

Scalable Holistic Linear Regression

Dimitris Bertsimas^a, Michael Lingzhi Li^b

^a*Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139*

^b*Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139*

Abstract

We propose a new scalable algorithm for holistic linear regression building on the research program initiated in Bertsimas and King (2016). Specifically, we develop new theory to model significance and multicollinearity and derive explicit constraints, as opposed to iteratively imposing them in a cutting plane framework, that account for significance and multicollinearity in linear regression. This results in an algorithm that scales with the number of samples n in the 10,000s, while the earlier framework scaled with n in the low 100s when accounting for significance and multicollinearity. We present computational results on real and synthetic datasets that suggest that the proposed approach has a significant computational edge in accuracy, false detection rate, computational time and scalability for holistic linear regression.

Keywords: Holistic Linear Regression, Multicollinearity and Significance in Linear Regression, Mixed-Integer Optimization

1. Introduction

In this paper, we continue the research program initiated in Bertsimas and King (2016) to develop an algorithmic approach for holistic linear regression in which we impose desirable properties simultaneously and a priori. Using mixed integer optimization (MIO) the earlier proposal modeled sparsity, pairwise collinearity and group sparsity using explicit constraints but

accounted for significance and multicollinearity through a cutting plane and bootstrap approach. The difficulty of using the cutting plane method is that it often requires a large number of iterations to ensure that the model has appropriate significance and does not exhibit multicollinearity. This results in an algorithm that does not scale beyond n , the number of samples, in the low 100s when accounting for significance and multicollinearity.

Our goal in this paper is to propose a new scalable algorithm for holistic linear regression. We propose a new way to impose significance and multicollinearity constraints explicitly that scales with n in the 10,000s. This allows us to build linear regression models much more effectively and accurately than in earlier works.

In our view, scalable holistic regression is important as it allows linear regression models to have interpretability, robustness, significance and accuracy in a systematic way. In contrast, today the practice of regression is more of an art than science. Continuing on the vision in Bertsimas and King (2016), the paper aspires to scale holistic regression further and make these methods easier to use in much larger problems.

The standard methodology for imposing significance in linear regression is to use the Student t -statistic. However, the test is carried out after the linear regression model has been calculated, and does not optimally select a subset of covariates that are significant a priori. In the Bertsimas and King (2016) framework, summarized in Section 2, significance is imposed iteratively leading to a cutting plane algorithm. Carrizosa et al. (2017) explored significance of coefficients by adding heuristic constraints to set lower bounds on the coefficients. Chung et al. (2017) used lazy constraints to ensure exact significance tests while deriving theoretical bounds for minimum power. In contrast to Carrizosa et al. (2017) and similar to Chung et al. (2017), we use lazy constraints to ensure minimum power. However, instead of using the t -statistic, we appeal to the asymptotic normality results instead.

For multicollinearity, in a landmark paper Hocking (1976) comprehensively reviewed the problem and concluded that there is no accepted way of dealing with this problem, citing “there is a lack of attention for this problem in the statistics community.” Various methods employed include principal component analysis to select the top k variables to avoid multicollinear combinations, and variance inflation factors Mansfield and Helms (1982) that provide a numerical quantity to determine how much the variance of a co-

efficient has been increased due to correlation with other variables. Tamura et al. (2016, 2017) explored incorporating multicollinearity constraints using variance inflation factors (VIFs) and condition numbers (CNs) respectively. However, both of these concepts are only approximations of true multicollinear relationships. It is true that multicollinear relationships are sufficient for high VIF and CN, but they are not necessary, as shown in O'Brien (2007) and Lazaridis (2007), respectively. That means constraining on VIF or CN would potentially produce extra constraints that are not needed for solving multicollinearity. In this paper, we introduce new theory that provides both necessary and sufficient guarantees in relation to detecting multicollinearity.

Specifically, our contributions in this paper are as follows:

1. We continue the program in Bertsimas and King (2016) and extend the formulation with significance constraints a priori.
2. We develop a new theory of detecting multicollinearity by connecting multicollinearity to the eigenvectors of the design matrix $\mathbf{X}^T \mathbf{X}$, where \mathbf{X} is the $n \times p$ matrix of the given data and use it to impose multicollinearity constraints within an MIO framework
3. We present computational results on real and synthetic datasets that suggest the overall algorithm for holistic regression scales with n in the 10,000s, while the method in Bertsimas and King (2016) scales with n in the low 100s when accounting for significance and multicollinearity.

