

# Stochastic Cutting Planes for Data-Driven Optimization

Dimitris Bertsimas · Michael Lingzhi Li

Received: 12/30/2020 / Accepted: date

**Abstract** We introduce a stochastic version of the cutting-plane method for a large class of data-driven Mixed-Integer Nonlinear Optimization (MINLO) problems. We show that under very weak assumptions the stochastic algorithm is able to converge to an  $\epsilon$ -optimal solution with high probability. Numerical experiments on several problems show that stochastic cutting planes is able to deliver a multiple order-of-magnitude speedup compared to the standard cutting-plane method. We further experimentally explore the lower limits of sampling for stochastic cutting planes in dimension  $n$  and show that for many problems, a sampling size of  $O(\sqrt[3]{n})$  appears to be sufficient for high quality solutions.

**Keywords** Stochastic Optimization · Mixed-Integer Optimization · Machine Learning · Cutting Planes · Outer Approximation · Scaling

---

D. Bertsimas  
Sloan School of Management and Operations Research Center  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
E-mail: dbertsim@mit.edu

M. L. Li  
Operations Research Center  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
E-mail: mlli@mit.edu

## 1 Introduction

The cutting plane method, or the outer approximation method, has been a popular method to solve Mixed Integer Non-Linear Optimization (MINLO) problems since its original introduction in Duran and Grossmann (1986) and Fletcher and Leyffer (1994). It tackles problems of the form:

$$\begin{aligned} & \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta} f(\mathbf{z}, \boldsymbol{\theta}) \\ & \text{subject to } g_i(\mathbf{z}, \boldsymbol{\theta}) \leq 0, \quad i \in [m], \end{aligned}$$

where  $[m] = \{1, \dots, m\}$ ,  $\mathbf{z}$  and  $\boldsymbol{\theta}$  are integer and continuous variables constrained to be in sets  $\mathcal{Z}$  and  $\Theta$ , respectively, and  $f, g_i$  are convex. The method aims to generate the tangent plane of both the objective and the constraints at an initial feasible solution  $(\mathbf{z}_0, \boldsymbol{\theta}_0)$ , and solve the resultant mixed integer linear optimization (MILO) instead in order to find an optimal solution  $(\mathbf{z}_1, \boldsymbol{\theta}_1)$ . Then, the process is repeated where the tangent plane around  $(\mathbf{z}_1, \boldsymbol{\theta}_1)$  is added until the solution stops updating (with a tolerance of  $\epsilon$ ). Fletcher and Leyffer (1994) proved that this method does indeed converge to the optimal solution in finite steps, and illustrated its practical effectiveness.

With the rise in availability of data in recent years, the cutting plane method has been applied to many data-driven optimization problems, ranging from risk minimization (Franc and Sonnenburg 2009), sparse regression Bertsimas and Van Parys (2020), matrix completion (Bertsimas and Li 2018), to multiple kernel learning and inference in graphical models (Franc et al. 2011).

However, as the data sizes continue to grow, cutting plane approaches face challenges in scaling as generating the full tangent planes usually involves evaluating the objective function and its derivative across all samples. With sample sizes easily reaching scales of  $10^6$  or higher, this presents a computational difficulty for utilizing cutting-plane methods in large-scale real world applications.

In this paper, we introduce a stochastic variant of the cutting plane algorithm that aims to greatly increase the scalability of the cutting plane problem by multiple orders of magnitude. Specifically, instead of aiming to generate the exact tangent plane in each iteration, we utilize an independent random subset of the data to generate an approximation of the tangent plane. A parallel analogy of this idea in continuous optimization is stochastic gradient descent, where only a subset of the samples is considered to calculate a gradient, resulting in large performance increases that enables the optimization of millions of variables (such as in a convolutional neural network).

Despite the simplicity of the idea, the authors could not find any detailed work exploring the theoretical properties and computational performance of this approach. In this work, we characterize the general convex data-driven

optimization problems that stochastic cutting planes can be applied to, and demonstrate that the proposed method is able to converge to a near optimal solution with an exponentially vanishing failure probability. We illustrate that on multiple distinct problems (sparse regression, support vector machines and stochastic knapsack problems), stochastic cutting planes record an order-of-magnitude speedup compared to standard cutting planes and compare competitively with state-of-the art specialized algorithms.

## 1.1 Structure

In Section 2, we introduce the general class of convex data-driven optimization problems to which we apply our framework. In Section 3, we prove that the stochastic cutting plane algorithm converges to an optimal solution of the original problem with exponentially vanishing failure probability under very weak conditions. In Section 4, we present computational results on a variety of problems to show that the algorithm considerably outperforms the standard cutting plane method, and extends scalability by multiple orders of magnitude.

## 1.2 Literature

Duran and Grossmann (1986) first introduced the exact cutting-plane method in the context of mixed integer non-linear optimization (MINLO) problems that are linear in the integer variables. Fletcher and Leyffer (1994) extended the framework to the general class of MINLO problems and provided theoretical guarantees for the case that the MINLO is convex. Parikh (1976) investigated an approximate cutting-plane method that involved generating a sequence of hyperplanes using non-optimal dual solutions to the inner NLO problem. The approximate cutting planes generated are always valid for the original objective function, but the approximation is only effective “near an optimal solution”.

Many other works entitled “stochastic cutting planes” or similar are primarily focused on randomized constraint generation in the presence of a large (possibly infinite) number of constraints, such as Volkov and Zavriev (1997) or Calafiore and Campi (2005) in robust optimization. Concretely, they consider problems of the form:

$$\begin{aligned} & \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta} f(\mathbf{z}, \boldsymbol{\theta}) \\ & \text{subject to } g_i(\mathbf{z}, \boldsymbol{\theta}) \leq 0, \quad i \in [m], \end{aligned}$$

where  $m$  is very large or an infinite set. The proposed algorithms generally first solves the problem with only a random finite subset of constraints, and

then add the remaining constraints iteratively (and randomly) till convergence. In particular, they assume that the function  $f(\mathbf{z}, \boldsymbol{\theta})$  is relatively simple to evaluate. This work focuses on the case where the objective function takes significant time to calculate and proves theoretical guarantees on optimal solution retrieval using stochastic cutting planes. Furthermore, we demonstrate real-world scalability with data beyond  $10^6$  samples.

In stochastic optimization, sample average approximation is frequently utilized to provide an approximation to the objective function (which is usually of the form of an expected value). For stochastic MINLO problems, there has been work (see e.g. Wei and Realff (2004), Kleywegt et al. (2002)) adapting the cutting planes methodology for the stochastic problem by solving the sample average approximation at every iteration using a new sample of size  $N$  multiple times. This work complements previous work by exploring how to solve the sample average approximation with sample size  $N$  efficiently, when  $N$  is large. In particular in Section 4, we demonstrate how stochastic cutting planes generates an 100x speedup to solve the sample average approximation version of the static stochastic knapsack problem (SSKP).

Bertsimas and Li (2018) utilized a similar idea to speed up the cutting-plane method utilized to solve a variant of the matrix completion problem. The theoretical guarantees in this paper were limited to the specific application and in addition required strong convexity, which is not required in the proof for the present paper.

## 2 Problem Formulation and Examples

We consider a general class of mixed-integer data-driven optimization problems with integer variables  $\mathbf{z} \in \mathbb{Z}^{p_1}$  and continuous variables  $\boldsymbol{\theta} \in \mathbb{R}^{p_2}$ . In addition, there is data  $\mathbf{d}_1, \dots, \mathbf{d}_N$  obtained either through historical results or simulations. As common in these settings, the goal is to find a representation of the data under some model driven by the variables  $\mathbf{z}, \boldsymbol{\theta}$ .

Utilizing the shorthand  $[N]$  to represent  $\{1, \dots, N\}$ , we can write the objective function of such data driven problem as:

$$f(\mathbf{z}, \boldsymbol{\theta}; [N]) := f(\mathbf{z}, \boldsymbol{\theta}; \mathbf{d}_1, \dots, \mathbf{d}_N),$$

where we use the notation  $f(\mathbf{z}, \boldsymbol{\theta}; [N])$  to denote that the objective function is created with samples in the set  $[N]$ . One of the most common examples of this type of objective function is data-driven risk minimization, where we have  $\mathbf{d}_i = (\mathbf{x}_i, \mathbf{y}_i)$ , and the objective function is:

$$f(\mathbf{z}, \boldsymbol{\theta}; [N]) = \frac{1}{N} \sum_{i=1}^N l(\mathbf{y}_i, m(\mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta})).$$

Here  $l(\cdot, \cdot)$  is some pre-defined loss function describing the difference between the target  $\mathbf{y}_i$  and the model fit  $m(\mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta})$ .

We consider the data-driven optimization problem:

$$\begin{aligned} \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta} f(\mathbf{z}, \boldsymbol{\theta}; [N]) \\ \text{subject to } g_i(\mathbf{z}, \boldsymbol{\theta}) \leq 0, \quad i \in [m], \end{aligned} \quad (1)$$

with the following assumptions:

**Assumption 1 (Convexity)**  $f(\mathbf{z}, \boldsymbol{\theta}; [N])$  and  $g_i(\mathbf{z}, \boldsymbol{\theta})$  are at least once differentiable and convex in  $\mathbf{z}, \boldsymbol{\theta}$  for all  $i \in [m]$ .

**Assumption 2 (Compactness)**  $\mathcal{Z}$  is a finite set in  $\mathbb{Z}^{p_1}$  and  $\Theta$  is a compact set in  $\mathbb{R}^{p_2}$ .

**Assumption 3 (Data Concentration)** Let  $S^n$  be a random subset of  $[N]$  of size  $n$  selected without replacement. Then with probability at least  $1 - \epsilon$ , we have:

$$\begin{aligned} |f(\mathbf{z}, \boldsymbol{\theta}; S^n) - f(\mathbf{z}, \boldsymbol{\theta}; [N])| &\leq M \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}}, \\ \|\nabla f(\mathbf{z}, \boldsymbol{\theta}; S^n) - \nabla f(\mathbf{z}, \boldsymbol{\theta}; [N])\| &\leq M' \sqrt{\frac{(p_1 + p_2) \log(\frac{1}{\epsilon})}{n}}, \end{aligned}$$

where  $M$  and  $M'$  are absolute constants independent of  $p_1, p_2, n, N$ .

The convexity and compactness assumptions are standard in the classic cutting-plane literature. In real-world problems, the convexity assumption might prove restrictive as many important problems do not admit convex formulations in both the integer and continuous variables. Nevertheless, the proposed algorithm can also be applied to problems that are non-convex in  $f$  and  $g$ , but the solution would not carry any theoretical guarantees.

