

BIG DATA, LEVERAGE SCORES, AND MINIMUM VOLUME COVERING ELLIPSOIDS: BRIDGING THEORY WITH PRACTICE

Anonymous authors

Paper under double-blind review

ABSTRACT

The Minimum Volume Covering Ellipsoid (MVCE) problem, characterized by n observations in d dimensions where $n \gg d$, can be computationally very expensive in the big data regime. We apply methods from randomized numerical linear algebra to develop a data-driven leverage score sampling algorithm for solving MVCE, and establish theoretical error bounds and a convergence guarantee. Assuming the leverage scores follow a power law decay, we show that the computational complexity of computing the approximation for MVCE is reduced from $\mathcal{O}(nd^2)$ to $\mathcal{O}(nd \log n + \text{poly}(d))$, which is a significant improvement in big data problems. Numerical experiments on large-scale synthetic data, as well as real-world data, demonstrate the efficacy of our new algorithm, showing that it substantially reduces computation time while yielding near-optimal solutions.

1 INTRODUCTION

The Minimum Volume Covering Ellipsoid (MVCE) problem arises in many applied and theoretical areas. Statistical applications include outlier detection (Titterton, 1978), clustering (Rosen, 1965), and the closely related D-optimal design problem (Silvey, 1980). In fact, the MVCE problem and the D-optimal design problem are dual to one another (Sibson, 1972; Titterton, 1975). Containing ellipsoids are used in parameter identification and control theory to describe uncertainty sets for parameters and state vectors (Schweppe, 1968; Chernousko, 2005). Minimum volume covering ellipsoids are also used in computational geometry and computer graphics (Eberly, 2006), in particular, for collision detection (Chen et al., 2016).

Many algorithms for computing MVCEs and D-optimal designs have been studied in the literature. These include Frank-Wolfe type algorithms (Frank et al., 1956; Wolfe, 1970; Atwood, 1973; Khachiyan, 1996; Kumar & Yildirim, 2005), interior point algorithms (Khachiyan & Todd, 1993; Nesterov & Nemirovskii, 1994), the Dual Reduced Newton algorithm (Sun & Freund, 2004), the Cocktail algorithm (Yu, 2011), the Randomized Exchange algorithm (Harman et al., 2020), and the Fixed Point algorithm (Cohen et al., 2019; Song et al., 2022; Woodruff & Yasuda, 2023; 2025). However, when these algorithms are applied to very large datasets, they may be computationally inefficient, and may exceed storage limitations (Harman et al., 2020). One solution is to combine such solution algorithms with active set or batching strategies (Sun & Freund, 2004; Källberg & Andrén, 2019; Kudela, 2019; Rosa & Harman, 2022). The main idea of this approach is to iteratively apply the solution algorithm to a smaller subset of points until convergence to the solution. The motivation for this approach is that an ellipsoid is only determined by at most $d(d+3)/2$ points on its boundary (John, 2014).

Instead of applying an active set strategy, we apply deterministic sampling to reduce the number of points considered by the algorithm. This deterministic sampling method selects points corresponding to the highest statistical leverage scores. The resulting compressed dataset approximately maintains many of the qualities of the original dataset, provided that the number of samples is sufficiently large (Papailiopoulos et al., 2014; McCurdy et al., 2019). However, there is no immediate guarantee of the quality of solution to the MVCE problem for this compressed dataset.

In this paper, we present a new simplified proof of the subspace embedding theorem from McCurdy et al. (2019). Using this theorem, we provide the first theoretical guarantees on the quality of initial

054 and final solution for the MVCE problem when using deterministic leverage score sampling. Further,
 055 we show that these guarantees still hold (with high probability) when approximate leverage scores
 056 are used. Assuming a power law decay on the leverage scores, we show that our method improves the
 057 theoretical computation time required to approximate the MVCE. We also demonstrate the efficiency
 058 of our approach on synthetic and real world datasets.

059 **Notation.** Throughout this paper, vectors and matrices are denoted by bold lowercase and bold
 060 uppercase letters, respectively (e.g. \mathbf{a} and \mathbf{A}). The i th entry of \mathbf{a} is denoted a_i , and the (i, j) th
 061 entry of \mathbf{A} is denoted $A_{i,j}$. Let \mathbf{A} and \mathbf{B} be symmetric positive definite matrices, then $\mathbf{A} \succeq \mathbf{B}$ if
 062 $\mathbf{A} - \mathbf{B}$ is symmetric positive semidefinite; and $\mathbf{A} \succ \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is symmetric positive definite. The
 063 determinant of a matrix \mathbf{A} is denoted $\det(\mathbf{A})$. Unless otherwise specified, we use $\mathbf{U} = \text{diag}(\mathbf{u})$,
 064 where \mathbf{U} is a matrix with the vector \mathbf{u} in the main diagonal. We use regular lowercase to denote
 065 scalar constants (e.g. c). Finally, \mathbf{e} denotes a vector of ones, and \mathbf{e}_i denotes the vector with one at
 066 position i , and zero otherwise.

069 2 MINIMUM VOLUME COVERING ELLIPSOID PROBLEM

070
 071 Let $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be a set of data points in \mathbb{R}^d . Then the minimum volume covering ellipsoid
 072 is the ellipsoid that covers \mathcal{X} , which attains the minimum volume of all covering ellipsoids of \mathcal{X} .
 073 We assume throughout that there exists a non-degenerate minimum volume covering ellipsoid.

074 We define the ellipsoid $\mathcal{E}(\mathbf{Q}, \mathbf{x}_c)$ as

$$075 \mathcal{E}(\mathbf{Q}, \mathbf{x}_c) := \{\mathbf{x} \in \mathbb{R}^d : (\mathbf{x} - \mathbf{x}_c)^\top \mathbf{Q} (\mathbf{x} - \mathbf{x}_c) \leq d\},$$

076
 077 where \mathbf{x}_c is the center of the ellipsoid, and \mathbf{Q} is a d -dimensional symmetric positive definite matrix.
 078 Its volume is given by

$$079 \text{vol}(\mathcal{E}(\mathbf{Q}, \mathbf{x}_c)) = d^{d/2} \Omega_d \det(\mathbf{Q})^{-1/2},$$

080
 081 where Ω_d is the volume of the unit ball in \mathbb{R}^d (e.g., see Todd (2016a)).