The structure of the paper is as follows. In Section 2, we review the work in Bertsimas and King (2016) on constructing a holistic framework for linear regression. In Section 3, we introduce the t -statistic formulation to model significance. In Section 4, we introduce a new formulation to model multicollinearity and present computational results with synthetic and real-world data that show the effectiveness of the method. In Section 5, we combine our proposals with the framework introduced in Bertsimas and King (2016), and compare its performance with recent models in the literature using real-world data.

2. The Framework of Bertsimas and King (2016)

Given data $(x_i, y_i), i = 1, \dots, n, x_i \in \mathbb{R}^p, y_i \in \mathbb{R}$, Bertsimas and King (2016) propose the following MIO:

$$\min_{\beta, z} \quad \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|^2 + \Gamma \|\beta\|_1 \quad (1)$$

$$\text{subject to: } -Mz_i \leq \beta_i \leq Mz_i, \quad i = 1, \dots, p \quad (2)$$

$$\sum_{i=1}^p z_i \leq k \quad (3)$$

$$z_i = z_j, \quad \forall i, j \in GS_m, \forall m \quad (4)$$

$$z_i + z_j \leq 1, \quad \forall i, j \in HC \quad (5)$$

$$z_i \in \{0, 1\}, \quad i = 1, \dots, p$$

The term $\Gamma \|\beta\|_1$ in the objective function (1) models robustness as seen in Bertsimas and Copenhaver (2018), who established the equivalence of the ℓ_1 penalization and robustness.

Constraints (2) and (3) model sparsity that at most k out of the p variables are selected in the linear regression model. Constraint (4) models group sparsity, i.e., variables in the set GS_m are either all selected or none is selected. Finally, pairwise collinearity is modeled in Constraint (5) where HC is the set

$$HC = \{(i, j) : |\text{Corr}(x_i, x_j)| \geq \rho\}$$

for some predefined correlation ρ cutoff.

Bertsimas and King (2016) apply the following iterative process to include constraints for significance and multicollinearity:

1. Solve the MIO (1)-(5) to obtain a subset S of the coefficients $\{\beta_1, \dots, \beta_p\}$.
2. For the set S the algorithm computes the significance levels for each of the variables via bootstrap methods, and calculates the condition number of the model. If a set S produces undesirable results – a condition number higher than desired, or a model with insignificant variables – the algorithm generates the constraint

$$\sum_{i \in S} z_i \leq |S| - 1$$

to exclude the set S from consideration. The algorithm adds the constraint to Problem (1) and repeats the process until no such set S is found.

Bertsimas and King (2016) report computational results that demonstrate that Model (1) is effective to solve problems up to n, p in the 1000s. However, when we include significance and multicollinearity constraints in a cutting plane methodology, the method scales up to n, p in the low 100s and some times no solution is found after considerable computation time.

3. Imposing Significance Constraints

Variable significance has long been one of the most important elements in linear regression, and has served as a proxy for variable selection and causality studies.

We first restate a standard result about the asymptotic guarantee of the normality of the least squares estimate of β to serve as the basis of our approach. For a linear regression problem:

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon.$$

We have the following theorem, as proven in Eicker et al. (1963):

Theorem 1. If ϵ is iid with $\mathbb{E}[\epsilon_i] = 0$ and $\mathbb{E}[\epsilon_i^2] = \sigma^2$ for all i , and $\lim_{n \rightarrow \infty} \frac{\mathbf{X}^T \mathbf{X}}{n} = \mathbf{Q}$ is invertible, then we have:

$$\frac{\sqrt{n}(\hat{\beta} - \beta)}{\sigma\sqrt{\mathbf{Q}}} \xrightarrow{d} N(0, 1)$$

Where $\hat{\beta}$ is the least squares estimate of β with

$$\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}.$$

Note that normality is not part of the assumption here - in contrast to the t -test statistics used in Chung et al. (2017). Therefore, using such asymptotic results, we would assume when n large enough, we have that:

$$\frac{\hat{\beta}_j - \beta_j}{\tilde{\sigma}\sqrt{\mathbf{K}_{jj}^{-1}}} \sim N(0, 1). \tag{6}$$

Where $\mathbf{K} = (\mathbf{X}^T \mathbf{X})^{-1}$ and

$$\tilde{\sigma} = \sqrt{\frac{\mathbf{Y}^T (\mathbf{I}_n - \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T) \mathbf{Y}}{n - p}}$$

is the least squares estimate of standard deviation σ .