The data concentration assumption essentially states that the objective function (and its derivative) created with a random subset of the data is concentrated around the value of the objective (and its derivative) under the full set of the data. In essence, this requires that a small instance of the problem is “similar” to the original problem. Note that this assumption in particular does *not* require  $f(\mathbf{z}; S^n)$  to be an unbiased estimator of  $f(\mathbf{z}; N)$ , as long the estimate  $f(\mathbf{z}; S^n)$  is sufficiently close. This is true for the vast majority of data-driven problems, and we now introduce a few prominent classes of these problems.

## 2.1 Data-Driven Risk Minimization

Given features  $\mathbf{x}_i$  and responses  $y_i$ ,  $i \in [N]$ , the data-driven risk minimization problem aims to solve the following unconstrained optimization problem:

$$\arg \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta} \frac{1}{N} \sum_{i=1}^N l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta}), \quad (2)$$

where  $l$  is a function convex in  $\mathbf{z}, \boldsymbol{\theta}$ . This structure contains a vast array of problems, including:

- *Machine Learning*: This includes for example support vector machines:

$$l(\cdot, \cdot, \cdot, \cdot) = \max\{1 - y_i \boldsymbol{\theta}^T \mathbf{x}_i, 0\},$$

least-squares regression:

$$l(\cdot, \cdot, \cdot, \cdot) = \|y_i - \boldsymbol{\theta}^T \mathbf{x}_i\|^2,$$

and many others.

- *Treatment Effect Estimation*: This includes a wide range of data-driven optimization problems in pricing and policy treatment effect estimation (see e.g. Rivers and Vuong 1988, Blundell and Powell 2004, Alley et al. 2019) where the loss is constructed from a parametric or semi-parametric loss function and there are further constraints on the parameters (non-negativity, bounds, etc). Many of these problems might involve non-convex loss functions, but they are usually solved by methods that assume convexity (e.g. gradient descent, cutting planes) or create successive convex approximations. For example, given a probit model  $\hat{f}(\mathbf{x}_i, \boldsymbol{\theta})$  for estimating propensity of purchasing a model, the form of the loss function would appear as:

$$l(\cdot, \cdot, \cdot, \cdot) = \frac{1}{N} \sum_{i=1}^N y_i \log(\hat{f}(\mathbf{x}_i, \boldsymbol{\theta})) + (1 - y_i) \log(1 - \hat{f}(\mathbf{x}_i, \boldsymbol{\theta})),$$

where  $N$  is the number of samples, and  $(\mathbf{x}_i, y_i)$  are the historical data.

- *Sample Average Approximation* : Sample Average Approximation (SAA) methods are commonly used in stochastic optimization problems where the master problem:

$$\min_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_{\mathbf{W} \sim P} [G(\boldsymbol{\theta}, \mathbf{W})]$$

is solved by generating  $(\mathbf{w}^i)_{i \in [N]}$ , a sample of size  $N$  (either taken from the distribution or created through historical samples), and solving the sample average approximation:

$$\min_{\boldsymbol{\theta} \in \Theta} \frac{1}{N} \sum_{i=1}^N G(\boldsymbol{\theta}, \mathbf{w}^i).$$

The approximation is often solved many times with different samples to create confidence intervals. It is often assumed that the resulting problem is tractable, but it is not necessarily so when  $N$  is large (Razaviyayn 2014). For a detailed discussion of the sample average approximation for stochastic approximation, please see e.g. Kleywegt et al. (2002).

We next show that Problem (2) satisfies the data concentration Assumption 3.

**Lemma 1** *Let  $l$  be a function that is at least once differentiable in  $\mathbf{z}$  and  $\boldsymbol{\theta}$ . Then  $f(\mathbf{z}, \boldsymbol{\theta}; [N]) = \frac{1}{N} \sum_{i=1}^N l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta})$  satisfies Assumption 3.*

*Remark 1* There is no need to assume convexity of the loss function  $l(\cdot, \cdot)$  for it to satisfy the data concentration Assumption 3. Such a convexity assumption is only necessary to prove global convergence results.

*Proof* Proof of Lemma 1: With  $f(\mathbf{z}, \boldsymbol{\theta}; [N]) = \frac{1}{N} \sum_{i=1}^N l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta})$ , we can invert Assumption 3 so that it states for all  $\epsilon > 0$ , we need to prove that there exists universal constants  $M, M'$  such that:

$$\mathbb{P} \left( \left| \frac{1}{N} \sum_{i=1}^N l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta}) - \frac{1}{n} \sum_{i \in S^n} l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta}) \right| > \epsilon \right) \leq \exp \left( -\frac{n\epsilon^2}{M} \right),$$

$$\mathbb{P} \left( \left\| \frac{1}{N} \sum_{i=1}^N \nabla l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta}) - \frac{1}{n} \sum_{i \in S^n} \nabla l(\mathbf{y}_i, \mathbf{x}_i, \mathbf{z}, \boldsymbol{\theta}) \right\| > \epsilon \right) \leq \exp \left( -\frac{n\epsilon^2}{M'(p_1 + p_2)} \right).$$

These two statements follow directly from Hoeffding's inequality in sampling without replacement (see e.g. Bardenet et al. (2015)).  $\square$

## 2.2 Sparse Data-Driven Risk Minimization

In Section 2.1 we included many key problems in data-driven optimization, but the format required the loss function to be separable in the space of samples. In this section, we give an example of a data-driven problem in which the objective function is not directly separable in the space of samples, and show that Assumption 3 still holds.

In particular, we consider the problem of sparse linear regression, where we seek at most  $k$  out of  $p$  factors in the design matrix  $\mathbf{X} \in \mathbb{R}^{N \times p}$  to explain the response data  $\mathbf{y} \in \mathbb{R}^N$ . The optimization problem can be written as:

$$\min_{\boldsymbol{\beta}: \|\boldsymbol{\beta}\|_0 \leq k} \frac{1}{N} \left( \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \frac{1}{\gamma} \|\boldsymbol{\beta}\|_2^2 \right). \quad (3)$$

Bertsimas and Van Parys (2020) showed that it can be reformulated as:

$$\min_{\mathbf{z} \in \mathcal{Z}} \frac{1}{N} \mathbf{y}^T \left( I_N + \gamma \sum_{i=1}^p z_i \mathbf{X}_i \mathbf{X}_i^T \right)^{-1} \mathbf{y}, \quad (4)$$

where here  $\mathcal{Z} = \{z \in \{0, 1\}^p \mid \sum_{i=1}^p z_i = k\}$ ,  $\mathbf{X}_i$  is the  $i$ th column of the design matrix, and  $\gamma$  controls the regularization. Note that this optimization function satisfies Assumptions 1 and 2 as the objective function is convex in  $\mathbf{z}$ , and the set  $\mathcal{Z}$  is indeed finite.

**Lemma 2** *The function  $f(\mathbf{z}, \boldsymbol{\theta}; [N]) = \frac{1}{N} \mathbf{y}^T (I_N + \gamma \sum_{i=1}^p z_i \mathbf{X}_i \mathbf{X}_i^T)^{-1} \mathbf{y}$  satisfies Assumption 3.*

The proof of the lemma is contained in Appendix A. The key characteristic that allowed this problem to also satisfy Assumption 3 is that Assumption 3 does not require the random sample estimate to be an *unbiased* estimate of the true loss function (and in this case it is indeed not true). Bertsimas et al. (2019) further discusses how to transform other sparse problems to a convex integer optimization problem that could be treated similarly. Bertsimas and Li (2018) discusses a similar problem in matrix completion where the data could be sampled in two dimensions.

### 3 Stochastic Cutting Planes and Theoretical Results

In this section, we present a formal description of the cutting plane algorithm, its stochastic cutting plane variant and show the global convergence of the stochastic cutting plane algorithm. The cutting plane algorithm introduced in Duran and Grossmann (1986) and generalized in Fletcher and Leyffer (1994) for Problem (1) iteratively rotates between solving a nonlinear subproblem and a master mixed-integer linear program (MILP) containing linear approximations of the functions involved. We first define the NLP subproblem as followed:

$$\begin{aligned} \text{NLP}(\mathbf{z}, [S_n]) &:= \min_{\boldsymbol{\theta} \in \Theta} f(\mathbf{z}, \boldsymbol{\theta}; [S_n]) \\ &\text{subject to } g_i(\mathbf{z}, \boldsymbol{\theta}) \leq 0, \quad i \in [m]. \end{aligned}$$

Then, we present the cutting plane algorithm in Algorithm 1. At every iteration, we first check if the current solution  $\mathbf{z}_t, \boldsymbol{\theta}_t, \eta_t$  is still a feasible solution for the  $t$ th iteration, and return  $(\mathbf{z}_t, \boldsymbol{\theta}_t, \eta_t)$  (as the optimal solution) if it is feasible. As the cutting plane iteratively adds constraints, if the previous solution is still feasible for the next iteration, then it is also optimal for the next iteration.



We further note that this guarantees that when the algorithm terminates at iteration  $t = T$ , the final cutting plane is at the optimal solution  $(\mathbf{z}^*, \boldsymbol{\theta}^*)$  of the  $T$ th iteration, i.e.,  $\mathbf{z}_T = \mathbf{z}^*, \boldsymbol{\theta}_T = \boldsymbol{\theta}^*$ , which would simplify the proof of some of our technical results. This is also how modern solvers implement cutting planes through lazy constraints Gurobi Optimization (2020).

