On model selection consistency of lasso. *Journal of Machine learning research* **7**(Nov), 2541–2563 (2006)
- [47] Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)* **67**(2), 301–320 (2005)

A Appendix: Proofs of Technical Results

Proof of Proposition 1: The interval relaxation of (2) is given by:

$$v_1 = \min_{\|\beta\|_\infty \leq M} g_1(\beta) := \frac{1}{2} \|y - X\beta\|_2^2 + \frac{\lambda_0}{M} \sum_{i \in [p]} |\beta_i| + \lambda_2 \beta_i^2. \quad (35)$$

The interval relaxation of (3) is (see [19] for a proof):

$$v_2 = \min_{\beta \in \mathbb{R}^p} g_2(\beta) := \frac{1}{2} \|y - X\beta\|_2^2 + \sum_{i \in [p]} 2\lambda_0 \mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0}), \quad (36)$$

where \mathcal{B} is defined in (10). Similarly, the interval relaxation of (4) can be written as follows:

$$v_3 = \min_{\|\beta\|_\infty \leq M} g_3(\beta) := \frac{1}{2} \|y - X\beta\|_2^2 + \sum_{i \in [p]} 2\lambda_0 \mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0}), \quad (37)$$

Define a function $t : \mathbb{R} \rightarrow \mathbb{R}$ as follows $t(\beta_i) = 2\lambda_0 \mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0}) - \frac{\lambda_0}{M} |\beta_i| - \lambda_2 \beta_i^2$. First, we prove (6).

Proof of (6): Suppose that $M \leq \frac{1}{2} \sqrt{\lambda_0/\lambda_2}$ and $|\beta_i| \leq M$. Then we have $|\beta_i| \leq \sqrt{\lambda_0/\lambda_2}$, and, therefore by the definition of \mathcal{B} , we have $\mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0}) = \sqrt{\lambda_2/\lambda_0} |\beta_i|$. This leads to:

$$t(\beta_i) = \left(2\sqrt{\lambda_0\lambda_2} - \frac{\lambda_0}{M}\right) |\beta_i| - \lambda_2 \beta_i^2 \leq 0, \quad (38)$$

where the inequality follows since $2\sqrt{\lambda_0\lambda_2} - \lambda_0/M \leq 0$ for $M \leq \frac{1}{2} \sqrt{\lambda_0/\lambda_2}$. Now, let β^\dagger be an optimal solution to (35). Then, the following holds:

$$v_1 - v_3 \geq g_1(\beta^\dagger) - g_3(\beta^\dagger) = - \sum_{i \in [p]} t(\beta_i^\dagger) \geq 0. \quad (39)$$

where the first inequality follows from $v_1 = g_1(\beta^\dagger)$ and $v_3 \leq g_3(\beta^\dagger)$ (which holds since β^\dagger is a feasible solution for (37)). The second inequality above follows by (38). Moreover, since $v_2 \leq v_3$, we also get that $v_1 \geq v_2$, which establishes (6).

To show (7) and (8), we first introduce the following useful lemma.

Lemma 1. *The following holds: $t(\beta_i) \geq 0$ for any $\beta_i \in \{\beta_i \mid |\beta_i| \leq M \text{ and } M \geq \sqrt{\lambda_0/\lambda_2}\}$.*

Proof of Lemma 1: Suppose that $M \geq \sqrt{\lambda_0/\lambda_2}$ and $|\beta_i| \leq M$. There are two cases to consider here: Case (i): $|\beta_i| \leq \sqrt{\lambda_0/\lambda_2}$ and Case (ii): $|\beta_i| \geq \sqrt{\lambda_0/\lambda_2}$.