3.1. Constructing Significance Constraints

For a test of size α , we first define the quantity $N_{sign} = \Phi^{-1}(1 - \frac{\alpha}{2})$, the inverse cdf of the $N(0, 1)$ distribution at point $1 - \frac{\alpha}{2}$. Then, we can impose the normality test by requiring:

$$\frac{|\beta_j|}{\tilde{\sigma} \sqrt{(\mathbf{X}_z^T \mathbf{X}_z)_{jj}^{-1}}} \geq N_{sign} z_j,$$

Where \mathbf{X}_z is the model matrix constrained to the columns where $z_i = 1$. This is equivalent to the big M -constraints:

$$\frac{\beta_j}{\tilde{\sigma} \sqrt{(\mathbf{X}_z^T \mathbf{X}_z)_{jj}^{-1}}} + M b_j \geq N_{sign} z_j \quad (7)$$

$$-\frac{\beta_j}{\tilde{\sigma} \sqrt{(\mathbf{X}_z^T \mathbf{X}_z)_{jj}^{-1}}} + M(1 - b_j) \geq N_{sign} z_j \quad (8)$$

$$b_j \in \{0, 1\}, \quad j = 1, \dots, p,$$

where M is a large constant. These two constraints are used to model significance of level α without the need of the bootstrap. As the model matrix \mathbf{X}_z changes with the selection of \mathbf{z} , in implementation these constraints are implemented as lazy constraints to only be enforced when a feasible integer solution is reached, in a similar fashion to Chung et al. (2017).

In interest of brevity, we defer computational experiments and present results when combined with the multicollinearity detection as illustrated below.

4. Multicollinearity Detection

Given data X , we would like the design matrix to be free of multicollinear relationships so that $\det(\mathbf{X}^T \mathbf{X})$ is not very close to 0. We denote the columns of \mathbf{X} as \mathbf{X}_j , $j = 1, \dots, p$.

We introduce the vector $(1, \dots, 1)^T$ into the design matrix as a new column (the intercept) and we can define the multicollinear relationship as:

Definition 1. A set of variables $\mathbf{X}_1, \dots, \mathbf{X}_p$ has an ϵ -multicollinear relationship if for some $\mathbf{a} \in \mathbb{R}^p$, $\|\mathbf{a}\| = 1$, we have that:

$$\left\| \sum_{j=1}^n a_j \mathbf{X}_j \right\| < \epsilon. \quad (9)$$

The structure of this section is as follows:

1. We first establish the key result that connects the existence of an ϵ -multicollinear relationship (9) to the existence of an eigenvector \mathbf{v} for the matrix $\mathbf{X}^T \mathbf{X}$ that has a small ($O(\sqrt{\epsilon})$) eigenvalue.
2. Using the previous key result, we find multicollinear relations ($\mathbf{a} = (a_1, \dots, a_p)$) using information from the small eigenvalues of the matrix $\mathbf{X}^T \mathbf{X}$. We introduce the idea of a *minimum-support multicollinear relationship*.
3. We propose an algorithm that uses the theory from the previous steps to identify all the multicollinear relationships.

4.1. Key Result

In this section, we establish a connection between the existence of an ϵ -multicollinear relationship and the existence of a eigenvector \mathbf{v} for the matrix $\mathbf{X}^T \mathbf{X}$ with a small ($O(\sqrt{\epsilon})$) eigenvalue:

Theorem 2. Let $V = \{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ be the set of orthonormal eigenvectors of $\mathbf{X}^T \mathbf{X} \in \mathbb{R}^{p \times p}$ such that the eigenvalues associated with V are less than ϵ . Then for $\mathbf{a} \in \mathbb{R}^p$, $\|\mathbf{a}\| = 1$:

- (a) If $\left\| \sum_{j=1}^p a_j \mathbf{X}_j \right\| < \epsilon$, then there exists a vector $\mathbf{b} \in \mathbb{R}^p$, $\|\mathbf{b}\| < (p - m)\sqrt{\epsilon}$ such that $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$.
- (b) If there exists a vector $\mathbf{b} \in \mathbb{R}^p$, $\|\mathbf{b}\| < \sqrt{\epsilon}$ such that $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$, then we have:

$$\left\| \sum_{j=1}^p a_j \mathbf{X}_j \right\| < \sqrt{(1 + \lambda_{m+1} + \dots + \lambda_p)\epsilon},$$

where $\lambda_{m+1}, \dots, \lambda_p$ are the eigenvalues associated with the set of orthonormal eigenvectors of M that have value greater or equal to ϵ .