082
 083 We can now write the mathematical formulation of the Minimum Volume Covering Ellipsoid prob-
 084 lem. Suppose we have a finite set of points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$. Then its minimum volume
 085 covering ellipsoid can be found by solving

$$086 \begin{aligned} & \underset{\mathbf{Q} \succ \mathbf{0}, \mathbf{x}_c \in \mathbb{R}^d}{\text{minimize}} && -\log \det(\mathbf{Q}) && (\text{P}_0) \\ & \text{subject to} && (\mathbf{x}_i - \mathbf{x}_c)^\top \mathbf{Q} (\mathbf{x}_i - \mathbf{x}_c) \leq d, && i = 1, \dots, n. \end{aligned}$$

087
 088
 089 Although the objective function is convex, (P₀) itself is not convex (e.g., see Todd (2016a)).

090
 091 We therefore reformulate (P₀) so that it is convex. At the cost of working in \mathbb{R}^{d+1} , we can calculate
 092 the centered minimum volume covering ellipsoid, and recover the solution to (P₀) (Titterton,
 093 1975). Therefore, we can set $\mathbf{x}_c = \mathbf{0}$, and obtain

$$094 \begin{aligned} & \underset{\mathbf{Q} \succ \mathbf{0}}{\text{minimize}} && f(\mathbf{Q}) := -\log \det(\mathbf{Q}) && (\text{P}) \\ & \text{subject to} && \mathbf{x}_i^\top \mathbf{Q} \mathbf{x}_i \leq d, && i = 1, \dots, n. \end{aligned}$$

095
 096
 097 We let \mathbf{Q}^* denote an optimal solution to (P). Then $\text{MVCE}(\mathcal{X}) := \mathcal{E}(\mathbf{Q}^*, \mathbf{0})$ is the minimum volume
 098 covering ellipsoid. We will refer to (P) as the MVCE problem.

099
 100 The dual problem to (P) is the D-optimal design problem. This problem is concave, and can be
 101 formulated as

$$102 \begin{aligned} & \underset{\mathbf{u} \in \mathbb{R}^n}{\text{maximize}} && g(\mathbf{u}) := \log \det \left(\sum_{i=1}^n u_i \mathbf{x}_i \mathbf{x}_i^\top \right) && (\text{D}) \\ & \text{subject to} && \sum_{i=1}^n u_i = 1, && \mathbf{u} \geq \mathbf{0}, \end{aligned}$$

where \mathbf{u} is called the design vector. We note that for every design vector \mathbf{u} , we can find its associated shape matrix

$$\mathbf{Q}(\mathbf{u}) := \left(\sum_{i=1}^n u_i \mathbf{x}_i \mathbf{x}_i^\top \right)^{-1},$$

provided that the inverse exists. Hence if we have an optimal solution \mathbf{u}^* to (D), then the optimal solution to (P) is $\mathbf{Q}(\mathbf{u}^*)$.

We note that (D) (and (P)) cannot usually be solved exactly, so we will focus on deriving approximate solutions. To ensure the chosen algorithm terminates with a guaranteed quality of solution, we will define some approximate optimality conditions. A feasible \mathbf{u} for (D) is called δ -primal feasible if $\mathbf{Q}(\mathbf{u})$ satisfies

$$\mathbf{x}_i^\top \mathbf{Q}(\mathbf{u}) \mathbf{x}_i \leq (1 + \delta) d,$$

for all $i = 1, \dots, n$. If $\mathbf{Q}(\mathbf{u})$ additionally satisfies

$$\mathbf{x}_i^\top \mathbf{Q}(\mathbf{u}) \mathbf{x}_i \geq (1 - \delta) d \text{ if } u_i > 0,$$

for all $i = 1, \dots, n$, then we say that \mathbf{u} is δ -approximately optimal. These optimality conditions ensure the optimality gap is small.

Proposition 1 (Todd (2016a), Proposition 2.9). *If we have a δ -primal feasible (or δ -approximately optimal) solution \mathbf{u} , then \mathbf{u} and $(1 + \delta)^{-1} \mathbf{Q}(\mathbf{u})$ are both within $d \log(1 + \delta)$ of being optimal in (D) and (P), respectively.*

3 SAMPLING USING LEVERAGE SCORES

The concept of statistical leverage scores has long been used in statistical regression diagnostics to identify outliers (Rousseeuw & Hubert, 2011). Given a data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ with $n > d$, consider any orthogonal matrix \mathbf{A} such that $\text{Range}(\mathbf{A}) = \text{Range}(\mathbf{X})$. The i th leverage score corresponding to the i th row of \mathbf{X} is defined as

$$\ell_i(\mathbf{X}) := \|\mathbf{A}(i, :)\|_2^2.$$

It can be easily shown that this is well defined in that the leverage score does not depend on the particular choice of the basis matrix \mathbf{A} . Furthermore, the i th leverage score is the i th diagonal entry of the Hat matrix, that is,

$$\ell_i(\mathbf{X}) = \mathbf{e}_i^\top \mathbf{H} \mathbf{e}_i, \quad i = 1, \dots, n,$$

where

$$\mathbf{H} := \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top.$$

The Hat matrix is symmetric and idempotent. We can use these properties to easily show that

$$0 \leq \ell_i(\mathbf{X}) \leq 1,$$

for all i , and

$$\sum_{i=1}^n \ell_i(\mathbf{X}) = \text{rank}(\mathbf{X}).$$

Since we only consider full rank matrices, the sum of the leverage scores is d .

3.1 DETERMINISTIC SAMPLING

When sampling deterministically, we sample the s rows from \mathbf{X} with highest leverage scores. Without loss of generality, assume that the leverage scores are ordered $\ell_1(\mathbf{X}) \geq \dots \geq \ell_n(\mathbf{X})$. We summarize the sampling procedure in Algorithm 1.

Using this algorithm, we obtain the following subspace embedding result.

Algorithm 1 Deterministic Leverage Score Sampling with Threshold (Papailiopoulos et al., 2014)**Require:** $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}, \varepsilon \in (0, 1)$ 1: Calculate leverage scores for each row in \mathbf{X} . Assume $\ell_1(\mathbf{X}) \geq \dots \geq \ell_n(\mathbf{X})$.2: Let $s = \arg \min_j \left(\sum_{i=1}^j \ell_i(\mathbf{X}) > d - \varepsilon \right)$.3: Let $\mathbf{R} = \mathbf{0} \in \mathbb{R}^{s \times n}$.4: **for** $i = 1 : s$ **do**5: Set row i of \mathbf{R} equal to \mathbf{e}_i .6: **end for****Ensure:** \mathbf{R}, s **Theorem 2** (McCurdy et al. (2019), Theorem 1). *Let $\varepsilon \in (0, 1)$. Use Algorithm 1 to construct the sampling matrix \mathbf{R} . Then*

$$(1 - \varepsilon)\mathbf{X}^\top \mathbf{X} \prec (\mathbf{R}\mathbf{X})^\top \mathbf{R}\mathbf{X} \preceq \mathbf{X}^\top \mathbf{X}.$$

We now present our new simplified proof of the lower bound.