For Case (i), from the definition of \mathcal{B} , we have $2\lambda_0 \mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0}) = 2\sqrt{\lambda_0\lambda_2} |\beta_i|$. Therefore,

$$t(\beta_i) = \left(2\sqrt{\lambda_0\lambda_2} - \frac{\lambda_0}{M}\right) |\beta_i| - \lambda_2 \beta_i^2.$$

In the above, it is easy to check that for $M \geq \sqrt{\lambda_0/\lambda_2}$ and $|\beta_i| \leq \sqrt{\lambda_0/\lambda_2}$, $t(\beta_i) \geq 0$. Now for Case (ii), we have $2\lambda_0\mathcal{B}(\beta_i\sqrt{\lambda_2/\lambda_0}) = \lambda_0 + \lambda_2\beta_i^2$, which leads to

$$t(\beta_i) = \lambda_0 \left(1 - \frac{|\beta_i|}{M}\right) \geq 0,$$

where the above is non-negative since we assume that $|\beta_i| \leq M$. This establishes lemma 1.

Proof of (7): Now let β^* be an optimal solution to (37) and assume that $M \geq \sqrt{\lambda_0/\lambda_2}$. Then, the following holds:

$$v_3 - v_1 \geq g_3(\beta^*) - g_1(\beta^*) = \sum_{i \in [p]} t(\beta_i^*) \geq 0. \quad (40)$$

where the first inequality holds since β^* is a feasible solution to (35) so $v_1 \leq g_1(\beta^*)$, and the last inequality is due to Lemma 1. This establishes (7).

Proof of (8): Suppose that $M \geq \sqrt{\lambda_0/\lambda_2}$ and $\lambda_2 \geq \lambda_2^*$, and let $\hat{\beta}$ be a solution in $\mathcal{S}(\lambda_2)$ (i.e., $\hat{\beta}$ is optimal for (36)), which satisfies $\|\hat{\beta}\|_\infty \leq M$. Since $v_2 \leq v_3$ and $\hat{\beta}$ is feasible for (37), then $\hat{\beta}$ is optimal for (37), i.e., $v_3 = g_2(\hat{\beta})$. But by (7) we have $v_3 \geq v_1$, which combined with $v_3 = g_2(\hat{\beta})$, leads to $g_2(\hat{\beta}) \geq v_1$. This which establishes (8) (since by definition $g_2(\hat{\beta}) = v_2$).

Proof of Theorem 1: The interval relaxation of (9) can be expressed as:

$$\min_{\beta} \left\{ \frac{1}{2} \|y - X\beta\|_2^2 + \sum_{i \in [p]} \min_{s_i, z_i} (\lambda_0 z_i + \lambda_2 s_i) \right\} \quad (41a)$$

$$s_i z_i \geq \beta_i^2, \quad i \in [p] \quad (41b)$$

$$-M z_i \leq \beta_i \leq M z_i, \quad i \in [p] \quad (41c)$$

$$z_i \in [0, 1], s_i \in \mathbb{R}_{\geq 0}, \quad i \in [p] \quad (41d)$$

Let $\omega(\beta_i; \lambda_0, \lambda_2, M) := \min_{s_i, z_i} (\lambda_0 z_i + \lambda_2 s_i)$ s.t. (41b), (41c), (41d). Next, we will derive the solutions of the inner minimization problem in (41a), whose objective is $\omega(\beta_i; \lambda_0, \lambda_2, M)$. At any optimal solution z_i^* , constraints (41b) and (41c) imply that $z_i^* = \max\{\frac{(\beta_i)^2}{s_i}, \frac{|\beta_i|}{M}\}$. Plugging the expression of z_i^* into $\omega(\beta_i; \lambda_0, \lambda_2, M)$ we arrive to the following equivalent problem:

$$\omega(\beta_i; \lambda_0, \lambda_2, M) = \min_{s_i} \max \left\{ \underbrace{\lambda_0 \frac{\beta_i^2}{s_i} + \lambda_2 s_i}_{\text{Case I}}, \underbrace{\lambda_0 \frac{|\beta_i|}{M} + \lambda_2 s_i}_{\text{Case II}} \right\} \quad \text{s.t. } s_i \geq \beta_i^2, s_i \geq 0 \quad (42)$$