Theorem 2 represents a weak equivalence between a small multicollinear relationship and the existence of a vector \mathbf{a} that is close to $\text{Span}(V)$, in the sense that there exists a small vector \mathbf{b} with $\|\mathbf{b}\| < O(\sqrt{\epsilon})$ such that $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$. The proof is as follows:

Proof.

- (a) If $m = p$, then every $\mathbf{a} \in \text{Span}(V)$. Thus, we assume $m < p$ and prove part (a) by contradiction. We assume there exists no $\mathbf{b} \in \mathbb{R}^p$ with $\|\mathbf{b}\| < (p - m)\sqrt{\epsilon}$ such that $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$. Let $\lambda_1, \dots, \lambda_p$ be the corresponding eigenvalues to eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_p$. Note that we have $0 \leq \lambda_1, \dots, \lambda_m < \epsilon$, and $\epsilon \leq \lambda_{m+1}, \dots, \lambda_p$. We write \mathbf{a} as:

$$\mathbf{a} = \alpha_1 \mathbf{v}_1 + \dots + \alpha_p \mathbf{v}_p.$$

Letting $\mathbf{b} = \alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p$, we have that $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$ by construction, which implies that:

$$\|\mathbf{b}\| = \|\alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p\| \geq (p - m)\sqrt{\epsilon}.$$

This implies that there exists a α_{j_0} , $j_0 \in \{m + 1, \dots, p\}$ such that $\|\alpha_{j_0}\| \geq \sqrt{\epsilon}$. Now,

$$\left\| \sum_{j=1}^p a_j \mathbf{X}_j \right\| = \|\mathbf{X}\mathbf{a}\| = \|\alpha_1 \mathbf{X}\mathbf{v}_1 + \dots + \alpha_p \mathbf{X}\mathbf{v}_p\| < \epsilon.$$

We have

$$\begin{aligned} \epsilon^2 &> \mathbf{a}^T \mathbf{X}^T \mathbf{X} \mathbf{a} \\ &= (\alpha_1 \mathbf{X}\mathbf{v}_1 + \dots + \alpha_p \mathbf{X}\mathbf{v}_p)^T (\alpha_1 \mathbf{X}\mathbf{v}_1 + \dots + \alpha_p \mathbf{X}\mathbf{v}_p) \\ &= \alpha_1^2 \lambda_1 + \dots + \alpha_p^2 \lambda_p \\ &\geq \alpha_{j_0}^2 \lambda_{j_0}. \end{aligned}$$

Since $|\alpha_{j_0}| \geq \sqrt{\epsilon}$, and $\lambda_{j_0} \geq \epsilon$, we have that $\epsilon^2 > \alpha_{j_0}^2 \lambda_{j_0} \geq \epsilon^2$, a contradiction.

- (b) If $m = p$, then $\mathbf{a} = \sum_{j=1}^p a_j \mathbf{v}_j$. Note $\|\mathbf{a}\|^2 = \sum_{j=1}^p a_j^2$, since \mathbf{v}_j are

orthonormal. Hence, for $\|\mathbf{a}\| = 1$

$$\begin{aligned}
\|\mathbf{X}\mathbf{a}\|^2 &= \mathbf{a}^T \mathbf{X}^T \mathbf{X} \mathbf{a} \\
&= \left(\sum_{j=1}^p a_j \mathbf{v}_j \right)^T \mathbf{X}^T \mathbf{X} \left(\sum_{j=1}^p a_j \mathbf{v}_j \right) \\
&= \left(\sum_{j=1}^p a_j \mathbf{v}_j \right)^T \left(\sum_{j=1}^p a_j \lambda_j \mathbf{v}_j \right) \\
&= \sum_{j=1}^p \lambda_j a_j^2 \\
&< \epsilon \|\mathbf{a}\|^2 = \epsilon.
\end{aligned}$$

leading to $\|\mathbf{X}\mathbf{a}\| < \sqrt{\epsilon}$. We assume $m < p$. We write \mathbf{a} as:

$$\mathbf{a} = \alpha_1 \mathbf{v}_1 + \dots + \alpha_p \mathbf{v}_p$$

and observe that:

$$\min_{\mathbf{u} \in \text{Span}(V)} \|\mathbf{a} - \mathbf{u}\| = \|\alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p\|. \quad (10)$$