---

**Algorithm 1** Cutting-plane algorithm.

---

```

1: procedure CUTPLANES( $\{\mathbf{d}_i\}_{i=1}^N$ )
2:    $t \leftarrow 1$ 
3:    $\mathbf{z}_1, \boldsymbol{\theta}_1 \leftarrow$  Initialization           # Heuristic Warm Start
4:    $\eta_1 \leftarrow lb$                              # Initialize feasible solution variable with lower bound
5:   while  $\eta_t < f(\mathbf{z}_t, \boldsymbol{\theta}_t; [N])$  or  $\exists j \in [m] 0 < g_j(\mathbf{z}_t, \boldsymbol{\theta}_t)$  do
6:     # while the current solution is not feasible for
     # the  $t$ th iteration
      $\mathbf{z}_{t+1}, \eta_{t+1} \leftarrow \arg \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta, \eta \geq lb} \eta$ 
7:     s.t.  $\eta \geq f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N]) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [t]$ 
            $0 \geq g_j(\mathbf{z}_i, \boldsymbol{\theta}_i) + \nabla g_j(\mathbf{z}_i, \boldsymbol{\theta}_i)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [t], \forall j \in [m]$ 
8:     # Solve the master problem
9:      $\boldsymbol{\theta}_{t+1} \leftarrow \arg_{\boldsymbol{\theta}} \text{NLP}(\mathbf{z}_{t+1}, [N])$  # Solve the NLP subproblem
10:     $t \leftarrow t + 1$ 
11:  end while
12:   $\mathbf{z}^*, \boldsymbol{\theta}^*, \eta^* \leftarrow \mathbf{z}_t, \boldsymbol{\theta}_t, \eta_t$ 
13:  return  $\eta^*, \mathbf{z}^*, \boldsymbol{\theta}^*$            # Return the filled matrix  $\mathbf{X}$ 
14: end procedure

```

---

In many examples such as risk minimization and sparse regression, the cutting plane algorithm (and its variants) has produced state-of-the-art results. However, one important drawback is its scalability with respect to data. It was noted in Franc and Sonnenburg (2008) that the scalability of the cutting-plane method for Support Vector Machines (SVMs), although significantly better than other exact methods, trails significantly behind stochastic methods such as stochastic gradient descent. This behavior is also observed in matrix completion in Bertsimas and Li (2018), where the full cutting plane algorithm scales less favorably than competing algorithms.

To understand the scalability of cutting plane algorithm, observe that at every iteration we are only solving a MILO problem, which has enjoyed great advances in recent years due to various algorithmic improvements (see Bixby and Rothberg (2007), Achterberg and Wunderling (2013) for details). As a result, the amount of time spent in the actual solver is usually not excessive. On the contrary, the cutting plane generation for  $f$  and  $g$  usually accounts for the majority of the running time. In data-driven optimization problems defined in Section 2 this is greatly exacerbated by the fact that the function  $f(\mathbf{z}, \boldsymbol{\theta}; [N])$  and its derivative  $\nabla f(\mathbf{z}, \boldsymbol{\theta}; [N])$  needs to be evaluated for all samples in  $[N]$ , which could easily exceed  $10^6$  samples or higher.

Therefore, with the objective to improve scalability for cutting-plane methods, we propose that instead of evaluating the full function and its derivative, we randomly select  $n$  samples without replacement in a set  $S^n$  to evaluate the function and its derivative at every iteration of the cutting plane. When  $n \ll N$ , we obtain a considerable speedup for the generation of the cutting plane. Concretely, we present the stochastic cutting plane algorithm in Algorithm 2.

---

**Algorithm 2** Stochastic Cutting-plane algorithm.

---

```

1: procedure STOCHASTIC CUTTING PLANES( $\{\mathbf{d}_i\}_{i=1}^N, \epsilon$ )
2:    $t \leftarrow 1$ 
3:    $\mathbf{z}_1, \boldsymbol{\theta}_1 \leftarrow$  Initialization           # Heuristic Warm Start
4:    $\eta_1 \leftarrow lb$                              # Initialize feasible solution variable with lower bound
5:    $S_1^n \leftarrow$  random  $n$ -sized subset of  $\{1, \dots, N\}$  # Initialize random subset of samples
6:   while  $\eta_t < f(\mathbf{z}_t, \boldsymbol{\theta}_t; [S_t^n])$  or  $\exists j \in [m] 0 < g_j(\mathbf{z}_t, \boldsymbol{\theta}_t)$  do
7:     # while the current solution is not feasible for
     # the  $t$ th iteration
      $\mathbf{z}_{t+1}, \eta_{t+1} \leftarrow \arg \min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta, \eta \geq lb} \eta$ 
8:     s.t.  $\eta \geq f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [t]$ 
            $0 \geq g_j(\mathbf{z}_i, \boldsymbol{\theta}_i) + \nabla g_j(\mathbf{z}_i, \boldsymbol{\theta}_i)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [t], \forall j \in [m]$ 
9:     # Solve the master problem
10:     $S_{t+1}^n \leftarrow$  random  $n$ -sized subset of  $\{1, \dots, N\}$ 
11:     $\boldsymbol{\theta}_{t+1} \leftarrow \arg_{\boldsymbol{\theta}} \text{NLP}(\mathbf{z}_{t+1}, [S_{t+1}^n])$  # Solve the NLP subproblem
12:     $t \leftarrow t + 1$ 
13:  end while
14:   $\mathbf{z}^*, \boldsymbol{\theta}^*, \eta^* \leftarrow \mathbf{z}_t, \boldsymbol{\theta}_t, \eta_t$ 
15:  return  $\eta^*, \mathbf{z}^*, \boldsymbol{\theta}^*$            # Return the filled matrix  $\mathbf{X}$ 
16: end procedure

```

---

### 3.1 Theoretical Guarantees

Fletcher and Leyffer (1994) proved that the cutting plane algorithm is able to converge to an optimal solution in finite steps. In this section, we show that the stochastic cutting plane algorithm, similarly, is able to converge to an  $\epsilon$ -optimal solution in finite steps with high probability. This is presented in the following theorem:

**Theorem 1** *Let  $\mathbf{d}_1, \dots, \mathbf{d}_n$  be given data, and assume that Problem (1) is feasible and has an optimal solution  $\mathbf{z}^*, \boldsymbol{\theta}^*$ . Under Assumptions 1-3, Algorithm 2 converges after a finite number of iterations  $K$  with probability one, and  $\mathbb{E}[K] \leq (e - 1)|\mathcal{Z}| + 1$ . When it terminates, it produces a feasible solution*

$\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*$  that satisfies:

$$f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; [N]) \leq f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) + \epsilon \quad (5)$$

with probability at least  $1 - (2K+1) \exp\left(\frac{-n\epsilon^2}{(2+\sqrt{p_1+p_2})^2 J}\right)$ , where  $J$  is an absolute constant.

The proof of Theorem 1 is contained in Appendix B. This theorem shows that the stochastic cutting plane algorithm finds an  $\epsilon$ -optimal solution with failure probability that is exponentially decaying with respect to the number of samples involved. We can see that the failure probability also decays exponentially in the optimality gap  $\epsilon$ , which indicates that the solutions to the stochastic cutting plane method are concentrated very close to the true solution, as long as a reasonable  $n$  is selected.

We further note that although in the worst case, the expected number of cutting planes needed could be very high  $((e-1)|\mathcal{Z}|+1)$ , this bound is only a small constant factor away from the worst case for the deterministic cutting plane algorithm ( $|\mathcal{Z}|$ ) as shown in (Duran and Grossmann 1986). Furthermore, although the number of cutting planes needed could be exponential in  $p_1$  (if  $\mathbf{z}$  are all binary, then  $\mathcal{Z} = 2^{p_1}$ ), in practice it usually converges much quicker (Fletcher and Leyffer 1994). Furthermore, in most application settings, we have  $p_1 \ll n \ll N$ , so even with the worst-case scenario of  $K \sim \exp(p_1)$ , the failure probability is still on the order of  $\exp(p_1 - O(n\epsilon^2)) \ll 1$ .

## 4 Computational Results

In this section, we present computational results from a variety of examples that demonstrate the versatility, scalability and accuracy of the stochastic cutting plane algorithm 2.

### 4.1 Sparse Regression

In this section, we apply Algorithm 2 to the sparse linear regression problem (4). As shown in Lemma 2, the sparse linear regression problem satisfies the data concentration assumption and thus the stochastic cutting plane algorithm 2 applies. We explore the scalability and accuracy of Algorithm 2.

To study the effectiveness of Algorithm 2 we generate data according to  $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$  where the entries of the design matrix  $X_{ij} \sim N(0, 1)$  independently. A random set  $S$  of size  $k$  is chosen from  $\{1, \dots, p\}$ , and the coefficients are

chosen so that  $\beta_j = 0$ ,  $\forall j \notin S$ , and  $\beta_j \sim N(0, 1)$ ,  $\forall j \in S$  independently. This ensures the exact sparsity of rank  $k$  for the resulting model. The noise is sampled independently with  $\epsilon_i \sim N(0, \sigma)$ .

For each parameter combination  $(N, p, k, \sigma)$ , we record the time  $T$  and the Mean Average Percentage Error (MAPE) of both the standard cutting-plane algorithm (Algorithm 1) and the stochastic cutting plane algorithm (SCP, Algorithm 2). The  $\gamma$  parameter for both algorithms are selected through cross-validation on a validation set of size  $N$  using the same setup as above. We set the tolerance in the cutting plane algorithm to  $\epsilon = 10^{-4}$ . The testing set is of size  $N$  sampled using the true  $\beta$  and independent  $X_{ij}$  and  $\epsilon_i$  with the same setup as above. Table 1 shows the mean results over 10 random generations of the dataset of the size  $(N, p, k, \sigma)$ , separated into 4 blocks where each variable is varied over a range of possibilities. For these experiments, we used  $n = \min\{N, 10\sqrt{N}\}$  as the sampling size for stochastic cutting planes. In Section 4.4, we detail further experiments on varying the sample size. The testing environment has a six-core i7-5820k CPU with 24GB of RAM. Both algorithms are implemented in pure Julia with Julia v1.4.2.

	$N$	$p$	$k$	$\sigma$	Cutting Planes		SCP	
					$T$	MAPE	$T$	MAPE
$N$	$10^3$	$10^2$	10	0.1	2.4s	4.4%	1.9s	4.4%
	$10^4$	$10^2$	10	0.1	5.7s	3.9%	3.7s	3.9%
	$10^5$	$10^2$	10	0.1	170s	4.5%	7.8s	4.5%
	$10^6$	$10^2$	10	0.1	1356s	3.1%	46s	3.1%
$p$	$10^5$	$10^2$	10	0.1	170s	4.5%	7.8s	4.5%
	$10^5$	$10^3$	10	0.1	2073s	4.1%	98.5s	4.1%
	$10^5$	$10^4$	10	0.1	29770s	4.7%	945s	4.7%
$k$	$10^5$	$10^2$	10	0.1	170s	4.5%	7.8s	4.5%
	$10^5$	$10^2$	20	0.1	1621s	2.9%	41.9s	2.9%
	$10^5$	$10^2$	50	0.1	30608s	1.8%	937s	1.8%
$\sigma$	$10^5$	$10^2$	10	0.1	170s	4.5%	7.8s	4.5%
	$10^5$	$10^2$	10	0.2	197s	6.7%	8.9s	6.7%
	$10^5$	$10^2$	10	0.3	231s	9.8%	10.7s	9.8%

**Table 1** Comparison of cutting planes and stochastic cutting planes on synthetic data for sparse regression.

We see that as  $N$  grows, the advantage of stochastic cutting planes in using  $n = O(\sqrt{N})$  samples to compute the cutting plane quickly increases. At the largest scale with  $N = 10^6$ , stochastic cutting planes generates over a  $30\times$  speedup compared to the standard cutting plane algorithm. Furthermore, we see that the stochastic cutting plane converges to the exact same solution as the standard cutting planes, as reflected the identical MAPE figures. This is a reflection of the exponentially vanishing failure probability as indicated in Theorem 1. Such order of magnitude speedup is also preserved when we vary  $p$ , sparsity  $k$ , or the amount of noise  $\sigma$  in the additional blocks of experiments without any loss of accuracy.