Suppose that Case I attains the maximum in (42). This holds if $s_i \leq |\beta_i|M$. The function $\lambda_0 \frac{\beta_i^2}{s_i} + \lambda_2 s_i$ is convex in s_i , and the first-order optimality conditions imply that the optimal solution

is $s_i^* = |\beta_i| \sqrt{\lambda_0/\lambda_2}$ if $|\beta_i| \leq \sqrt{\lambda_0/\lambda_2} \leq M$ and $s_i^* = \beta_i^2$ if $\sqrt{\lambda_0/\lambda_2} \leq |\beta_i| \leq M$. Plugging s_i^* into $\omega(\beta_i; \lambda_0, \lambda_2, M)$ leads to $\omega(\beta_i; \lambda_0, \lambda_2, M) = 2\lambda_0 \mathcal{B}(\beta_i \sqrt{\lambda_2/\lambda_0})$, assuming $\sqrt{\lambda_0/\lambda_2} \leq M$.

Now suppose Case II attains the maximum in (42). This holds if $s_i \geq |\beta_i|M$. The function $\lambda_0 \frac{|\beta_i|}{M} + \lambda_2 s_i$ is monotonically increasing in s_i , where s_i is lower bounded by $|\beta_i|M$ and β_i^2 . But we always have $|\beta_i|M \geq \beta_i^2$ (since $|\beta_i| \leq M$), which implies that the optimal solution is $s_i^* = |\beta_i|M$. Substituting s_i^* into (42) we get $\omega(\beta_i; \lambda_0, \lambda_2, M) = (\lambda_0/M + \lambda_2 M)|\beta_i|$, and this holds as long as $\sqrt{\lambda_0/\lambda_2} \geq M$.

Proof of Proposition 2: We first show (14). Note that v_1 is given by Problem (35). The following holds for $\sqrt{\lambda_0/\lambda_2} \geq M$:

$$\begin{aligned} v_4 - v_1 &\geq F(\beta^*) - g_1(\beta^*) = \sum_{i \in [p]} \left\{ \psi_2(\beta_i^*; \lambda_0, \lambda_2, M) - \frac{\lambda_0}{M} |\beta_i^*| - \lambda_2 (\beta_i^*)^2 \right\} \\ &= \lambda_2 \sum_{i \in [p]} (M |\beta_i^*| - (\beta_i^*)^2) \end{aligned}$$

where the inequality holds since β^* is a feasible solution to (35). This establishes (14).

We now show (15). Since $|\beta_i^*| \leq M$ and $\sqrt{\lambda_0/\lambda_2} \geq M$, we must have $\psi_1(\beta_i^*; \lambda_0, \lambda_2) = 2|\beta_i^*| \sqrt{\lambda_0 \lambda_2}$ (this corresponds to the first case in (10)). The following then holds:

$$\begin{aligned} v_4 - v_3 &\geq F(\beta^*) - G(\beta^*) = \sum_{i \in [p]} \left\{ \psi_2(\beta_i^*; \lambda_0, \lambda_2, M) - \psi_1(\beta_i^*; \lambda_0, \lambda_2) \right\} \\ &= h(\lambda_0, \lambda_2, M) \|\beta^*\|_1, \end{aligned} \tag{43}$$

where the inequality holds since β^* is feasible for (13). The first inequality in (15) follows from (43); and the second inequality in (15) follows from the observation that $v_2 \leq v_3$.