Since by assumption there exists a \mathbf{b} with $\|\mathbf{b}\| < \sqrt{\epsilon}$ and $\mathbf{a} - \mathbf{b} \in \text{Span}(V)$, the vector $\mathbf{a} - \mathbf{b}$ is a feasible solution to problem (10), and thus taking $\mathbf{u} = \mathbf{a} - \mathbf{b}$ we have:

$$\|\mathbf{a} - (\mathbf{a} - \mathbf{b})\| = \|\mathbf{b}\| \geq \|\alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p\|,$$

leading to:

$$\|\alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p\| < \sqrt{\epsilon}.$$

Since

$$\|\alpha_{m+1} \mathbf{v}_{m+1} + \dots + \alpha_p \mathbf{v}_p\|^2 = \sum_{j=m+1}^p \alpha_j^2 < \epsilon,$$

Therefore, we have $|\alpha_j| < \sqrt{\epsilon}$ for all $j = m+1, \dots, p$. Thus, we have:

$$\begin{aligned}
\mathbf{a}^T \mathbf{X}^T \mathbf{X} \mathbf{a} &= (\alpha_1 \mathbf{X} \mathbf{v}_1 + \dots + \alpha_p \mathbf{X} \mathbf{v}_p)^T (\alpha_1 \mathbf{X} \mathbf{v}_1 + \dots + \alpha_p \mathbf{X} \mathbf{v}_p) \\
&= \alpha_1^2 \lambda_1 + \dots + \alpha_p^2 \lambda_p \\
&= (\alpha_1^2 \lambda_1 + \dots + \alpha_m^2 \lambda_m) + (\alpha_{m+1}^2 \lambda_{m+1} + \dots + \alpha_p^2 \lambda_p) \\
&\leq \epsilon \cdot \sum_{j=1}^m \alpha_j^2 + \epsilon \cdot \sum_{j=m+1}^p \lambda_j \\
&\leq (1 + \lambda_{m+1} + \dots + \lambda_p) \epsilon,
\end{aligned}$$

leading to $\|\mathbf{X}\mathbf{a}\| \leq \sqrt{(1 + \lambda_{m+1} + \dots + \lambda_p)\epsilon}$ as required.

□

Theorem 2 implies that if we are able to describe $\text{Span}(V)$, then we would be able to identify multicollinear relationships \mathbf{a} that exist in the design matrix \mathbf{X} , as Theorem 2(b) implies that every vector within $\sqrt{\epsilon}$ distance away from $\text{Span}(V)$ represents a $O(\sqrt{\epsilon})$ multicollinear relationship.

4.2. Identifying Multicollinear Relations

For $\dim(V) = r$, we have $r - 1$ linearly independent multicollinear relationships. There are infinite number of ways the basis of the $r - 1$ multicollinear relationships could be constructed, and different ways of constructing such bases lead to different constraints.

For example, assume that we have six variables $x_1, x_2, x_3, x_4, x_5, x_6$, and we know that $x_1 + x_2 = x_3$ and $x_4 + x_5 = x_6$. Letting $\mathbf{a}_1 = (1, 1, -1, 0, 0, 0)^T$ and $\mathbf{a} = (0, 0, 0, 1, 1, -1)^T$, we have $V = \text{Span}(\mathbf{a}_1, \mathbf{a})$. Using Theorem 2 and ignoring \mathbf{b} as $\|\mathbf{b}\| = O(\sqrt{\epsilon})$, we can identify the two multicollinear relationships as \mathbf{a}_1 and \mathbf{a} . Then, we add the constraints

$$z_1 + z_2 + z_3 \leq 2, \quad z_4 + z_5 + z_6 \leq 2$$

to Model (1). However, there are alternative ways to characterize V in terms of two linearly independent vectors. Letting $\bar{\mathbf{a}}_1 = (1, 1, -1, 1, 1, -1)^T$ and $\bar{\mathbf{a}} = (1, 1, -1, -1, -1, 1)^T$, then V is also $V = \text{Span}(\bar{\mathbf{a}}_1, \bar{\mathbf{a}})$. Given this representation of V we would impose the constraints

$$z_1 + z_2 + z_3 + z_4 + z_5 + z_6 \leq 4$$

to Model (1). Note that the two sets of constraints are not equivalent.