## 4.2 Support Vector Machines

In this section, we apply Algorithm 2 to the problem of support vector machines for classification. In this particular case, the variables are all continuous, and therefore we focus on utilizing the cutting-plane algorithm to solve the NLP subproblem as defined in Algorithm 2 since the set of integer variables  $\mathcal{Z} = \emptyset$ .

Specifically, we consider a dataset where we have  $(\mathbf{x}_i, y_i)$ ,  $i \in [N]$  and  $y_i \in \{-1, +1\}$ , and we wish to create a classifier of the form  $f(\mathbf{x}) = \text{Sign}(\boldsymbol{\theta}^T \mathbf{x})$ . Then the SVM problem with  $\ell_2$  regularization can be written as:

$$\min_{\boldsymbol{\theta}, \epsilon \geq 0} \frac{1}{2} \|\boldsymbol{\theta}\|^2 + CR(\boldsymbol{\theta}), \quad (6)$$

where  $C$  is a user-defined constant controlling the degree of regularization, and:

$$R(\boldsymbol{\theta}) = \max_{c \in \{0,1\}^N} \left\{ \frac{1}{N} \sum_{i=1}^N c_i - \frac{1}{N} \sum_{i=1}^N c_i y_i (\mathbf{w}^T \mathbf{x}_i) \right\}. \quad (7)$$

In this setting all variables are continuous. We follow the formulation in classical papers of utilizing cutting planes in SVMs (e.g. Franc and Sonnenburg (2008), Joachims (2006)) by using cutting planes (i.e. linearizing the objective) only on the  $R(\boldsymbol{\theta})$  portion of the objective. It is easy to see that  $R(\boldsymbol{\theta})$  satisfies the data concentration assumption by recognizing the form of  $R(\boldsymbol{\theta})$  as a maximum over a finite set of expressions in the form of those introduced in Section 2.1, which satisfies Assumption 3 by Lemma 1.

For any  $\boldsymbol{\theta}$ , the following expression is a valid subgradient for  $R(\boldsymbol{\theta})$ :

$$\partial R(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N c_i y_i \mathbf{x}_i, \quad (8)$$

with

$$c_i = \begin{cases} 1, & y_i(\boldsymbol{\theta}^T \mathbf{x}_i) < 1, \\ 0, & \text{otherwise.} \end{cases}$$

To conduct the experiment, we utilize the forest covertype dataset from the UCI Machine Learning Repository Dua and Graff (2017) that has 580,012 samples with 54 features, and the objective is to predict the type of forest cover. We binarize the objective variable to predict if the forest cover belongs to Class 2 (the most frequent class at 48.9%) or not, following a similar treatment in Joachims (2006).

We fix  $C = 10^6$  as defined in Equation (6), and set the tolerance in the cutting plane to  $\epsilon = 10^{-4}$ . Then, we randomly sample two disjoint datasets, each of size

$N$ , for training and testing from the master dataset. We experiment with different sizes of  $N$  to observe the speedup and performance of stochastic cutting planes (SCP) over standard cutting planes. We again set  $n = \min\{N, 10\sqrt{N}\}$  as the sampling size for stochastic cutting planes. The results, recording time  $T$  and the out-of-sample accuracy on the testing set  $ACC$  over 10 random runs are contained in Table 2. We further compare to LIBSVM (Chang and Lin 2011), an established reference SVM implementation using Sequential Minimal Optimization (SMO), and stochastic gradient descent (SGD), which is widely regarded as the fastest algorithm for solving SVM problems (at the expense of some accuracy). The testing environment has a six-core i7-5820k CPU with 24GB of RAM. Except for SGD, all algorithms are implemented in pure Julia with Julia v1.4.2 to minimize the difference in conditions. For SGD, we utilize the optimized formulation contained in scikit-learn V0.22.2 and Python V3.7.4.

$N$	$p$	Cutting Planes		SCP		LIBSVM		SGD	
		$T$	ACC	$T$	ACC	$T$	ACC	$T$	ACC
$10^3$	54	107.6s	74.9%	7.5s	72.3%	0.75s	74.2%	0.01s	71.8%
$10^4$	54	145.7s	75.5%	10.7s	74.8%	30.9s	75.0%	0.09s	73.4%
$10^5$	54	439.2s	75.9%	32.1s	75.3%	960.5s	75.7%	0.45s	74.3%

**Table 2** Comparison of cutting planes and stochastic cutting planes on the covertype dataset for support vector machines.

We see that stochastic cutting planes achieve a significant order-of-magnitude speedup compared to the standard cutting plane while maintaining a comparable accuracy with minimum degradation. At the largest scales it is over  $13\times$  faster than the standard cutting planes formulation and  $30\times$  faster than LIBSVM with only a degradation of  $\sim 0.5\%$  in out-of-sample accuracy. This reflects the scalability of the stochastic cutting plane method versus other methods in large-scale data. Stochastic gradient descent is still significantly faster, but at the further expense of accuracy of about 1%. We note that however, the running time of SCP appears to grow slower (roughly proportional to  $\sqrt{n}$ ) compared to SGD (roughly proportional to  $n$ ).

We note that at the largest scale with  $N = 10^5$ , each stochastic cutting plane uses  $n = 10^{3.5}$  samples, and yet it manages to be  $3\times$  faster than the full cutting plane algorithm when ran on  $10^3$  samples while its accuracy is closer to the cutting plane algorithm when ran on  $10^4$  samples. This shows that the randomness of the data samples selected by each cutting plane is key, and the performance of the stochastic cutting plane using a random sample of size  $n$  cannot be simply replicated by choosing  $n$  samples and running the standard cutting plane algorithm.

### 4.3 The Stochastic Knapsack Problem

We further apply stochastic cutting planes to SAA versions of stochastic optimization problems. In particular, we test stochastic cutting planes on the classic *static stochastic knapsack problem* (SSKP). A SSKP is a resource allocation problem where a decision maker has to choose a subset of  $k$  known alternative projects to take on. For this purpose, a known quantity  $q$  of relatively low-cost resource is available, while any further resource required can be obtained at a known cost of  $c$  per unit of resource. The amount  $W_i$  of resource required by project  $i$  is not known exactly but the decision maker has an estimate of the probability distribution. Each project  $i$  has an expected net reward of  $r_i$ . Thus, the optimization problem can be formulated as:

$$\max_{\mathbf{z} \in \{0,1\}^k} \sum_{i=1}^k r_i z_i - c \mathbb{E} \left[ \sum_{i=1}^k W_i z_i - q \right]^+$$

Where  $[\cdot]^+ := \max(\cdot, 0)$  and the expectation is taken over the probability distribution of  $\mathbf{W} = [W_1, \dots, W_k]$ . This objective function appears in many real-life problems such as airline crew scheduling, shortest path problems, and others, often with further constraints on  $\mathbf{z}$  to limit the optimization over  $\mathbf{z} \in \mathcal{Z} \subset \{0,1\}^k$ .

We here consider the SAA version of SSKP where  $N$  realizations of  $\mathbf{W}$  are available (either by simulation or historical data) in the form of  $[\mathbf{W}^1, \dots, \mathbf{W}^N]$ . Then the SAA version of SSKP can be written as:

$$\max_{\mathbf{z} \in \{0,1\}^k} \sum_{i=1}^k r_i z_i - \frac{c}{N} \sum_{j=1}^N \left[ \sum_{i=1}^k W_i^j z_i - q \right]^+ \quad (9)$$

This problem is an integer concave optimization problem, and we can utilize cutting planes on the cost part of the objective:

$$C(\mathbf{z}) = \frac{c}{N} \sum_{j=1}^N \left[ \sum_{i=1}^k W_i^j z_i - q \right]^+.$$

$C(\mathbf{z})$  satisfies the data concentration assumption since it satisfies the form in Lemma 1. Its derivative is easily calculated as:

$$(\nabla C(\mathbf{z}))_i = \frac{c}{N} \sum_{j=1}^N W_i^j \mathbb{1} \left\{ \sum_{i=1}^k W_i^j z_i - q \geq 0 \right\}.$$

To generate the specific instances for this problem, we follow the setup in Kleywegt et al. (2002). The rewards follow  $r_i \sim \text{Unif}[10, 20]$ , while  $W_i \sim N(\mu_i, \sigma_i)$  where  $\mu_i \sim \text{Unif}[20, 30]$  and  $\sigma_i \in \text{Unif}[5, 15]$ . For all instances, we set the per unit penalty  $c = 4$ , and set the initial available inventory  $q =$

$\max(k, 20)$  (the lower bound is so that the optimal solution is not the trivial 0 vector). We investigate different combinations of  $(N, k)$  to uncover how the performance of stochastic cutting plane scales with problem size and samples generated.