Proof of Proposition 3: Fix some $i \in \text{Supp}(\hat{\beta})^c$. Recall that $\min_{|\beta_i| \leq M} F(\hat{\beta}_1, \dots, \beta_i, \dots, \hat{\beta}_p)$ is equivalent to Problem (17), where $\tilde{\beta}_i = \langle y - \sum_{j \neq i} X_j \hat{\beta}_j, X_i \rangle$. Since $\hat{\beta}_i = 0$, we have $\tilde{\beta}_i = \langle \hat{r}, X_i \rangle$. For $\sqrt{\lambda_0/\lambda_2} \leq M$, (18) implies that the solution of (17) is nonzero iff $|\tilde{\beta}_i| > 2\sqrt{\lambda_0 \lambda_2}$. Similarly, for $\sqrt{\lambda_0/\lambda_2} \geq M$, by (18), the solution of (17) is nonzero iff $|\tilde{\beta}_i| > \lambda_0/M + \lambda_2 M$. Using these observations in the definition of \mathcal{V} in Algorithm 2, leads to the result of the proposition.

Proof of Proposition 4: Fix some $i \in \mathcal{V}$. Note that

$$|\langle \hat{r}, X_i \rangle - \langle r^0, X_i \rangle| = |\langle X\beta^0 - X\hat{\beta}, X_i \rangle| \leq \|X\beta^0 - X\hat{\beta}\|_2 \|X_i\|_2 \leq \epsilon.$$

Using the triangle inequality and the bound above, we get:

$$|\langle \hat{r}, X_i \rangle| \leq |\langle r^0, X_i \rangle| + |\langle \hat{r}, X_i \rangle - \langle r^0, X_i \rangle| \leq |\langle r^0, X_i \rangle| + \epsilon.$$

Therefore, if $i \in \mathcal{V}$, i.e., $|\langle \hat{r}, X_i \rangle| > c(\lambda_0, \lambda_2, M)$, then $|\langle r^0, X_i \rangle| + \epsilon > c(\lambda_0, \lambda_2, M)$, implying that $i \in \hat{\mathcal{V}}$.

Sketch of Proof for Theorem 2: Problem (11) can be written equivalently as:

$$\min_{\beta \in \mathbb{R}^p, r \in \mathbb{R}^n} \frac{1}{2} \|r\|_2^2 + \sum_{i \in [p]} \psi(\beta_i; \lambda_0, \lambda_2, M) \quad \text{s.t.} \quad r = y - X\beta, \quad |\beta_i| \leq M, \quad \forall i \in [p] \quad (44)$$

Case of $\sqrt{\lambda_0/\lambda_2} \leq M$: We first consider the case when $\sqrt{\lambda_0/\lambda_2} \leq M$. We dualize all the constraints, leading to the following Lagrangian dual:

$$\max_{\alpha \in \mathbb{R}^n, \eta \in \mathbb{R}_{\geq 0}^p} \min_{\beta \in \mathbb{R}^p, r \in \mathbb{R}^n} L(\beta, r, \alpha, \eta) \quad (45)$$

where $L(\beta, r, \alpha, \eta)$ is the Lagrangian function defined as follows:

$$L(\beta, r, \alpha, \eta) := \frac{1}{2} \|r\|_2^2 + \sum_{i \in [p]} \psi(\beta_i; \lambda_0, \lambda_2, M) + \alpha^T (r - y + X\beta) + \sum_{i \in [p]} \eta_i (|\beta_i| - M).$$

Solving the inner minimization problem in (45) leads to:

$$\max_{\alpha \in \mathbb{R}^n, \eta \in \mathbb{R}_{\geq 0}^p} -\frac{1}{2} \|\alpha\|_2^2 - \alpha^T y - \sum_{i=1}^p \left[\frac{(|\alpha^T X_i| - \eta_i)^2}{4\lambda_2} - \lambda_0 \right]_+ - \sum_{i=1}^p M\eta_i. \quad (46)$$

Note we can replace the term $M\eta_i$ in (46) by $M|\eta_i|$, and drop the constraint $\eta \geq 0$. We then perform a change of variable $\eta_i = \gamma_i \text{sign}(\alpha^T X_i)$, where $\gamma_i \in \mathbb{R}$. Therefore, we have $(|\alpha^T X_i| - \eta_i)^2 = (\alpha^T X_i - \gamma_i)^2$ which leads to the dual formulation in (24).