Proof. We use the fact that for any i ,

$$\mathbf{x}_i \mathbf{x}_i^\top \preceq \ell_i(\mathbf{X}) \mathbf{X}^\top \mathbf{X},$$

see, for example, the proof of Lemma 4 in Cohen et al. (2015). Hence

$$\mathbf{X}^\top \mathbf{X} - (\mathbf{R}\mathbf{X})^\top \mathbf{R}\mathbf{X} = \sum_{i=s+1}^n \mathbf{x}_i \mathbf{x}_i^\top \preceq \sum_{i=s+1}^n \ell_i(\mathbf{X}) \mathbf{X}^\top \mathbf{X} = \left(\sum_{i=s+1}^n \ell_i(\mathbf{X}) \right) \mathbf{X}^\top \mathbf{X} \prec \varepsilon \mathbf{X}^\top \mathbf{X},$$

since $\sum_{i=s+1}^n \ell_i(\mathbf{X}) < \varepsilon$, by construction of Algorithm 1. Rearranging, we obtain our lower bound. \square

Algorithm 1 has time complexity $\mathcal{O}(nd^2)$, due to the cost of calculating the leverage scores exactly. This can be improved by using approximate leverage score algorithms. For example, the algorithm of Drineas et al. (2012) uses a fast sampled randomized Hadamard transform (SRHT), and has time complexity $\mathcal{O}(nd \log n)$.¹ Moreover, we can use these approximate leverage score algorithms to prove an analogous result to Theorem 2 (see Appendix A).

We note that for many datasets, the number of points selected after applying Algorithm 1 may still be large. In this case, we can apply Algorithm 1 multiple times, to obtain a similar subspace embedding result (see Appendix B).

4 OUR APPROACH

We are interested in calculating $\text{MVCE}(\mathcal{X})$, where $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$, $n \gg d$. For very large n (and d), this can be very computationally expensive. Instead of calculating $\text{MVCE}(\mathcal{X})$ directly, we will calculate the MVCE on a much smaller subset of \mathcal{X} . We select our subset by using leverage score based sampling, as introduced in Section 3. Let this subset be \mathcal{X}_s , where \mathcal{X}_s contains up to s points sampled from \mathcal{X} . For ease of calculation, let the points from \mathcal{X} be stored in the rows of the data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$. Then we sample s rows from \mathbf{X} , which we store in the rows of $\mathbf{X}_s \in \mathbb{R}^{s \times d}$. We then use \mathbf{X}_s instead of \mathbf{X} in our solution algorithm. That is, our solution algorithm solves the concave problem

$$\begin{aligned} & \underset{\mathbf{u} \in \mathbb{R}^s}{\text{maximize}} && g_s(\mathbf{u}) := \log \det(\mathbf{X}_s^\top \mathbf{U} \mathbf{X}_s) && (\text{D}_s) \\ & \text{subject to} && \sum_{i=1}^s u_i = 1, \quad \mathbf{u} \geq \mathbf{0}. \end{aligned}$$

¹In a recent preprint, Eshragh et al. (2023) developed the Sequential Approximate Leverage-Score Algorithm (SALSA) with time complexity $\mathcal{O}(nd)$.

216 4.1 COMPUTATIONAL COMPLEXITY

217
218 We compare the computational complexity of our algorithm to the current state of the art algo-
219 rithm, the Wolfe-Atwood (WA) algorithm (Wolfe, 1970; Atwood, 1973), which computes a δ -
220 approximately optimal solution to (D). It uses the Kumar-Yildirim initialization (Kumar & Yildirim,
221 2005), which places equal weight on a small subset of points. These algorithms have time complex-
222 ity $\mathcal{O}(nd^2(\log \log d + \delta^{-1}))$ (Todd & Yildirim, 2007) and $\mathcal{O}(nd^2)$ (Kumar & Yildirim, 2005),
223 respectively. Together, this costs $\mathcal{O}(nd^2(\log \log d + \delta^{-1}))$.

224 The computational complexity of our approach is as follows. We calculate approximate leverage
225 scores in $\mathcal{O}(nd \log n)$ time, which, with high probability, satisfy $\hat{\ell}_i(\mathbf{X}) = (1 \pm \beta)\ell_i(\mathbf{X})$, for some
226 $\beta \in (0, 1/2]$ (Drineas et al., 2012). Then Algorithm 1 also runs in $\mathcal{O}(nd \log n)$ time. We then use
227 the WA algorithm with Kumar-Yildirim initialization (Kumar & Yildirim, 2005) to find an optimal
228 solution to (D_s) . Since $\mathbf{X}_s \in \mathbb{R}^{s \times d}$, these algorithms have time complexity $\mathcal{O}(sd^2 \log \log d)$ and
229 $\mathcal{O}(sd^2)$, respectively. Thus, the total time complexity is $\mathcal{O}(nd \log n + sd^2(\log \log d + \delta^{-1}))$.
230 Further, if the leverage scores exhibit a power law decay, then $s = \text{poly}\left(d(1 + \beta), \frac{1}{(1-\beta)\epsilon}\right)$ (see
231 Appendix A).
232

233 4.2 COMPARISON WITH THE WORK OF COHEN ET. AL.