We compare stochastic cutting planes (with sampling size  $n = \min\{N, 10\sqrt{N}\}$ ) to the standard cutting plane algorithm, and further with the standard linear reformulation of (9) using auxiliary variables  $\mathbf{x} \in \mathbb{R}^n$  as outlined in Kleywegt et al. (2002):

$$\max_{\mathbf{z} \in \{0,1\}^k, \mathbf{x} \geq 0} \sum_{i=1}^k r_i z_i - \frac{c}{N} \sum_{j=1}^N z_j \quad (10)$$

$$s.t. \quad x_j \geq \sum_{i=1}^k W_i^j z_i - q, \quad \forall j \in [N]. \quad (11)$$

The testing environment has a six-core i7-5820k CPU with 24GB of RAM. All algorithms are implemented in pure Julia with Julia v1.4.2 to minimize the difference in conditions, and we utilize Gurobi V9.0.1 as the integer programming solver. We record two statistics: the time taken to solve the formulation ( $T$ ), and the objective value ( $Obj$ ). For ease of comparison, we normalized the objective value of the optimal solution to 100%. Each statistic is shown as an average of 20 randomly generated instances:

$N$	$k$	Cutting Planes		SCP		Linear Reformulation	
		$T$	Obj	$T$	Obj	$T$	Obj
$10^3$	10	0.33s	100%	0.24s	99.7%	7.0s	100%
$10^4$	10	0.89s	100%	0.27s	99.9%	70.3s	100%
$10^5$	10	3.6s	100%	0.35s	100%	3760s	100%
$10^6$	10	61.3s	100%	0.97s	100%	N/A	N/A
$10^3$	20	0.42s	100%	0.25s	99.9%	8.5s	100%
$10^4$	20	1.1s	100%	0.39s	99.9%	437s	100%
$10^5$	20	9.3s	100%	0.46s	100%	7328s	100%
$10^6$	20	124s	100%	1.7s	100%	N/A	N/A
$10^3$	50	1.3s	100%	1.0s	100%	6.1s	100%
$10^4$	50	5.3s	100%	1.1s	100%	260s	100%
$10^5$	50	51.8s	100%	2.5s	100%	4005s	100%
$10^6$	50	706s	100%	6.9s	100%	N/A	N/A

**Table 3** Comparison of cutting planes and stochastic cutting planes for SSKP. N/A means that no instance completed running under 10000s.

We first observe that both the standard and the stochastic cutting plane formulation outperform the linearized reformulation significantly in time needed. This is not surprising as the linear reformulation requires  $N$  auxiliary variables to model the  $[\cdot]^+$  function, which severely impacts solve time.

We further notice that stochastic cutting planes is able to achieve a near optimal solution across all combinations of  $N$  and  $k$  tested. In the largest




instances (where  $N = 10^6$ ), it was able to almost always return the same solution as the general cutting plane algorithm while spending two orders of magnitude less time. This again demonstrates the significant speedup that stochastic cutting plane provides.

#### 4.4 Sample Size Selection

The key parameter in the method of stochastic cutting planes is the sample size  $n$ , as it drives both the accuracy in Theorem 1, and the speedup compared to the full cutting plane algorithm. Therefore, in this section, we explore how the sample size  $n$  affects the quality of solutions in various problems.

Specifically, we test different combinations of  $(N, n)$  on sparse regression with  $p = 100$ ,  $k = 10$  using the same setup as in Section 4.1 over 10 runs each, and record the percentage of cases that the stochastic cutting plane algorithm returned the same solution as the optimal solution. From Figure 1 we can




**Fig. 1** Percentage of Runs where the Stochastic Cutting Planes agrees with true solution on different  $(N, n)$  combinations for sparse regression with  $p = 100$  and  $k = 10$ .

see that for the task of sparse regression, the sample size  $n$  needed so that the stochastic solution always coincides with the true solution seems to grow slower than  $O(\sqrt{N})$  but faster than  $O(\log(N))$ . Therefore, in cases where accuracy of the solution is important, it would seem to be prudent, at least for the

task of sparse regression, to select  $O(\sqrt{N})$  samples to ensure that the solution does indeed coincide, as we showed in Table 1. However, in scenarios where accuracy is less important than scalability, then the practitioner can explore even smaller  $n$  such as  $O(\sqrt[3]{N})$  as it is clear from the graph that  $O(\sqrt{N})$  is a conservative bound as  $N \rightarrow \infty$ .

We further explore this sample size relationship in the context of the SSKP problem. We follow the same setup as in Section 4.3, and set  $k = 50$ . We record the objective value of SCP as a percentage of the optimal solution, averaged over 20 runs.



**Fig. 2** Average performance of the solution reached by Stochastic Cutting Planes as a percentage of true optimal objective on different  $(N, n)$  combinations for SSKP with  $k = 50$ . Nodes with  $< 99.9\%$  of optimality are colored red.

Similar to sparse regression, we see from Figure 2 that the chosen sample size  $n = 10\sqrt{N}$  is a quite conservative and the average objective across 20 runs is always within 0.01% of the optimal solution. We further see that even at scales of  $n = 10\log(N)$ , the objective value of the solution reached by SCP is always within 0.5% of the true solution. Therefore, in problems where absolute optimality is not necessary, further running time gains (compared to what is demonstrated in this work) could be achieved by using a smaller sample size.

## 5 Conclusions

We have presented stochastic cutting planes for data-driven convex MINLPs. Our theoretical results show that the proposed methodology is able to obtain the optimal solution with exponentially vanishing failure probability. Our numerical experiments demonstrate that on a wide variety of problems, the stochastic cutting planes algorithm is able to achieve an order-of-magnitude speedup compared to the standard cutting planes algorithm without a significant sacrifice in accuracy, and is competitive with state-of-the-art algorithms.

## References

- Achterberg T, Wunderling R (2013) Mixed integer programming: Analyzing 12 years of progress. In: Facets of combinatorial optimization, Springer, pp 449–481
- Alley M, Biggs M, Hariss R, Herrmann C, Li M, Perakis G (2019) Pricing for heterogeneous products: Analytics for ticket reselling. Available at SSRN 3360622
- Bardenet R, Maillard OA, et al. (2015) Concentration inequalities for sampling without replacement. *Bernoulli* 21(3):1361–1385
- Bertsimas D, Li ML (2018) Interpretable matrix completion: A discrete optimization approach. arXiv preprint arXiv:181206647
- Bertsimas D, Van Parys B (2020) Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *The Annals of Statistics* 48(1):300–323
- Bertsimas D, Cory-Wright R, Pauphilet J (2019) A unified approach to mixed-integer optimization: Nonlinear formulations and scalable algorithms. arXiv preprint arXiv:190702109
- Bixby R, Rothberg E (2007) Progress in computational mixed integer programming—a look back from the other side of the tipping point. *Annals of Operations Research* 149(1):37–41
- Blundell RW, Powell JL (2004) Endogeneity in semiparametric binary response models. *The Review of Economic Studies* 71(3):655–679
- Boucheron S, Lugosi G, Massart P (2013) Concentration inequalities: A nonasymptotic theory of independence. Oxford university press
- Calafore G, Campi MC (2005) Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming* 102(1):25–46
- Chang CC, Lin CJ (2011) Libsvm: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)* 2(3):1–27
- Dua D, Graff C (2017) UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>
- Duran MA, Grossmann IE (1986) An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical programming* 36(3):307–339
- Feller W (2008) An introduction to probability theory and its applications, vol 2. John Wiley & Sons
- Fletcher R, Leyffer S (1994) Solving mixed integer nonlinear programs by outer approximation. *Mathematical programming* 66(1-3):327–349
- Franc V, Sonnenburg S (2008) Optimized cutting plane algorithm for support vector machines. In: Proceedings of the 25th international conference on Machine learning, pp 320–327

- Franc V, Sonnenburg S (2009) Optimized cutting plane algorithm for large-scale risk minimization. *Journal of Machine Learning Research* 10(10)
- Franc V, Sonnenburg S, Werner T (2011) Cutting plane methods in machine learning. *Optimization for Machine Learning* pp 185–218
- Gurobi Optimization L (2020) Gurobi optimizer reference manual. URL <http://www.gurobi.com>
- Joachims T (2006) Training linear svms in linear time. In: *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp 217–226
- Kleywegt AJ, Shapiro A, Homem-de Mello T (2002) The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2):479–502
- Parikh S (1976) Approximate cutting planes in nonlinear programming. *Mathematical Programming* 11(1):194–198
- Razaviyayn M (2014) Successive convex approximation: Analysis and applications. PhD thesis, University of Minnesota
- Rivers D, Vuong QH (1988) Limited information estimators and exogeneity tests for simultaneous probit models. *Journal of econometrics* 39(3):347–366
- Stewart GW (1990) *Matrix perturbation theory*. Citeseer
- Tropp JA (2012) User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics* 12(4):389–434
- Volkov Y, Zavriev S (1997) A general stochastic outer approximations method. *SIAM journal on control and optimization* 35(4):1387–1421
- Wei J, Realff MJ (2004) Sample average approximation methods for stochastic minlps. *Computers & Chemical Engineering* 28(3):333–346
- Woodbury MA (1949) *The stability of out-input matrices*. Chicago, IL 93

## A Proof of Lemma 2

We only prove the lemma for the function  $f(\mathbf{z}, \boldsymbol{\theta}; [N])$  as the derivative  $\nabla f$  follows the exact same procedure just with additional notation.

First, by the matrix inversion lemma (Woodbury 1949), for any feasible solution  $\mathbf{z} \in \mathcal{Z}$ , we can rewrite  $f(\mathbf{z}, \boldsymbol{\theta}; [N])$  as:

$$\begin{aligned} f(\mathbf{z}, \boldsymbol{\theta}; [N]) &= \frac{1}{N} \mathbf{y}^T \left( I_N - \mathbf{X}_z \left( \frac{I_k}{\gamma} + \mathbf{X}_z^T \mathbf{X}_z \right)^{-1} \mathbf{X}_z \right) \mathbf{y} \\ &= \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{N} \mathbf{y}^T \mathbf{X}_z \left( \frac{I_k}{\gamma} + \mathbf{X}_z^T \mathbf{X}_z \right)^{-1} \mathbf{X}_z \mathbf{y} \\ &= \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{N} \mathbf{y}^T \mathbf{V} \left( \frac{I_k}{\gamma} + \mathbf{V}^T \mathbf{V} \right)^{-1} \mathbf{V} \mathbf{y}, \end{aligned} \quad (\text{A1})$$

where  $\mathbf{V} = \mathbf{X}_z \in \mathbb{R}^{N \times k}$  is the subset of the design matrix  $\mathbf{X}$  with only the columns  $\{i \mid z_i = 1\}$ . For simplicity of notation, we use  $\mathbf{V}$  to suppress the dependence on  $\mathbf{z}$ .