Note that $r^* = \arg \min_r L(\beta^*, r, \alpha^*, \eta^*) = -\alpha^*$. Moreover, by complementary slackness, if $|\beta_i^*| < M$ then $\eta_i^* = 0$. Thus, $\gamma_i^* = 0$ for every i for which $|\beta_i^*| < M$. Finally, note that if $|\beta_i^*| = M$, then we have $\gamma^* = \arg \max_{\gamma} h_1(\alpha^*, \gamma)$, and it is easy to check that the solution of the latter problem is given by $\gamma_i^* = \alpha^{*T} X_i - 2M\lambda_2 \text{sign}(\alpha^{*T} X_i)$.

Case of $\sqrt{\lambda_0/\lambda_2} \geq M$: For this case, the dual can be derived using the same steps outlined above (i.e., dualize all the constraints in (44) and then simplify the resulting dual formulation to eliminate the non-negativity constraints on the dual variables).

The following lemma is useful for proving Theorem 3.

Lemma 2. Suppose $\sqrt{\lambda_0/\lambda_2} \leq M$. Let $\hat{\beta}$ be a solution from Algorithm 2 and $(\hat{\alpha}, \hat{\gamma})$ be the corresponding dual solution defined in (27). Let β^* and r^* be as defined in Theorem 2, and define the primal gap $\epsilon = \|X(\beta^* - \hat{\beta})\|_2$. Then, for every $i \in \text{Supp}(\hat{\beta})^c$, we have $v(\hat{\alpha}, \hat{\gamma}_i) = 0$; and for every $i \in \text{Supp}(\hat{\beta})$, we have

$$v(\hat{\alpha}, \hat{\gamma}_i) \leq c_i \epsilon + (4\lambda_2)^{-1} \epsilon^2 + v(\alpha^*, \gamma_i^*), \quad (47)$$

where $c_i = (2\lambda_2)^{-1}$ if $|\beta_i^*| < M$, and $c_i = M$ if $|\beta_i^*| = M$.

Proof of Lemma 2: Fix some $i \in \text{Supp}(\hat{\beta})^c$. Since $\hat{\beta}$ is the output of Algorithm 2, we have: $|\hat{r}^T X_i| \leq 2\sqrt{\lambda_0 \lambda_2}$ and $\hat{\gamma}_i = 0$. Thus,

$$(\hat{\alpha}^T X_i - \hat{\gamma}_i)^2 / (4\lambda_2) = (\hat{r}^T X_i)^2 / (4\lambda_2) \leq \lambda_0,$$

which implies that $v(\hat{\alpha}, \hat{\gamma}_i) = 0$.

Now fix some $i \in \text{Supp}(\hat{\beta})$, and let $a := (\hat{\alpha} - \alpha^*)^T X_i$ and $b := \alpha^{*T} X_i - \gamma_i^*$. By (27), we have $\hat{\gamma}_i = \arg \min_{\gamma_i} v(\hat{\alpha}, \gamma_i)$, which leads to $v(\hat{\alpha}, \hat{\gamma}_i) \leq v(\hat{\alpha}, \gamma_i^*)$. An upper bound on $v(\hat{\alpha}, \hat{\gamma}_i)$ can be then obtained as follows:

$$\begin{aligned} v(\hat{\alpha}, \hat{\gamma}_i) &\leq v(\hat{\alpha}, \gamma_i^*) = \left[\frac{(\hat{\alpha}^T X_i - \gamma_i^*)^2}{4\lambda_2} - \lambda_0 \right]_+ + M|\gamma_i^*| \\ &= \left[\frac{((\hat{\alpha} - \alpha^*)^T X_i + \alpha^{*T} X_i - \gamma_i^*)^2}{4\lambda_2} - \lambda_0 \right]_+ + M|\gamma_i^*| \\ &= \left[\frac{a^2 + 2ab + b^2}{4\lambda_2} - \lambda_0 \right]_+ + M|\gamma_i^*| \\ &\leq \frac{a^2 + 2|a||b|}{4\lambda_2} + \left[\frac{b^2}{4\lambda_2} - \lambda_0 \right]_+ + M|\gamma_i^*| \\ &\leq \frac{a^2 + 2|a||b|}{4\lambda_2} + v(\alpha^*, \gamma_i^*). \end{aligned} \quad (48)$$