234 Recently, Cohen et al. (2019) developed the Fixed Point algorithm, which computes a δ -primal
235 feasible solution to (D). (Woodruff & Yasuda, 2023) extend this result; by sampling rows of \mathbf{X}
236 with probabilities proportionate to the weights calculated by the Fixed Point algorithm, they obtain
237 a δ -approximately optimal solution to (D).
238

239 The computational complexity of the Fixed Point algorithm is as follows. We examine Algorithm
240 2 from Cohen et al. (2019), which uses sketching techniques from randomized numerical linear
241 algebra to speed up each iteration. Theorem C.7 from Cohen et al. (2019) states that Algorithm
242 2 takes at most $\mathcal{O}(\delta^{-1} \log(\frac{n}{d}))$ iterations to complete. In each iteration there are $\mathcal{O}(\delta^{-1})$ linear
243 systems of the form $\mathbf{A}^T \mathbf{W} \mathbf{A} \mathbf{x} = \mathbf{b}$ to be solved. Assuming \mathbf{A} is dense, solving each of these linear
244 systems costs $\mathcal{O}(nd)$. This gives a total time complexity of $\mathcal{O}(\delta^{-2} nd \log(\frac{n}{d}))$ for dense matrices.
245

246 In big data regimes with $n \gg d$, it is reasonable to assume that $\delta < \frac{1}{d}$ is desirable. This assumption
247 holds for many problem instances considered in the MVCE literature, since these problems typically
248 have dimension $d \leq 200$, and are solved until tolerance $\delta = 10^{-6}$ (or stricter) is achieved (see, e.g.
249 Sun & Freund (2004); Damla Ahipasaoglu et al. (2008); Yu (2011); Källberg & Andr n (2019);
250 Kudela (2019)). With this assumption, Algorithm 2 has total time complexity $\mathcal{O}(nd^3 \log(nd))$ and
251 our algorithm has total time complexity $\mathcal{O}(nd \log n + sd^3)$, where $s \leq n$. Thus, in the context of
252 tall data matrices with dense input, our algorithm outperforms the Fixed Point algorithm theoret-
253 ically. (This also holds numerically, see Appendix E.)
254

255 5 INITIAL OPTIMALITY GAP

256 We would like to know how well $\text{MVCE}(\mathcal{X}_s)$ approximates $\text{MVCE}(\mathcal{X})$. We will first provide an
257 upper bound for the initial optimality gap, for a particular choice of initial \mathbf{u} . This \mathbf{u} must be feasible
258 for (D_s) , that is, we must have $\sum_{i=1}^s u_i = 1$, and $\mathbf{u} \geq \mathbf{0}$. We choose
259

$$260 \mathbf{u}_0 = \left[\frac{1}{s}, \dots, \frac{1}{s}\right]^T \in \mathbb{R}^s.$$

261 Then the initial optimality gap is given by
262

$$263 g^* - g_s(\mathbf{u}_0),$$

264 where g^* is the optimal objective value when the full set \mathcal{X} is considered.
265

266 To derive our bound, we will compare our initial solution with an initialization due to Khachiyan
267 (1996). Khachiyan’s initialization \mathbf{u}_K puts equal weight on all n points of \mathcal{X} , and guarantees that
268

$$269 g^* - g(\mathbf{u}_K) \leq d \log n,$$

as shown by Khachiyan (1996). Therefore, our bound can be found by exploiting the fact that

$$g^* - g_s(\mathbf{u}_0) = (g^* - g(\mathbf{u}_K)) + (g(\mathbf{u}_K) - g_s(\mathbf{u}_0)).$$

Upon simplification,

$$g(\mathbf{u}_K) = \log \det (\mathbf{X}^\top \mathbf{X}) - d \log n, \quad (1)$$

and, similarly,

$$g_s(\mathbf{u}_0) = \log \det (\mathbf{X}_s^\top \mathbf{X}_s) - d \log s. \quad (2)$$

Hence

$$g^* - g_s(\mathbf{u}_0) \leq \log \det (\mathbf{X}^\top \mathbf{X}) - \log \det (\mathbf{X}_s^\top \mathbf{X}_s) + d \log s. \quad (3)$$

This suggests the need to compare $\log \det (\mathbf{X}_s^\top \mathbf{X}_s)$ with $\log \det (\mathbf{X}^\top \mathbf{X})$.

Suppose we construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$ as in Algorithm 1. Then Theorem 2 guarantees that

$$(1 - \varepsilon)\mathbf{X}^\top \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{X}_s.$$

Now, for positive semidefinite matrices \mathbf{A}, \mathbf{B} , with $\mathbf{A} \preceq \mathbf{B}$, we have

$$\det \mathbf{A} \leq \det \mathbf{B}. \quad (4)$$

Since the logarithm is monotonic, it then follows that $\log \det \mathbf{A} \leq \log \det \mathbf{B}$. Hence

$$\log \det (\mathbf{X}^\top \mathbf{X}) < \log \det (\mathbf{X}_s^\top \mathbf{X}_s) - d \log(1 - \varepsilon). \quad (5)$$

Combining Equations (3) and (5), we obtain the bound presented in Theorem 3.

Theorem 3. *Select $\mathbf{X}_s = \mathbf{R}\mathbf{X}$ using Algorithm 1, and let our initial solution be given by $\mathbf{u}_0 = \frac{1}{s}\mathbf{e}$. Then*

$$g^* - g_s(\mathbf{u}_0) < d \log \left(\frac{s}{1 - \varepsilon} \right).$$

Proof. See Appendix C. □

6 FINAL OPTIMALITY GAP

We now provide an upper bound for the final optimality gap. The final optimality gap is given by

$$g^* - g_s^*,$$

where g_s^* is the optimal objective value when only \mathcal{X}_s is considered. Note that for any \mathbf{u} feasible for D_s , we have

$$g^* - g_s^* \leq g^* - g_s(\mathbf{u}),$$

since g_s is concave.

Consider the feasible solution $\tilde{\mathbf{u}}_s$ for D_s , given by

$$\tilde{\mathbf{u}}_s = \frac{1}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{u}_s,$$

where \mathbf{u}_s contains the first s entries of \mathbf{u}^* . We would like a bound similar to the one in Theorem 2, with \mathbf{X} replaced with a rescaled version \mathbf{Y} . More precisely, let

$$\mathbf{Y} := \sqrt{U^*} \mathbf{X}.$$

To apply Theorem 2, we require

$$\sum_{i=s+1}^n \ell_i(\mathbf{Y}) < \varepsilon, \quad (6)$$

for some $\varepsilon \in (0, 1)$. To show this, we use the following result.