We let  $S$  be a subset of  $[N]$  of size  $n$  (we suppress the dependence of  $S$  on  $n$  to reduce notation complexity). Then, we have, similar to Equation (A1):

$$f(\mathbf{z}, \boldsymbol{\theta}; S) = \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S - \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \frac{I_k}{\gamma} + \mathbf{V}_S^T \mathbf{V}_S \right)^{-1} \mathbf{V}_S \mathbf{y}_S, \quad (\text{A2})$$

where  $\mathbf{y}_S$  is the subvector of  $\mathbf{y}$  with only the entries in set  $S$ , and similar for  $\mathbf{V}$ . Therefore, what we need to prove is that:

$$\begin{aligned} \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{N} \mathbf{y}^T \mathbf{V} \left( \frac{I_k}{\gamma} + \mathbf{V}^T \mathbf{V} \right)^{-1} \mathbf{V} \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S + \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \frac{I_k}{\gamma} + \mathbf{V}_S^T \mathbf{V}_S \right)^{-1} \mathbf{V}_S \mathbf{y}_S \right| \leq M \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}} \right) \\ \geq 1 - \epsilon. \end{aligned} \quad (\text{A3})$$

We now introduce and prove lemmas that allow us to prove Equation (A3):

**Lemma 1** *Let  $S$  be a random subset of  $[N]$  of size  $n$ . Then, we have*

$$\mathbb{P} \left( \left| \frac{\mathbf{y}_S^T \mathbf{y}_S}{n} - \frac{\mathbf{y}^T \mathbf{y}}{N} \right| \leq \sqrt{\frac{A \log(\frac{1}{\epsilon})}{n}} \right) \geq 1 - \epsilon, \quad (\text{A4})$$

$$\mathbb{P} \left( \left\| \frac{\mathbf{y}_S^T \mathbf{V}_S}{n} - \frac{\mathbf{y}^T \mathbf{V}}{N} \right\| \leq \sqrt{\frac{B \log(\frac{1}{\epsilon})}{n}} \right) \geq 1 - \epsilon. \quad (\text{A5})$$

*Proof* Proof: Notice that we have:

$$\begin{aligned} \mathbf{y}^T \mathbf{y} &= \sum_{i=1}^N y_i^2, & \mathbf{y}^T \mathbf{V} &= \sum_{i=1}^N y_i v_i, \\ \mathbf{y}_S^T \mathbf{y}_S &= \sum_{i \in S} y_i^2, & \mathbf{y}_S^T \mathbf{V}_S &= \sum_{i \in S} y_i v_i. \end{aligned}$$

Therefore, if we treat  $y_1^2, \dots, y_N^2$  as a finite population, then  $\mathbf{y}_S^T \mathbf{y}_S$  is a random sample of  $n$  points drawn without replacement from that set, and similarly for  $\mathbf{y}_S^T \mathbf{V}_S$ . Then, the required inequality follows directly from Hoeffding's Inequality in the case without replacement, stated below:

**Proposition 1 (Hoeffding's Inequality)** Let  $\mathcal{X} = (x_1, \dots, x_n)$  be a finite population of  $N$  points and  $X_1, \dots, X_n$  be a random sample drawn without replacement from  $\mathbf{X}$ . Let:

$$a = \min_{1 \leq i \leq n} x_i, \quad \text{and} \quad b = \max_{1 \leq i \leq n} x_i.$$

Then, for all  $\epsilon > 0$ , we have:

$$\mathbb{P} \left( \left| \frac{\sum_{i=1}^n X_i}{n} - \mu \right| \geq \epsilon \right) \leq 2 \exp \left( -\frac{2n\epsilon^2}{(b-a)^2} \right). \quad (\text{A6})$$

For proof of this proposition, see for example Boucheron et al. (2013).  $\square$

Then, we have the following lemma:

**Lemma 2**

$$\mathbb{P} \left( \left\| \left( \frac{I_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} - \left( \frac{I_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \right\| \leq \sqrt{\frac{H \log(\frac{1}{\epsilon})}{n}} \right) \geq 1 - \epsilon. \quad (\text{A7})$$

To prove this, we first introduce a matrix analog of the well-known Chernoff bound, the proof of which can be found in Tropp (2012):

**Lemma 3** Let  $\mathcal{X} \in \mathbb{R}^{k \times k}$  be a finite set of positive-semidefinite matrices, and suppose that:

$$\max_{\mathbf{X} \in \mathcal{X}} \lambda_{\max}(\mathbf{X}) \leq D,$$

where  $\lambda_{\min}/\lambda_{\max}$  is the minimum/maximum eigenvalue function. Sample  $\{\mathbf{X}_1, \dots, \mathbf{X}_\ell\}$  uniformly at random without replacement. Let

$$\mu_{\min} := \ell \cdot \lambda_{\min}(\mathbb{E} \mathbf{X}_1), \quad \mu_{\max} := \ell \cdot \lambda_{\max}(\mathbb{E} \mathbf{X}_1).$$

Then

$$\mathbb{P} \left\{ \lambda_{\min} \left( \sum_j \mathbf{X}_j \right) \leq (1 - \delta) \mu_{\min} \right\} \leq k \cdot \exp \left( \frac{-\delta^2 \mu_{\min}}{4D} \right), \quad \text{for } \delta \in [0, 1),$$

$$\mathbb{P} \left\{ \lambda_{\max} \left( \sum_j \mathbf{X}_j \right) \leq (1 + \delta) \mu_{\max} \right\} \leq k \cdot \exp \left( \frac{-\delta^2 \mu_{\max}}{4D} \right), \quad \text{for } \delta \geq 0.$$

Now we proceed with the proof.

*Proof* Proof of Lemma 2: First, we do a QR decomposition of the matrix  $\mathbf{V} = \mathbf{Q}\mathbf{R}$  where  $\mathbf{Q} \in \mathbb{R}^{N \times k}$  has orthogonal columns and  $\mathbf{R} \in \mathbb{R}^{k \times k}$  is upper triangular. We further normalize the decomposition so we have that  $\frac{\mathbf{Q}^T \mathbf{Q}}{N} = \mathbf{I}_k$ . Then, we have that  $\mathbf{V}_S = \mathbf{Q}_S \mathbf{R}$ , and we note that:

$$\mathbf{Q}^T \mathbf{Q} = \sum_{i=1}^N \mathbf{q}_i \mathbf{q}_i^T,$$

$$\mathbf{Q}_S^T \mathbf{Q}_S = \sum_{i \in S} \mathbf{q}_i \mathbf{q}_i^T,$$

where  $\mathbf{q}_i \mathbf{q}_i^T \in \mathbb{R}^{k \times k}$  rank-one positive semi-definite matrices. Therefore, we can take  $\mathbf{Q}_S^T \mathbf{Q}_S$  as a random sample of size  $n$  from the set  $\mathcal{X} = \{\mathbf{q}_i \mathbf{q}_i^T\}_{i \in [N]}$ , which satisfies the conditions in Lemma 3 with  $D = O(k)$ . Therefore, we use Lemma 3 to bound  $\mathbf{Q}_S^T \mathbf{Q}_S$  that:

$$\begin{aligned} \mathbb{P} \left\{ \lambda_{\min} \left( \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} \right) \leq (1 - \delta) \right\} &\leq k \cdot \exp \left( \frac{-\delta^2 n}{kD} \right), \\ \mathbb{P} \left\{ \lambda_{\max} \left( \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} \right) \geq (1 + \delta) \right\} &\leq k \cdot \exp \left( \frac{-\delta^2 n}{kD} \right), \end{aligned}$$

for some absolute constant  $D$ . Some rearrangement gives:

$$\mathbb{P} \left\{ \lambda_{\min} \left( \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} \right) \geq 1 - \sqrt{\frac{kD \log \left( \frac{2k}{\epsilon} \right)}{n}} \text{ and } \lambda_{\max} \left( \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} \right) \leq 1 + \sqrt{\frac{kD \log \left( \frac{2k}{\epsilon} \right)}{n}} \right\} \geq 1 - \epsilon. \quad (\text{A8})$$

Now since  $\frac{\mathbf{Q}^T \mathbf{Q}}{N} = \mathbf{I}_k$ , we have:

$$\lambda_{\min} \left( \frac{\mathbf{Q}^T \mathbf{Q}}{N} \right) = \lambda_{\max} \left( \frac{\mathbf{Q}^T \mathbf{Q}}{N} \right) = 1. \quad (\text{A9})$$

Combining equation (A9) and (A8) we obtain

$$\mathbb{P} \left\{ \left\| \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} - \frac{\mathbf{Q}^T \mathbf{Q}}{N} \right\| \leq \sqrt{\frac{kD \log \left( \frac{2k}{\epsilon} \right)}{n}} \right\} \geq 1 - \epsilon. \quad (\text{A10})$$

Thus, we have

$$\begin{aligned} \mathbb{P} \left\{ \left\| \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} - \frac{\mathbf{V}^T \mathbf{V}}{N} \right\| \leq \|\mathbf{R}\|^2 \sqrt{\frac{kD \log \left( \frac{2k}{\epsilon} \right)}{n}} \right\} \\ \geq \mathbb{P} \left\{ \left\| \frac{\mathbf{Q}_S^T \mathbf{Q}_S}{n} - \frac{\mathbf{Q}^T \mathbf{Q}}{N} \right\| \leq \sqrt{\frac{kD \log \left( \frac{2k}{\epsilon} \right)}{n}} \right\} \geq 1 - \epsilon. \end{aligned}$$

We let  $D' = \log(2) \|\mathbf{R}\|^4 D$  and obtain

$$\mathbb{P} \left\{ \left\| \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} - \frac{\mathbf{V}^T \mathbf{V}}{N} \right\| \leq \sqrt{\frac{kD' \log \left( \frac{k}{\epsilon} \right)}{n}} \right\} \geq 1 - \epsilon. \quad (\text{A11})$$

From this, we can easily see that there exists a constant  $G$  such that:

$$\mathbb{P} \left\{ \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right) - \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right) \right\| \leq \sqrt{\frac{kG \log \left( \frac{k}{\epsilon} \right)}{n}} \right\} \geq 1 - \epsilon, \quad (\text{A12})$$

as  $\frac{\mathbf{I}_k}{n\gamma}$  and  $\frac{\mathbf{I}_k}{N\gamma}$  are of order of magnitude  $O\left(\frac{1}{n}\right)$ . Now let us introduce another lemma from matrix perturbation theory (for proof, see e.g. Stewart (1990)):

**Lemma 4** *Let  $\mathbf{A}, \mathbf{B}$  be invertible matrices. Then, we have the following bound:*

$$\|\mathbf{A}^{-1} - \mathbf{B}^{-1}\| \leq \|\mathbf{A}^{-1}\| \|\mathbf{B}^{-1}\| \|\mathbf{A} - \mathbf{B}\|. \quad (\text{A13})$$

Thus, now let  $H = G \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \right\|^2 \left\| \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \right\|^2$ , be a constant. Then we have, using Lemma 4 and Equation (A12):

$$\begin{aligned} & \mathbb{P} \left\{ \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} - \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \right\| \leq \sqrt{\frac{kH \log\left(\frac{k}{\epsilon}\right)}{n}} \right\} \\ & \geq \mathbb{P} \left\{ \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right) - \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right) \right\| \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \right\| \left\| \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \right\| \leq \sqrt{\frac{kH \log\left(\frac{k}{\epsilon}\right)}{n}} \right\} \\ & = \mathbb{P} \left\{ \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right) - \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right) \right\| \leq \sqrt{\frac{kG \log\left(\frac{k}{\epsilon}\right)}{n}} \right\} \\ & \geq 1 - \epsilon. \end{aligned}$$