Next, we obtain upper bounds on $|a|$ and $|b|$. By the Cauchy-Schwarz inequality, we have $|a| \leq \epsilon \|X_i\|_2$. Since the columns of X are normalized, the bound simplifies to $|a| \leq \epsilon$. Moreover, note that by Theorem 2, we have $\alpha^* = -r^*$, which leads to the following

$$\frac{1}{2} \|\alpha^*\|_2^2 = \frac{1}{2} \|r^*\|_2^2 \leq F(\beta^*) \leq F(0) = \frac{1}{2} \|y\|_2^2 = \frac{1}{2},$$

implying that $\|\alpha^*\|_2 \leq 1$.

We now present a bound on $|b| = |\alpha^{*T} X_i - \gamma_i^*|$ by considering two cases:

Case 1: ($|\beta_i^*| < M$): here $\gamma_i^* = 0$, which leads to $|b| = |\alpha^{*T} X_i| \leq \|\alpha^*\|_2 \leq 1$.

Case 2: ($|\beta_i^*| = M$): here $\gamma_i^* = \alpha^{*T} X_i - 2M\lambda_2 \text{sign}(\alpha^{*T} X_i)$, which leads to $|b| = 2M\lambda_2$.

Plugging $|a| \leq \epsilon$ and the bounds on $|b|$ (above) into (48), we arrive to (47).

Proof of Theorem 3: First, we consider the case of $\sqrt{\lambda_0/\lambda_2} \leq M$, i.e., we will establish the bound in (30). Note that the following holds:

$$\begin{aligned} \frac{1}{2}\|\hat{\alpha}\|_2^2 + \hat{\alpha}^T y &= \frac{1}{2}\|\hat{\alpha} - \alpha^* + \alpha^*\|_2^2 + (\hat{\alpha} - \alpha^* + \alpha^*)^T y \\ &= \frac{1}{2}\|\hat{\alpha} - \alpha^*\|_2^2 + (\hat{\alpha} - \alpha^*)^T \alpha^* + \frac{1}{2}\|\alpha^*\|_2^2 + (\hat{\alpha} - \alpha^*)^T y - \alpha^{*T} y \\ &\leq \frac{1}{2}\epsilon^2 + \epsilon\|\alpha^*\|_2 + \frac{1}{2}\|\alpha^*\|_2^2 + \epsilon\|y\|_2 + \alpha^{*T} y, \end{aligned} \quad (49)$$

where (49) follows by Cauchy-Schwarz. Plugging $\|y\|_2 = 1$ and $\|\alpha^*\|_2 \leq 1$ (this was established in the proof of Lemma 2) in (49), we get:

$$\frac{1}{2}\|\hat{\alpha}\|_2^2 + \hat{\alpha}^T y \leq \frac{1}{2}\epsilon^2 + 2\epsilon + \frac{1}{2}\|\alpha^*\|_2^2 + \alpha^{*T} y. \quad (50)$$

Plugging (50) into (24):

$$h_1(\hat{\alpha}, \hat{\gamma}) \geq -\frac{1}{2}\|\alpha^*\|_2^2 - \alpha^{*T} y - 2\epsilon - \frac{1}{2}\epsilon^2 - \sum_{i \in [p]} v(\hat{\alpha}, \hat{\gamma}_i). \quad (51)$$