Proposition 4. Use Algorithm 1 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$, with $\varepsilon \in (0, 1)$. Suppose that an optimal solution \mathbf{u}^* of (D) satisfies $u_i^* > 0$ for all $i = 1, \dots, s$. Define $\mathbf{Y} := \sqrt{\mathbf{U}^*}\mathbf{X}$. Then

$$\sum_{i=s+1}^n \ell_i(\mathbf{Y}) \leq \sum_{i=s+1}^n \ell_i(\mathbf{X}).$$

Otherwise, let \mathbf{u} be a δ -feasible solution to (D). Define $\mathbf{Y} := \sqrt{\mathbf{U}}\mathbf{X}$. Then

$$\sum_{i=s+1}^n \ell_i(\mathbf{Y}) \leq \sum_{i=s+1}^n \ell_i(\mathbf{X}).$$

Proof. See Appendix D. □

Then (6) necessarily holds, by construction of Algorithm 1. Therefore, we may apply Theorem 2 with \mathbf{Y} instead of \mathbf{X} , to obtain the bound

$$(1 - \varepsilon)\mathbf{X}^\top \mathbf{U}^* \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{U}_s \mathbf{X}_s.$$

That is, the feasible solution $\tilde{\mathbf{u}}_s$ satisfies

$$(1 - \varepsilon)\mathbf{X}^\top \mathbf{U}^* \mathbf{X} \preceq \frac{1 - \varepsilon}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{X}^\top \mathbf{U}^* \mathbf{X} \prec \frac{1}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{X}_s^\top \mathbf{U}_s \mathbf{X}_s = \mathbf{X}_s^\top \tilde{\mathbf{U}}_s \mathbf{X}_s.$$

Hence

$$\begin{aligned} g^* - g_s^* &\leq g^* - g_s(\tilde{\mathbf{u}}_s) \\ &= \log \det(\mathbf{X}^\top \mathbf{U}^* \mathbf{X}) - \log \det(\mathbf{X}_s^\top \tilde{\mathbf{U}}_s \mathbf{X}_s) \\ &< \log \det(\mathbf{X}^\top \mathbf{U}^* \mathbf{X}) - (d \log(1 - \varepsilon) + \log \det(\mathbf{X}^\top \mathbf{U}^* \mathbf{X})) \\ &= d \log\left(\frac{1}{1 - \varepsilon}\right). \end{aligned}$$

Theorem 5. Use Algorithm 1 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$, with $\varepsilon \in (0, 1)$. If there exists an optimal solution \mathbf{u}^* of (D) satisfying $u_i^* > 0$ for all $i = 1, \dots, s$, then

$$g^* - g_s^* < d \log\left(\frac{1}{1 - \varepsilon}\right).$$

Otherwise,

$$g^* - g_s^* < d \log\left(\frac{1 + \delta}{1 - \varepsilon}\right),$$

where the parameter $\delta > 0$ can be chosen to be arbitrarily small.

Proof. See Appendix D. □

7 NUMERICAL RESULTS

We generate three large datasets of size $n = 10^7$, $d = 100$. Without loss of generality, we assume the leverage scores are sorted in descending order, that is, $\ell_1(\mathbf{X}) \geq \dots \geq \ell_n(\mathbf{X})$. The first dataset is Rotated Cauchy (Todd, 2016a). The points are generated so that they have rotational symmetry, and the distances of the points from the origin are Cauchy. The leverage scores of these points quickly decay. The second dataset is Lognormal, which has a shallower leverage score decay. The third dataset is Gaussian, which has leverage scores that are close to uniform. We provide additional numerical results in Appendix E.

All computations are performed on a personal laptop with a 64 bit MacOS 13 operating system, and a 2.4 GHz Quad-Core Intel Core i5 processor with 8 GB of RAM. The algorithms are run using MATLAB (R2021a).

We use the WA algorithm (Wolfe, 1970; Atwood, 1973) to calculate the MVCEs of the two datasets, with $\delta = 10^{-9}$. We initialize using the Kumar-Yildirim initialization (Kumar & Yildirim, 2005). We use Todd’s Matlab implementation of these algorithms (Todd, 2016b). Then, for s varying from 0.1 to 10% of n , we sample from each dataset in three ways: deterministic leverage score sampling, uniform sampling, and randomized leverage score sampling (sampling with probabilities proportionate to the leverage scores (Drineas et al., 2008)). In these results, we use the exact leverage scores.

Let \tilde{g}^* and \tilde{g}_s^* be the optimal values obtained when using the WA algorithm on the full and sampled datasets respectively. In Figure 1a, we summarize the calculated optimality gaps $\tilde{g}^* - \tilde{g}_s^*$ for the Rotated Cauchy dataset. The deterministic and randomized leverage score sampling performed similarly, with near zero optimality gap for all values of s . Uniform sampling performed poorly, its optimality gap decreasing with increasing s . In Figure 1b we summarize the total computation time for calculating the MVCEs. For the sampled datasets, this also includes the computation time for calculating leverage scores (if applicable), and sampling from the dataset. The total computation time for the full dataset was 2 450 seconds, that is, just over 40 minutes. Uniform sampling was faster than the leverage score sampling methods (due to the calculation of the leverage scores), but had very large optimality gaps (see Figure 1a).

In Figure 2a, we summarize the calculated optimality gaps $\tilde{g}^* - \tilde{g}_s^*$ for the Lognormal dataset. The deterministic and randomized leverage score sampling performed similarly, with zero optimality gap for all values of s . Uniform sampling performed poorly, its optimality gap decreasing with increasing s . In Figure 2b we summarize the total computation time for calculating the MVCEs. The total computation time for the full dataset was 2564 seconds, that is, just over 42 minutes. Uniform sampling was generally faster than the leverage score sampling methods (due to the calculation of the leverage scores), but had very large optimality gaps (see Figure 2a).

In Figure 3a, we summarize the calculated optimality gaps $\tilde{g}^* - \tilde{g}_s^*$ for the Gaussian dataset. Only the deterministic leverage score sampling performed well, with near zero optimality gap for all s . The uniform and randomized leverage score sampling performed similarly, with optimality gaps decreasing with increasing s . This is unsurprising, since the leverage scores for a Gaussian dataset are close to uniform. In Figure 3b we summarize the total computation time for calculating the MVCEs. The total computation time for the full dataset was very large, at 141 840 seconds, that is, over 39 hours. As with the other dataset, uniform sampling was the fastest, but at the cost of a larger optimality gap (see Figure 3a). Additionally, both leverage score sampling methods had similar runtimes, but only the deterministic sampling had near zero optimality gap (see Figure 3a).

Overall, the deterministic leverage score sampling performs the best, achieving both a small optimality gap and greatly decreasing computation time on all three datasets.

8 CONCLUSION

In this paper, we have provided the first theoretical guarantees on the quality of initial and final solutions for the MVCE problem, when sampling points deterministically according to their statistical leverage scores. We proved this approach is efficient, assuming the leverage scores exhibit a power law decay. Numerical results show that our data-driven leverage score sampling algorithm performs even better than the established theoretical error bounds, even in cases where the leverage scores are close to uniform distribution, which could be a by-product of our analysis. Future work could include extending these results to other data-driven sampling methods, including randomized leverage score sampling, and also implementing this algorithm extensively on real-world large-scale datasets.