As required. We then absorb  $H' = k \log(k)H$  as  $k$  is a dimension independent from the variable dimension  $p$  and the sample dimension  $n$ .  $\square$

With Lemma 1 and 2, we are now ready to prove the main result in Equation (A3). First, let  $M$  be a constant such that:

$$M \geq \max \left\{ \sqrt{4A}, 36 \left\| \left( \frac{\mathbf{I}_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \frac{\mathbf{V} \mathbf{y}}{N} \right\|_B, 36 \left\| \frac{\mathbf{y}_S^T \mathbf{V}_S}{n} \right\| \left\| \frac{\mathbf{V} \mathbf{y}}{N} \right\|^2, H, 36 \left\| \left( \frac{\mathbf{I}_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \frac{\mathbf{V}_S \mathbf{y}_S}{n} \right\|_B^2 \right\}.$$

We introduce the Frechet inequalities, which state that:

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n) \geq \mathbb{P}(A_1) + \mathbb{P}(A_2) + \dots + \mathbb{P}(A_n) - (n-1).$$

Then we have, by Frechet's inequalities:

$$\begin{aligned} & \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{N} \mathbf{y}^T \mathbf{V} \left( \frac{\mathbf{I}_k}{\gamma} + \mathbf{V}^T \mathbf{V} \right)^{-1} \mathbf{V} \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S + \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \frac{\mathbf{I}_k}{\gamma} + \mathbf{V}_S^T \mathbf{V}_S \right)^{-1} \mathbf{V}_S \mathbf{y}_S \right| \leq M \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) \\ & \geq \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S \right| \leq \frac{M}{2} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) + \\ & + \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{V} \left( \frac{\mathbf{I}_k}{\gamma} + \mathbf{V}^T \mathbf{V} \right)^{-1} \mathbf{V} \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \frac{\mathbf{I}_k}{\gamma} + \mathbf{V}_S^T \mathbf{V}_S \right)^{-1} \mathbf{V}_S \mathbf{y}_S \right| \leq \frac{M}{2} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) - 1. \end{aligned}$$



Again, by Frechet inequalities, we can further breakdown the expression above to:

$$\begin{aligned}
&\geq \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S \right| \leq \frac{M}{2} \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left| \left( \frac{1}{N} \mathbf{y}^T \mathbf{V} - \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \right) \left( \frac{I_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} \frac{\mathbf{V} \mathbf{y}}{N} \right| \leq \frac{M}{6} \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left| \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \left( \frac{I_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} - \left( \frac{I_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \right) \frac{\mathbf{V} \mathbf{y}}{N} \right| \leq \frac{M}{6} \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left| \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \left( \frac{I_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \left( \frac{\mathbf{V} \mathbf{y}}{n} - \frac{\mathbf{V}_S \mathbf{y}_S}{n} \right) \right| \leq \frac{M}{6} \sqrt{\frac{\log(\frac{1}{\epsilon})}{n}} \right) - 3
\end{aligned}$$

By the definition of  $M$ , we can simplify the expressions above to:

$$\begin{aligned}
&\geq \mathbb{P} \left( \left| \frac{1}{N} \mathbf{y}^T \mathbf{y} - \frac{1}{n} \mathbf{y}_S^T \mathbf{y}_S \right| \leq \sqrt{\frac{A \log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left\| \frac{1}{N} \mathbf{y}^T \mathbf{V} - \frac{1}{n} \mathbf{y}_S^T \mathbf{V}_S \right\| \leq \sqrt{\frac{B \log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left\| \left( \frac{I_k}{N\gamma} + \frac{\mathbf{V}^T \mathbf{V}}{N} \right)^{-1} - \left( \frac{I_k}{n\gamma} + \frac{\mathbf{V}_S^T \mathbf{V}_S}{n} \right)^{-1} \right\| \leq \sqrt{\frac{H \log(\frac{1}{\epsilon})}{n}} \right) \\
&+ \mathbb{P} \left( \left\| \frac{\mathbf{V} \mathbf{y}}{n} - \frac{\mathbf{V}_S \mathbf{y}_S}{n} \right\| \leq \sqrt{\frac{B \log(\frac{1}{\epsilon})}{n}} \right) - 3.
\end{aligned}$$

Now we apply Lemma 1 and 2 to evaluate these expressions:

$$\begin{aligned}
&\geq 4 - 4\epsilon - 3 \\
&= 1 - 4\epsilon.
\end{aligned}$$

Thus, taking  $\epsilon' = 4\epsilon$ , and  $M' = \frac{M}{\log(2)}$  gives the result we want.

## B Proof of Theorem 1

### B.1 Proof of Finiteness

We first prove that the outer approximation algorithm indeed converges in finite number of steps  $C$  with probability 1. To do so, we note that the set of feasible solutions for the integer variables  $\mathcal{Z}$  is finite, and we prove that the series of feasible integer solutions  $z_1, z_2, \dots$  created during Algorithm 2 never repeats indefinitely, with probability 1.

First, we show that if we have  $z_a = z_b$  for some  $a + 1 \leq b$  where we have not terminated at iteration  $b$ , then we must have  $f(z_b, \theta_b; [S_b]) > f(z_a, \theta_a; [S_a])$ . At iteration  $a$ , the following cutting plane was added:

$$\eta \geq f(z_a, \theta_a; [S_a]) + \nabla f(z_a, \theta_a; [S_a])^T \begin{pmatrix} z - z_a \\ \theta - \theta_a \end{pmatrix}. \quad (\text{A14})$$

Therefore, using (A14) at iteration  $b - 1 \geq a$  (note  $z_b, \eta_b$  is the optimal solution in the  $b - 1$ th iteration, so it must satisfy the inequality in (A14)) gives:

$$\eta_b \geq f(z_a, \theta_a; [S_a]) + \nabla f(z_a, \theta_a; [S_a])^T \begin{pmatrix} z_b - z_a \\ \theta_b - \theta_a \end{pmatrix} \quad (\text{A15})$$

$$= f(z_a, \theta_a; [S_a]) + \nabla f(z_a, \theta_a; [S_a])^T (\theta_b - \theta_a) \quad (\text{A16})$$

$$\geq f(z_a, \theta_a; [S_a]). \quad (\text{A17})$$

The last inequality follows because  $\theta_a$  is the optimal solution to  $\text{NLP}(\theta_a; [S_a])$ , so no feasible direction can exist at  $\theta_a$ . Now consider the cutting plane problem at iteration  $b$ , where the following cutting plane was added:

$$\eta \geq f(z_b, \theta_b; [S_b]) + \nabla f(z_b, \theta_b; [S_b])^T \begin{pmatrix} z - z_b \\ \theta - \theta_b \end{pmatrix}. \quad (\text{A18})$$

Assume that we have  $f(z_b, \theta_b; [S_b]) \leq f(z_a, \theta_a; [S_a])$ . Then  $\eta = \eta_b$ ,  $z = z_b$  is a feasible solution for the cutting plane problem at iteration  $b$  since it satisfies the newly added constraint:  $\eta_b \geq f(z_a, \theta_a; [S_a]) \geq f(z_b, \theta_b; [S_b]) + \nabla f(z_b, \theta_b; [S_b])^T \begin{pmatrix} z_b - z_b \\ \theta_b - \theta_b \end{pmatrix} = f(z_b, \theta_b; [S_b])$ . By construction of the algorithm, we therefore satisfy the termination condition, so the algorithm terminates, a contradiction. Therefore, we must have  $f(z_b, \theta_b; [S_b]) > f(z_a, \theta_a; [S_a])$ .

Then consider Algorithm 2 after  $Q = |\mathcal{Z}|(q - 1) + 1$  iterations where  $q > 0$  is a positive integer, with the corresponding series of feasible integer solutions being  $z_1, \dots, z_Q$ . Then since there are only  $|\mathcal{Z}|$  feasible solutions, by the pigeonhole principle, there must be  $z \in \mathcal{Z}$  that has appeared at least  $q$  times in the series, and denote these as  $z_{l_1}, \dots, z_{l_q}$ . By the observation above, we know that if Algorithm 2 has not terminated, we must have:

$$f(z_{l_1}, \theta_{l_1}; [S_{l_1}]) < f(z_{l_2}, \theta_{l_2}; [S_{l_2}]) < \dots < f(z_{l_q}, \theta_{l_q}; [S_{l_q}])$$

Now note that  $z_{l_1} = z_{l_2} = \dots = z_{l_q}$ . Given  $z_t$ ,  $f(z_t, \theta_t; [S_t])$  is a random variable that only depends on  $S_t$  (by construction of the algorithm,  $\theta_t$  is the optimal solution to  $\text{NLP}(z_t, [S_t])$ ). Therefore, we have that  $f(z_{l_1}, \theta_{l_1}; [S_{l_1}]), \dots, f(z_{l_q}, \theta_{l_q}; [S_{l_q}])$  are  $q$  independent realizations (since  $S_{l_1}, S_{l_2}, \dots$  are independent) of an identically distributed random variable, and thus by standard probability results (see e.g. Feller (2008)) we have that:

$$\mathbb{P}(f(z_{l_1}, \theta_{l_1}; [S_{l_1}]) < f(z_{l_2}, \theta_{l_2}; [S_{l_2}]) < \dots < f(z_{l_q}, \theta_{l_q}; [S_{l_q}])) \leq \frac{1}{q!}$$

The result follows by noting that  $f(z_{l_1}, \theta_{l_1}; [S_{l_1}]), \dots, f(z_{l_q}, \theta_{l_q}; [S_{l_q}])$  are exchangeable random variables so any ordering is equally likely, and only at most 1 out of all  $q!$  permutations

of  $f(\mathbf{z}_{l_1}, \boldsymbol{\theta}_{l_1}; [S_{l_1}]), \dots, f(\mathbf{z}_{l_q}, \boldsymbol{\theta}_{l_q}; [S_{l_q}])$  satisfy the monotone relationship. Then we have that:

$$\begin{aligned} & \mathbb{P}(\text{Algorithm has not terminated after } Q \text{ iterations} \mid \mathbf{z}_1, \dots, \mathbf{z}_Q) \\ & \leq \mathbb{P}(f(\mathbf{z}_{l_1}, \boldsymbol{\theta}_{l_1}; [S_{l_1}]) < f(\mathbf{z}_{l_2}, \boldsymbol{\theta}_{l_2}; [S_{l_2}]) < \dots < f(\mathbf{z}_{l_q}, \boldsymbol{\theta}_{l_q}; [S_{l_q}])) \\ & \leq \frac{1}{q!} \end{aligned}$$

Therefore, by taking expectations on both sides, we have:

$$\mathbb{P}(\text{Algorithm has not terminated after } Q \text{ iterations}) \leq \frac{1}{q!},$$

which converges to 0 as  $Q \rightarrow \infty$ .