By Lemma 2, for every $i \in \text{Supp}(\hat{\beta})^c$, we have $v(\hat{\alpha}, \hat{\gamma}_i) = 0$, which implies $v(\hat{\alpha}, \hat{\gamma}_i) \leq v(\alpha^*, \gamma_i^*)$ (since v is a non-negative function). Let

$$k_1 = |\{i \in \text{Supp}(\hat{\beta}) \mid |\hat{\beta}_i| < M\}| \quad \text{and} \quad k_2 = |\{i \in \text{Supp}(\hat{\beta}) \mid |\hat{\beta}_i| = M\}|.$$

Plugging $v(\hat{\alpha}, \hat{\gamma}_i) \leq v(\alpha^*, \gamma_i^*)$ for every $i \in \text{Supp}(\hat{\beta})^c$, and (47) for every $i \in \text{Supp}(\hat{\beta})$, in (51), we get (after simplifying):

$$h_1(\hat{\alpha}, \hat{\gamma}) \geq h_1(\alpha^*, \gamma^*) - \epsilon(2 + k_1(2\lambda_2)^{-1} + k_2 M) - \epsilon^2 \left(\frac{1}{2} + \frac{k}{4\lambda_2} \right), \quad (52)$$

which leads to (30).

Now, we consider the case of $\sqrt{\lambda_0/\lambda_2} \geq M$, where we will establish the bound in (31). By the same argument used in deriving (51), we have:

$$h_2(\hat{\rho}, \hat{\mu}) \geq -\frac{1}{2}\|\rho^*\|_2^2 - \rho^{*T} y - 2\epsilon - \frac{1}{2}\epsilon^2 - M\|\hat{\mu}\|_1. \quad (53)$$

Next, we fix some $i \in \text{Supp}(\hat{\beta})$ and we upper bound $\hat{\mu}_i$ as follows:

$$\begin{aligned} \hat{\mu}_i &= \left[|\hat{\rho}^T X_i| - \lambda_0/M - \lambda_2 M \right]_+ \leq \left[|(\hat{\rho} - \rho^*)^T X_i| + |\rho^{*T} X_i| - \lambda_0/M - \lambda_2 M \right]_+ \\ &\leq |(\hat{\rho} - \rho^*)^T X_i| + \left[|\rho^{*T} X_i| - \lambda_0/M - \lambda_2 M \right]_+ \\ &\leq \epsilon + \mu_i^*, \end{aligned} \tag{54}$$

where in the last step above we use Cauchy-Schwarz for the first term, and for the second term we use the following fact (which is a consequence of the optimality conditions):

$$|\rho^{*T} X_i| \leq \lambda_0/M + \lambda_2 M \quad \text{if } |\hat{\beta}_i| < M \quad \text{and} \quad |\rho^{*T} X_i| \geq \lambda_0/M + \lambda_2 M \quad \text{if } |\hat{\beta}_i| = M.$$

For $i \in \text{Supp}(\hat{\beta})^c$, we have $\hat{\mu}_i = 0$, which implies $\hat{\mu}_i \leq \mu_i^*$. Plugging the latter inequality (for all $i \in \text{Supp}(\hat{\beta})^c$) and (54) (for all $i \in \text{Supp}(\hat{\beta})$) into (53), and simplifying, we get:

$$h_2(\hat{\rho}, \hat{\mu}) \geq h_2(\rho^*, \mu^*) - \epsilon(2 + Mk) - \epsilon^2/2, \tag{55}$$

which leads to (31).

B Appendix: Additional Experimental Details

Computing Setup and Software: We used six machines in our experiments, each using 8 cores of an AMD EPYC 7571 processor and 30GB of RAM. We run Gurobi and MOSEK using their Python APIs, and [12]’s OA approach using the SubsetSelectionCIO toolkit [10] in Julia. Relevant software versions: Ubuntu 18.04.3, Python 3.7.6, R 3.6.3, Julia 1.0.5, Gurobi 9.0.1, MOSEK 9.2.3, and L0Learn 1.2. For Julia, exceptionally, we use Gurobi 8.0.1 due to compatibility issues with Gurobi 9.