REFERENCES

- Corwin L. Atwood. Sequences converging to D-optimal designs of experiments. *The Annals of Statistics*, 1(2):342–352, 1973. doi: 10.1214/aos/1176342371.
- Rajen Bhatt and Abhinav Dhall. Skin Segmentation. UCI Machine Learning Repository, 2012. DOI: <https://doi.org/10.24432/C5T30C>. License: CC BY 4.0.




Figure 1: Rotated Cauchy: Calculated optimality gap summary (a) and time summary (b).




Figure 2: Lognormal: Calculated optimality gap summary (a) and time summary (b).




Figure 3: Gaussian: Calculated optimality gap summary (a) and time summary (b).

- 486 Yu-Jen Chen, Ming-Yi Ju, and Kao-Shing Hwang. A virtual torque-based approach to kinematic
487 control of redundant manipulators. *IEEE Transactions on Industrial Electronics*, 64(2):1728–
488 1736, 2016. doi: 10.1109/TIE.2016.2548439.
- 489 F. L. Chernousko. Ellipsoidal state estimation for dynamical systems. *Nonlinear Analysis: Theory,
490 Methods & Applications*, 63(5-7):872–879, 2005. doi: 10.1016/j.na.2005.01.009.
- 491 Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron
492 Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 2015 Conference on
493 Innovations in Theoretical Computer Science*, pp. 181–190, 2015.
- 494 Michael B Cohen, Ben Cousins, Yin Tat Lee, and Xin Yang. A near-optimal algorithm for approxi-
495 mating the John ellipsoid. In *Conference on Learning Theory*, pp. 849–873. PMLR, 2019.
- 496 S Damla Ahipasaoglu, Peng Sun, and Michael J Todd. Linear convergence of a modified Frank–
497 Wolfe algorithm for computing minimum-volume enclosing ellipsoids. *Optimisation Methods
498 and Software*, 23(1):5–19, 2008.
- 501 P. Drineas, M. Magdon-Ismail, M.W. Mahoney, and D.P. Woodruff. Fast approximation of matrix
502 coherence and statistical leverage. *Journal of Machine Learning Research*, 13(Dec):3475–3506,
503 2012.
- 504 Petros Drineas, Michael W. Mahoney, and Shan Muthukrishnan. Relative-error CUR matrix decom-
505 positions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, 2008.
- 507 David Eberly. *3D game engine design: a practical approach to real-time computer graphics*. CRC
508 Press, 2006. doi: 10.1201/b18212.
- 509 Ali Eshragh, Luke Yerbury, Asef Nazari, Fred Roosta, and Michael W. Mahoney. SALSA: Se-
510 quential approximate leverage-score algorithm with application in analyzing big time series data.
511 *arXiv preprint arXiv:2401.00122*, 2023.
- 512 Jordi Fonollosa. Gas sensor array under dynamic gas mixtures. UCI Machine Learning Repository,
513 2015. DOI: <https://doi.org/10.24432/C5WP4C>. License: CC BY 4.0.
- 514 Marguerite Frank, Philip Wolfe, et al. An algorithm for quadratic programming. *Naval research
515 logistics quarterly*, 3(1-2):95–110, 1956. doi: 10.1002/nav.3800030109.
- 516 Radoslav Harman, Lenka Filová, and Peter Richtárik. A randomized exchange algorithm for com-
517 puting optimal approximate designs of experiments. *Journal of the American Statistical Associa-
518 tion*, 115(529):348–361, 2020. doi: 10.1080/01621459.2018.1546588.
- 519 Fritz John. Extremum problems with inequalities as subsidiary conditions. In *Traces and
520 Emergence of Nonlinear Programming*, pp. 197–215. Springer Basel, 2014. doi: 10.1007/
521 978-3-0348-0439-4_9.
- 522 Linus Källberg and Daniel Andrén. Active set strategies for the computation of minimum-volume
523 enclosing ellipsoids. Technical report, Mälardalen Real-Time Research Centre, Mälardalen Uni-
524 versity, November 2019. URL <http://www.es.mdu.se/publications/5680->.
- 525 Leonid G. Khachiyan. Rounding of polytopes in the real number model of computation. *Mathemat-
526 ics of Operations Research*, 21(2):307–320, 1996. doi: 10.1287/moor.21.2.307.
- 527 Leonid G. Khachiyan and Michael J. Todd. On the complexity of approximating the maximal in-
528 scribed ellipsoid for a polytope. *Mathematical Programming*, 61(1-3):137–159, 1993.
- 529 Jakub Kudela. Minimum-volume covering ellipsoids: Improving the efficiency of the Wolfe-
530 Atwood algorithm for large-scale instances by pooling and batching. In *MENDEL*, volume 25,
531 pp. 19–26, 2019.
- 532 Piyush Kumar and E. Alper Yildirim. Minimum-volume enclosing ellipsoids and core sets. *Journal
533 of Optimization Theory and applications*, 126(1):1–21, 2005. doi: 10.1007/s10957-005-2653-6.
- 534 Shannon R McCurdy, Vasilis Ntranos, and Lior Pachter. Deterministic column subset selection for
535 single-cell rna-seq. *Plos one*, 14(1):e0210571, 2019.