It is therefore important to identify the characterization of V that leads to the most stringent constraints to prevent multicollinearity. Towards this objective and ignoring the vector \mathbf{b} in Theorem 2, we introduce the idea of identifying a vector $\mathbf{a} \in \text{Span}(V)$ that has minimum support. We first compute the set $V = \{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ of orthonormal eigenvectors with corresponding eigenvalues less than ϵ . According to Theorem 2, all multicollinear relationships (up to a perturbation of ϵ) are included in this space. Now, we

want to find a vector $\mathbf{a} \in \text{Span}(V)$ of minimum support. This is computed as follows:

$$\min \quad \sum_{j=1}^m z_j \quad (11)$$

$$\text{subject to} \quad \mathbf{a} = \sum_{i=1}^m \theta_i \mathbf{v}_i$$

$$|a_j| \leq M \cdot z_j, \quad j = 1, \dots, m \quad (12)$$

$$\sum_{i=1}^m |\theta_i| \geq \delta$$

$$z_j \in \{0, 1\}, \quad j = 1, \dots, m,$$

Note that (12) can be modeled as Special Ordered Sets (SOS) of type 1, which does not need an explicit value of M . In the experiments, however, we utilize the big M formulation. We provide the following procedure for determining M . We reformulate (12) to read as:

$$\left\| \sum_{i=1}^m \theta_i \mathbf{v}_i \right\|_1 \leq M \sum_{i=1}^m z_j \leq Mm. \quad (13)$$

Using the fact that \mathbf{v}_i 's are orthonormal, we have:

$$\left\| \sum_{i=1}^m \theta_i \mathbf{v}_i \right\|_1 \leq \sqrt{m} \|\boldsymbol{\theta}\|_2. \quad (14)$$

Taking $\|\boldsymbol{\theta}\|_2 = 1$, we can select $M = \frac{\sqrt{m}}{m} = \frac{1}{\sqrt{m}}$. We note this is the tightest possible M with equality at $\mathbf{v}_i = \mathbf{e}_i$ and $\boldsymbol{\theta} = (\frac{1}{\sqrt{m}}, \dots, \frac{1}{\sqrt{m}})$.

Here δ is a positive constant that ensures that $\mathbf{a} \neq 0$. Once the vector \mathbf{a} has been identified, we add the constraint

$$\sum_{i \in \text{Supp}(\mathbf{a})} z_i \leq |\text{Supp}(\mathbf{a})| - 1 \quad (15)$$

to Problem (1). To continue the process of identifying new linearly independent multicollinear relationships, we add Eq. (15) to Problem (11), resolve

the problem to identify a new multicollinear relationship, add the corresponding constraint (15) to (1). We continue solving Problem (11) until the problem becomes infeasible, which means that we identified all linearly independent multicollinear relationships. Algorithm 1 determines all multicollinear relationships.

Algorithm 1 Iterative MIO for finding all linearly independent multicollinear relationships.

```

1: procedure ITERATIVEMC( $E$ )  $\triangleright E$ , the set of  $\mathbf{X}^T \mathbf{X}$  eigenvectors with
   eigenvalues  $< \epsilon$ 
2:    $S \leftarrow \emptyset$   $\triangleright$  Initialize Output
3:    $i \leftarrow 0$   $\triangleright$  Initialize count of found multicollinear relationships
4:   while  $i < |E|$  do  $\triangleright$  We cannot find more if  $i = |E| - 1$ 
5:      $\mathbf{a}_0 =$  Solution of (11)
6:     if  $\mathbf{v}_0 \neq \emptyset$  then
7:        $S \leftarrow S \cup \mathbf{v}_0$ 
8:       Add the constraint  $\sum_{i \in \text{Supp}(\mathbf{a}_0)} z_i \leq |\text{Supp}(\mathbf{a}_0)| - 1$  to (11)
9:        $i \leftarrow i + 1$ 
10:    else  $\triangleright$  If (11) is Infeasible
11:      Exit Loop
12:    end if
13:  end while
14:  return  $S$   $\triangleright$  Return the characterization of  $V$ 
15: end procedure

```

4.3. Computational Results

In this section, we use synthetic data to evaluate the performance of Algorithm 1.

We model the design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ such that $X_{ij} \sim N(0, 1)$ independently for each $i \in \{1, \dots, n\}$, $j \in \{1, \dots, p\}$. Then we randomly select certain number of columns to be replaced by linear combinations of other columns $\sum_{i_j \in S} \gamma_j \mathbf{X}_j$. The parameters γ_i are selected randomly from the uniform distribution $[-10, 10]$, and we control S as follows:

1. We first determine the number q of variables we want to involve in this multicollinear relationship.
2. We randomly select q numbers from $\{1, \dots, p\}$ without replacement, and denote that set S .

We add noise $\tilde{\mathbf{X}}$ according to the distribution indicated in Table 1, and evaluate the performance of Algorithm 1 on $\mathbf{X} + \tilde{\mathbf{X}}$.