Now we can bound the number of iterations  $K$  when the algorithm terminates. Then by we have:

$$\begin{aligned} \mathbb{E}[K] &= \sum_{i=0}^{\infty} \mathbb{P}(K > i) \\ &\leq \sum_{q=1}^{\infty} |\mathcal{Z}| \mathbb{P}(K > |Z|(q-1) + 1) + 1 \\ &= \sum_{q=1}^{\infty} |\mathcal{Z}| \frac{1}{q!} + 1 \\ &= (e-1)|\mathcal{Z}| + 1 \end{aligned}$$

We have thus proved that the algorithm terminates in finite number of iterations with probability 1, and furthermore that its expected value is bounded by  $(e-1)|\mathcal{Z}| + 1$ .

## B.2 Proof of Feasibility and Optimality

Fletcher and Leyffer (1994) has shown that when the cutting plane algorithm terminates, it returns infeasibility or the optimal (feasible) solution. In particular, the cutting plane algorithm returns infeasibility if and only if the original problem is infeasible. Now note that we have assumed the problem is indeed feasible and the formulation of constraints in the stochastic cutting plane algorithm is unchanged from the standard cutting plane algorithm. Therefore, Algorithm 2 returns a feasible solution when it terminates.

Therefore, we focus on proving that the solution is  $\epsilon$ -optimal with high probability.

As in the theorem, let  $(\mathbf{z}^*, \boldsymbol{\theta}^*)$  be an optimal solution for the original problem and  $(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*)$  be the solution returned by Algorithm 2. Now let us consider the MILO at the iteration when the algorithm terminated at the  $K$ th iteration:

$$\min_{\mathbf{z} \in \mathcal{Z}, \boldsymbol{\theta} \in \Theta, \eta \geq lb} \eta \tag{A19}$$

$$\text{s.t. } \eta \geq f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [K], \tag{A20}$$

$$0 \geq g_j(\mathbf{z}_i, \boldsymbol{\theta}_i) + \nabla g_j(\mathbf{z}_i, \boldsymbol{\theta}_i)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [K], \quad \forall j \in [m]. \tag{A21}$$

For notational simplicity, we define  $\eta^*(\mathbf{z}, \boldsymbol{\theta})$  to be the optimal value for  $\eta$  under the solution  $(\mathbf{z}, \boldsymbol{\theta})$  for the problem above. In particular, if  $(\mathbf{z}, \boldsymbol{\theta})$  is feasible, then we have:

$$\eta^*(\mathbf{z}, \boldsymbol{\theta}) = \max_{i \in [K]} f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z} - \mathbf{z}_i \\ \boldsymbol{\theta} - \boldsymbol{\theta}_i \end{pmatrix}. \quad (\text{A22})$$

By construction of the algorithm, there must exist a cutting plane at the optimal point  $(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*)$ . Without loss of generality, we assume it is the final cutting plane, with the form:

$$\eta \geq f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n) + \nabla f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n)^T \begin{pmatrix} \mathbf{z} - \tilde{\mathbf{z}}^* \\ \boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}^* \end{pmatrix}. \quad (\text{A23})$$

Now we introduce a few lemmas:

**Lemma 5**  $\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) \geq f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n)$ .

*Proof* Proof: By feasibility of  $(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*)$ ,  $(\mathbf{z}, \boldsymbol{\theta}) = (\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*)$  into Equation (A23), the solution returned by Algorithm 2 must satisfy

$$\eta^*(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*) \geq f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n). \quad (\text{A24})$$

Now by optimality of  $(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*)$ , we must have:

$$\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) \geq \eta^*(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*). \quad (\text{A25})$$

Combining equation (A24) and (A25) gives the required statement.  $\square$

**Lemma 6** Under the randomization scheme in Algorithm 2, we have, for some absolute constant  $D$ :

$$\mathbb{P}(\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon) \geq 1 - 2K \exp\left(-\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 D}\right). \quad (\text{A26})$$

*Proof* Proof: Note that we can write  $\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*)$  as:

$$\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) = \max_{i \in [K]} f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix}.$$

Now by Frechet's inequalities:

$$\begin{aligned} & \mathbb{P}(\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon) \\ &= \mathbb{P}\left(\max_{i \in [K]} f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon\right) \\ &= \mathbb{P}\left(f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon \quad \forall i \in [K]\right) \\ &\geq \sum_{i=1}^K \mathbb{P}\left(f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon\right) - (K-1). \end{aligned} \quad (\text{A27})$$

We focus on each term:

$$\mathbb{P}\left(f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon\right)$$

By the convexity assumption on  $f$  (Assumption 1), we have that:

$$f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \geq f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N]) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix}, \quad \forall i \in [K]. \quad (\text{A28})$$

Therefore, by substituting Equation (A28), we have, for all  $i \in [K]$ :

$$\mathbb{P} \left( f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) + \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon \right) \quad (\text{A29})$$

$$\begin{aligned} &= \mathbb{P} \left( (f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) - f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])) + (\nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T - \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])^T) \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} \leq \epsilon \right) \\ &\geq \mathbb{P} \left( (f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) - f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])) \leq \frac{\epsilon}{1 + \sqrt{p_1 + p_2}} \right. \\ &\quad \left. \cap (\nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T - \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])^T) \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} \leq \frac{(\sqrt{p_1 + p_2})\epsilon}{1 + \sqrt{p_1 + p_2}} \right). \end{aligned} \quad (\text{A30})$$

Using the Frechet inequalities, we can further simplify the expression above:

$$\begin{aligned} (\text{A30}) &\geq \mathbb{P} \left( (f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n) - f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])) \leq \frac{\epsilon}{1 + \sqrt{p_1 + p_2}} \right) \\ &\quad + \mathbb{P} \left( (\nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; S_i^n)^T - \nabla f(\mathbf{z}_i, \boldsymbol{\theta}_i; [N])^T) \begin{pmatrix} \mathbf{z}^* - \mathbf{z}_i \\ \boldsymbol{\theta}^* - \boldsymbol{\theta}_i \end{pmatrix} \leq \frac{(\sqrt{p_1 + p_2})\epsilon}{1 + \sqrt{p_1 + p_2}} \right) - 1. \end{aligned} \quad (\text{A31})$$

Now by Assumption 3, these expressions lead to :

$$\geq 1 - \exp \left( -\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 M^2} \right) + 1 - \exp \left( -\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 M'^2} \right) - 1 \quad (\text{A32})$$

$$= 1 - 2 \exp \left( -\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 D} \right), \quad (\text{A33})$$

where  $D$  is an absolute constant. Now we substitute Equation (A33) into Equation (A27) to get that:

$$\begin{aligned} &\mathbb{P}(\eta^*(\mathbf{z}^*, \boldsymbol{\theta}^*) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon) \\ &\geq K - 2K \exp \left( -\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 D} \right) - (K - 1) \\ &= 1 - 2K \exp \left( -\frac{n\epsilon^2}{(1 + \sqrt{p_1 + p_2})^2 D} \right). \end{aligned}$$

as needed.  $\square$

We now combine the lemmas to derive the final inequality:

$$\begin{aligned} &\mathbb{P} \left( f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; [N]) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; [N]) \leq \epsilon \right) \\ &\geq \mathbb{P} \left( f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; [N]) - f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n) \leq \frac{\epsilon}{2 + \sqrt{p_1 + p_2}} \cap f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; N) \leq \frac{(1 + \sqrt{p_1 + p_2})\epsilon}{2 + \sqrt{p_1 + p_2}} \right). \end{aligned}$$

By the Frechet inequalities, we can simplify the expression above to:

$$\begin{aligned} &\geq \mathbb{P} \left( f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; [N]) - f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n) \leq \frac{\epsilon}{2 + \sqrt{p_1 + p_2}} \right) \\ &+ \mathbb{P} \left( f(\tilde{\mathbf{z}}^*, \tilde{\boldsymbol{\theta}}^*; S_K^n) - f(\mathbf{z}^*, \boldsymbol{\theta}^*; N) \leq \frac{(1 + \sqrt{p_1 + p_2})\epsilon}{2 + \sqrt{p_1 + p_2}} \right) - 1. \end{aligned}$$

Using Lemma 5, we can bound  $f(\tilde{z}^*, \tilde{\theta}^*; S_K^n)$ :

$$\begin{aligned} &\geq \mathbb{P} \left( f(\tilde{z}^*, \tilde{\theta}^*; [N]) - f(\tilde{z}^*, \tilde{\theta}^*; S_K^n) \leq \frac{\epsilon}{2 + \sqrt{p_1 + p_2}} \right) \\ &+ \mathbb{P} \left( \eta^*(z^*, \theta^*) - f(z^*, \theta^*; N) \leq \frac{(1 + \sqrt{p_1 + p_2})\epsilon}{2 + \sqrt{p_1 + p_2}} \right) - 1. \end{aligned}$$

Using Assumption 3, we can evaluate the first expression:

$$\geq 1 - \exp \left( \frac{-n\epsilon^2}{(2 + \sqrt{p_1 + p_2})^2 M^2} \right) + \mathbb{P} \left( \eta^*(z^*, \theta^*) - f(z^*, \theta^*; N) \leq \frac{(1 + \sqrt{p_1 + p_2})\epsilon}{2 + \sqrt{p_1 + p_2}} \right) - 1.$$

Substituting Lemma 6 gives us the formula for the second expression:

$$\begin{aligned} &\geq 1 - \exp \left( \frac{-n\epsilon^2}{(2 + \sqrt{p_1 + p_2})^2 M^2} \right) + 1 - 2K \exp \left( \frac{-n\epsilon^2}{(2 + \sqrt{p_1 + p_2})^2 D} \right) - 1 \\ &= 1 - (2K + 1) \exp \left( \frac{-n\epsilon^2}{(2 + \sqrt{p_1 + p_2})^2 J} \right), \end{aligned}$$

where  $J$  is an absolute constant, as required.  $\square$