- 540 Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex program-*
541 *ming*. SIAM, 1994.
- 542
- 543 Bruno Ordozgoiti, Antonis Matakos, and Aristides Gionis. Generalized leverage scores: Geometric
544 interpretation and applications. In *International Conference on Machine Learning*, pp. 17056–
545 17070. PMLR, 2022.
- 546 Dimitris Papailiopoulos, Anastasios Kyriillidis, and Christos Boutsidis. Provable deterministic lever-
547 age score sampling. In *Proceedings of the 20th ACM SIGKDD international conference on*
548 *Knowledge discovery and data mining*, pp. 997–1006, 2014.
- 549 Samuel Rosa and Radoslav Harman. Computing minimum-volume enclosing ellipsoids for large
550 datasets. *Computational Statistics & Data Analysis*, 171:107452, 2022. doi: 10.1016/j.csda.
551 2022.107452.
- 552
- 553 Judah Ben Rosen. Pattern separation by convex programming. *Journal of Mathematical Analysis*
554 *and Applications*, 10(1):123–134, 1965. doi: 10.1016/0022-247X(65)90150-2.
- 555 Peter J. Rousseeuw and Mia Hubert. Robust statistics for outlier detection. *Wiley interdisciplinary*
556 *reviews: Data mining and knowledge discovery*, 1(1):73–79, 2011. doi: 10.1002/widm.2.
- 557
- 558 Fred Schweppe. Recursive state estimation: Unknown but bounded errors and system inputs. *IEEE*
559 *Transactions on Automatic Control*, 13(1):22–28, 1968. doi: 10.1109/TAC.1968.1098790.
- 560 R. Sibson. Discussion of Dr Wynn’s and of Dr Laycock’s papers. *Journal of the Royal Statisti-*
561 *cal Society: Series B (Methodological)*, 34(2):181–183, 1972. doi: 10.1111/j.2517-6161.1972.
562 tb00898.x.
- 563 S. D. Silvey. *Optimal design: and introduction to the theory for parameter estimation*. London,
564 Chapman and Hall/CRC, 1980. doi: 10.1007/978-94-009-5912-5.
- 565
- 566 Zhao Song, Xin Yang, Yuanyuan Yang, and Tianyi Zhou. Faster algorithm for structured John
567 ellipsoid computation. *arXiv preprint arXiv:2211.14407*, 2022.
- 568
- 569 Peng Sun and Robert M. Freund. Computation of minimum-volume covering ellipsoids. *Operations*
570 *Research*, 52(5):690–706, 2004. doi: 10.1287/opre.1040.0115.
- 571 D. M. Titterton. Optimal design: some geometrical aspects of D-optimality. *Biometrika*, 62(2):
572 313–320, 1975. doi: 10.2307/2335366.
- 573
- 574 D. M. Titterton. Estimation of correlation coefficients by ellipsoidal trimming. *Journal of the*
575 *Royal Statistical Society: Series C (Applied Statistics)*, 27(3):227–234, 1978. doi: 10.2307/
576 2347157.
- 577 Michael J. Todd. *Minimum-volume ellipsoids: Theory and algorithms*. SIAM, 2016a. doi: 10.1137/
578 1.9781611974386.
- 579
- 580 Michael J. Todd. Minvol. SIAM, 2016b. URL: <http://archive.siam.org/books/mo23/>.
- 581 Michael J. Todd and E. Alper Yıldırım. On Khachiyan’s algorithm for the computation of minimum-
582 volume enclosing ellipsoids. *Discrete Applied Mathematics*, 155(13):1731–1744, 2007. doi:
583 10.1016/j.dam.2007.02.013.
- 584 Philip Wolfe. Convergence theory in nonlinear programming. *Integer and nonlinear programming*,
585 pp. 1–36, 1970. doi: 10.1007/BF00932858.
- 586
- 587 David P Woodruff and Taisuke Yasuda. New subset selection algorithms for low rank approximation:
588 Offline and online. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*,
589 pp. 1802–1813, 2023.
- 590 David P. Woodruff and Taisuke Yasuda. John ellipsoids via lazy updates. *arXiv preprint*
591 *arXiv:2501.01801*, 2025.
- 592
- 593 Yaming Yu. D-optimal designs via a cocktail algorithm. *Statistics and Computing*, 21(4):475–481,
2011. doi: 10.1007/s11222-010-9183-2.

A DETERMINISTIC APPROXIMATE LEVERAGE SCORE SAMPLING

The computation time in Algorithm 1 is dominated by the cost of calculating the leverage scores exactly, which is $\mathcal{O}(nd^2)$. We can instead use approximate leverage scores, which are more computationally efficient to calculate. We demonstrate the approach using the algorithm of Drineas et al. (2012).

Theorem 6 (Drineas et al. (2012), Theorem 2). *Fix a constant $\beta \in (0, 1/2]$. Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be full rank, with $n \gg d$. Let the leverage scores of the rows of \mathbf{X} be given by $\ell_1(\mathbf{X}), \dots, \ell_n(\mathbf{X})$. Then there exists a randomized algorithm that calculates approximate leverage scores $\hat{\ell}_1(\mathbf{X}), \dots, \hat{\ell}_n(\mathbf{X})$ such that with probability at least 0.8, simultaneously for all $i = 1, \dots, n$, $\hat{\ell}_i(\mathbf{X}) = (1 \pm \beta) \ell_i(\mathbf{X})$. This algorithm has time complexity $\mathcal{O}(nd \log n)$.*

We then sample deterministically according to these approximate leverage scores. We summarize the modified sampling procedure in Algorithm 2.

Algorithm 2 Deterministic Approximate Leverage Score Sampling with Threshold

Require: $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$, $\beta \in (0, 1/2]$, $\varepsilon \in (0, 1)$

- 1: Calculate approximate leverage scores for each row in \mathbf{X} , using Theorem 6. Assume $\hat{\ell}_1(\mathbf{X}) \geq \dots \geq \hat{\ell}_n(\mathbf{X})$.
 - 2: Let $t = \sum_{i=1}^n \hat{\ell}_i(\mathbf{X})$.
 - 3: Let $s = \arg \min_j \left(\sum_{i=1}^j \hat{\ell}_i(\mathbf{X}) \geq t - (1 - \beta) \varepsilon \right)$.
 - 4: Let $\mathbf{R} = \mathbf{0} \in \mathbb{R}^{s \times n}$.
 - 5: **for** $i = 1 : s$ **do**
 - 6: Set row i of \mathbf{R} equal to \mathbf{e}_i .
 - 7: **end for**
- Ensure:** \mathbf{R}, s
-

We have chosen s carefully, to ensure the following subspace embedding result.

Corollary 7. *Let $\beta \in (0, 1/2]$, and $\varepsilon \in (0, 1)$. Use Algorithm 2 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$. Then, with probability at least 0.8, we have*

$$(1 - \varepsilon) \mathbf{X}^\top \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{X}_s \preceq \mathbf{X}^\top \mathbf{X}.$$

Proof. For this proof, we only need to show that $\sum_{i=s+1}^n \ell_i(\mathbf{X}) < \varepsilon$ holds with probability at least 0.8. We can then apply Theorem 2 to obtain our bound.