Algorithm 1 performance is evaluated on the accuracy and the false positive rate of the multicollinear relationships found, along with the time taken for the algorithm to converge.

In Table 1, $MR(q)$ indicates the number of multicollinear relationships involving q variables that have been introduced into the data. For $MR(4+)$, we randomly selected a number within $\{5, 6, 7, 8, 9, 10\}$ to be the number of variables involved in the multicollinear relationship. We created 10 random instances and report the average statistics across those 10 instances.

n	p	$MR(3)$	$MR(4)$	$MR(4+)$	Noise	ACC	FPR	Time
1000	100	3	1	1	$N(0, 0.01)$	100%	0%	0.27s
1000	500	3	1	1	$N(0, 0.01)$	100%	0%	2.37s
1000	1000	3	1	1	$N(0, 0.01)$	100%	5%	20.23s
1000	500	5	3	2	$N(0, 0.01)$	100%	0%	33.40s
1000	1000	5	3	2	$N(0, 0.01)$	100%	24%	5940.56s
1000	500	3	1	1	$N(0, 0.03)$	100%	0%	2.29s
1000	1000	3	1	1	$N(0, 0.03)$	100%	11%	32.17s

Table 1: Performance of Algorithm 1 for multicollinearity detection.

Table 1 shows that Algorithm 1 scales up to n, p in thousands and could detect multicollinearity with high accuracy and low false positive rates.

5. Holistic Linear Regression Framework Evaluation

In this section, we combine the results of the previous two sections with the framework introduced in Bertsimas and King (2016) on five different datasets randomly selected from the UCI Machine Learning Repository (Dua and Graff (2017)). The whole framework is:

$$\begin{aligned} & \min \quad (1) \\ & \text{subject to } (2) - (5), (7), (8), (15) \end{aligned}$$

We select $\delta = 10^{-6}$ for the multicollinearity detection. We compare our formulation with a MISDO formulation by Tamura et al. (2017) and a MIQO formulation by Chung et al. (2017). We also used Lasso (Tibshirani (1996)), and the framework in Bertsimas and King (2016) as baselines. Samples were randomized and we utilized a 60%/20%/20% for training, validation, and testing, where the validation set was used for tuning of hyperparameters, such as the Γ in the framework (1). We refer to our framework below in Table 2 as *Holistic*. The results are compared across the following dimensions:

- **Sparsity** (k) - Number of non-zero variables in the model.
- **Regression Loss** (Loss) - Mean squared error on the test set.
- **Significance** - Percentage of non-zero coefficients in the model that are significant on the 5% level using bootstrap to evaluate.
- **Time** (T) - Total time used by the model. Time spent detecting multicollinearity for the Holistic model is shown in brackets.
- **Multicollinearity Accuracy** (MA) - Let V be the set as defined in Theorem 2 and $V_{\mathbf{z}}$ be the corresponding set in the final model using $\epsilon = 10^{-2}$. Then we calculate $100\% \times \left(1 - \frac{\dim(V_{\mathbf{z}})}{\dim(V)}\right)$, the percentage of multicollinearity relations "avoided" in the final model.

Dataset	n	p	Method	k	Loss	Sign.	T	MA
Airfoil	1502	5	<i>Holistic</i>	3	558	100%	50s	-
			<i>Tamura</i>	4	562	75%	19s	-
			<i>Chung</i>	3	558	64%	107s	-
			<i>Lasso</i>	4	570	75%	7s	-
			<i>Bootstrap</i>	4	564	75%	39870s	-
Cancer	568	29	<i>Holistic</i>	7	1.71	100%	54s	-
			<i>Tamura</i>	11	1.90	63%	410s	-
			<i>Chung</i>	7	1.71	100%	310s	-
			<i>Lasso</i>	23	0.72	31%	10s	-
			<i>Bootstrap</i>	12	2.22	60%	60000s	-
Parkinsons	5875	16	<i>Holistic</i>	1	533	100%	403s (15.2s)	100%
			<i>Tamura*</i>	1	533	100%	60000s	100%
			<i>Chung</i>	3	549	100%	876s	50%
			<i>Lasso</i>	3	522	33%	14s	50%
			<i>Bootstrap</i>	3	571	33%	60000s	50%
Air Quality	9358	12	<i>Holistic</i>	4	89.2	100%	380s	-
			<i>Tamura</i>	N/A	N/A	N/A	N/A	-
			<i>Chung</i>	6	96.1	100%	770s	-
			<i>Lasso</i>	9	83.7	33%	11s	-
			<i>Bootstrap</i>	5	89.6	80%	58146s	-
Crime	2215	125	<i>Holistic</i>	9	180	100%	725s (120.3s)	100%
			<i>Tamura</i>	N/A	N/A	N/A	N/A	N/A
			<i>Chung</i>	11	207	100%	1506s	40%
			<i>Lasso</i>	19	172	47%	21s	20%
			<i>Bootstrap</i>	12	195	74%	60000s	20%

Table 2: Comparison of the holistic framework with Lasso for five real world data sets. - means that there are no true multicollinear relationships in the data. N/A means the algorithm did not return a feasible solution within 60000s.

The results in Table 2 show that in real data situations, the entire framework could reasonably scale up to 10,000 in n and at least 100 in p , while both the MISDO and the MIQO framework scaled slower. In particular, the MISDO formulation failed to return a feasible solution for the largest datasets, and encountered numerical issues in the process (which was also identified in the original paper). The MISDO formulation also does not consider significance constraints, which was reflected in that some of the variables it selected were insignificant. The MIQO formulation scaled better compared

to the MISDO formulation, but was still over $2x$ slower than holistic regression in all datasets. The MIQO formulation does not explicitly model multicollinearity, and this meant the final model in the Parkinsons and the Crime dataset did not avoid more than 50% of multicollinear relationships. In comparison, the holistic regression successfully detected 100% all multicollinear relationships within the data and avoided choosing all variables within that relationship in the final result.

Compared to the Lasso baseline, the holistic framework achieved comparable loss with Lasso among most tasks, while using many fewer variables to do so (usually less than half), and all of the selected variables from the framework are significant at the 5% level. The original bootstrap method proposed in Bertsimas and King (2016) quickly timed out as p increased, resulting in suboptimal performance.

The computational results suggest that the proposed holistic linear regression algorithm greatly increases its scalability in detecting significance and avoiding multicollinearity. Using both real and synthetic data, we have demonstrated that the approach produces high quality linear regression models in realistic timelines.

- Bertsimas D, Copenhaver MS (2018) Characterization of the equivalence of robustification and regularization in linear and matrix regression. *European Journal of Operational Research* 270:931–942.
- Bertsimas D, Copenhaver MS, Mazumder R (2017) The trimmed lasso: Sparsity and robustness. *arXiv preprint arXiv:1708.04527* .
- Bertsimas D, King A (2016) An algorithmic approach to linear regression. *Operations Research* 64(1):2–16.
- Carrizosa E, Olivares-Nadal AV, Ramirez-Cobob P (2017) Enhancing interpretability by tightening linear regression models. Technical report, Technical report.
- Chung S, Park YW, Cheong T (2017) A mathematical programming approach for integrated multiple linear regression subset selection and validation. *arXiv preprint arXiv:1712.04543* .
- Dua D, Graff C (2017) UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>.
- Eicker F (1966) A multivariate central limit theorem for random linear vector forms. *Ann. Math. Statist.* 37(6):1825–1828.
- Eicker F, et al. (1963) Asymptotic normality and consistency of the least squares estimators for families of linear regressions. *The Annals of Mathematical Statistics* 34(2):447–456.

- Hocking RR (1976) A biometrics invited paper. the analysis and selection of variables in linear regression. *Biometrics* 32(1):1–49.
- Lazaridis A (2007) A note regarding the condition number: the case of spurious and latent multicollinearity. *Quality & Quantity* 41(1):123–135.
- Mann HB, Wald A (1943) On stochastic limit and order relationships. *Ann. Math. Statist.* 14(3):217–226.
- Mansfield ER, Helms BP (1982) Detecting multicollinearity. *The American Statistician* 36(3a):158–160.
- O'Brien RM (2007) A caution regarding rules of thumb for variance inflation factors. *Quality & quantity* 41(5):673–690.
- Stewart GW (1990) Matrix perturbation theory.
- Tamura R, Kobayashi K, Takano Y, Miyashiro R, Nakata K, Matsui T (2016) Mixed integer quadratic optimization formulations for eliminating multicollinearity based on variance inflation factor. optimization online. *Optimization Online* .
- Tamura R, Kobayashi K, Takano Y, Miyashiro R, Nakata K, Matsui T (2017) Best subset selection for eliminating multicollinearity. *Journal of the Operations Research Society of Japan* 60(3):321–336.
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* 58(1):267–288.