Assume that $\hat{\ell}_i(\mathbf{X}) \geq (1 - \beta) \ell_i(\mathbf{X})$ holds simultaneously for all $i = 1, \dots, n$. By construction of Algorithm 2, we have

$$\sum_{i=s+1}^n \hat{\ell}_i(\mathbf{X}) < (1 - \beta) \varepsilon.$$

Then since $\hat{\ell}_i(\mathbf{X}) \geq (1 - \beta) \ell_i(\mathbf{X})$, for all $i = 1, \dots, n$, we have

$$(1 - \beta) \varepsilon > \sum_{i=s+1}^n \hat{\ell}_i(\mathbf{X}) \geq \sum_{i=s+1}^n (1 - \beta) \ell_i(\mathbf{X}) = (1 - \beta) \sum_{i=s+1}^n \ell_i(\mathbf{X}),$$

that is,

$$\sum_{i=s+1}^n \ell_i(\mathbf{X}) < \varepsilon.$$

□

Corollary 8. *If the leverage scores exhibit a power law decay, then $s = \text{poly} \left(d(1 + \beta), \frac{1}{(1 - \beta)\varepsilon} \right)$.*

648 *Proof.* Papailiopoulos et al. (2014) show that if the leverage scores exhibit a power law decay, then
 649 Algorithm 1 outputs

$$650$$

$$651 \quad s = \arg \min_j \left(\sum_{i=1}^j \ell_i(\mathbf{X}) > d - \varepsilon \right) = \max \left\{ \left(\frac{2d}{\varepsilon} \right)^{\frac{1}{1+\eta}} - 1, \left(\frac{2d}{\eta\varepsilon} \right)^{\frac{1}{\eta}} - 1, d \right\} = \text{poly} \left(d, \frac{1}{\varepsilon} \right).$$

$$652$$

$$653$$

654 Then, comparing Algorithm 1 with Algorithm 2, we need only replace d with t , and ε with $(1 - \beta)\varepsilon$.
 655 That is, Algorithm 2 outputs

$$656 \quad s = \arg \min_j \left(\sum_{i=1}^j \hat{\ell}_i(\mathbf{X}) > t - (1 - \beta)\varepsilon \right) = \text{poly} \left(t, \frac{1}{(1 - \beta)\varepsilon} \right).$$

$$657$$

$$658$$

659 But

$$660 \quad t = \sum_{i=1}^n \hat{\ell}_i(\mathbf{X}) \leq \sum_{i=1}^n (1 + \beta) \ell_i(\mathbf{X}) = (1 + \beta) \sum_{i=1}^n \ell_i(\mathbf{X}) = (1 + \beta) d.$$

$$661$$

$$662$$

663 Hence $s = \text{poly} \left(d(1 + \beta), \frac{1}{(1 - \beta)\varepsilon} \right)$. □
 664
 665

666 B REPEATED APPLICATIONS OF LEVERAGE SCORE SAMPLING

$$667$$

668 For many datasets, a single application of Algorithm 1 does not sufficiently reduce the number of
 669 points considered by the MVCE algorithm. We show how multiple applications of Algorithm 1 can
 670 still guarantee a subspace embedding result. We summarize the modified sampling procedure in
 671 Algorithm 3.

672 **Algorithm 3** Repeated Deterministic Leverage Score Sampling with Threshold

674 **Require:** $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$, $\varepsilon \in (0, 1)$, $m \in \mathbb{N}$

- 675 1: Let $\mathbf{X}_{(0)} = \mathbf{X}$.
- 676 2: **for** $k = 1 : m$ **do**
- 677 3: Calculate leverage scores for each row in $\mathbf{X}_{(k)}$. Assume $\ell_1(\mathbf{X}_{(k)}) \geq \dots \geq \ell_n(\mathbf{X}_{(k)})$.
- 678 4: Let $s = \arg \min_j \left(\sum_{i=1}^j \ell_i(\mathbf{X}_{(k)}) > d - \varepsilon \right)$.
- 679 5: Let $\mathbf{R}_{(k)} = \mathbf{0} \in \mathbb{R}^{s \times n}$.
- 680 6: **for** $i = 1 : s$ **do**
- 681 7: Set row i of $\mathbf{R}_{(k)}$ equal to \mathbf{e}_i .
- 682 8: **end for**
- 683 9: Let $\mathbf{X}_{(k+1)} = \mathbf{R}_{(k)} \mathbf{X}_{(k)}$.
- 684 10: **if** $\text{rank}(\mathbf{X}_{(k+1)}) < d$ **then**
- 685 11: Redefine $m := k$.
- 686 12: **break**
- 687 13: **end if**
- 688 14: **end for**

689 **Ensure:** $\mathbf{X}_s = \mathbf{X}_{(m)}$

691 We now present the subspace embedding result.

692 **Corollary 9.** Let $\varepsilon \in (0, 1)$, and $m \in \mathbb{N}$. Use Algorithm 3 to construct \mathbf{X}_s . Then

$$693 \quad (1 - \varepsilon)^m \mathbf{X}^\top \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{X}_s.$$

$$694$$

$$695$$

696 *Proof.* The proof relies on the repeated application of Theorem 2. Essentially, we prove

$$697 \quad (1 - \varepsilon)^m \mathbf{X}_{(0)}^\top \mathbf{X}_{(0)} \prec (1 - \varepsilon)^{m-1} \mathbf{X}_{(1)}^\top \mathbf{X}_{(1)} \prec \dots \prec (1 - \varepsilon) \mathbf{X}_{(m-1)}^\top \mathbf{X}_{(m-1)} \prec \mathbf{X}_{(m)}^\top \mathbf{X}_{(m)},$$

$$698 \quad (7)$$

$$699$$

700 which we can show by induction. When $m = 1$, Theorem 2 implies

$$701 \quad (1 - \varepsilon)^1 \mathbf{X}_{(0)}^\top \mathbf{X}_{(0)} = (1 - \varepsilon) \mathbf{X}^\top \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{X}_s = \mathbf{X}_{(1)}^\top \mathbf{X}_{(1)},$$

702 proving the base case. Now suppose (7) holds for some $m = k > 1$. That is, assume

$$703 (1 - \varepsilon)^k \mathbf{X}_{(0)}^\top \mathbf{X}_{(0)} \prec \cdots \prec \mathbf{X}_{(k)}^\top \mathbf{X}_{(k)}. \quad (8)$$

704 That is, we have

$$705 (1 - \varepsilon)^{k+1} \mathbf{X}_{(0)}^\top \mathbf{X}_{(0)} \prec \cdots \prec (1 - \varepsilon) \mathbf{X}_{(k)}^\top \mathbf{X}_{(k)}. \quad (9)$$

706 Additionally, if we sample from $\mathbf{X}_{(k)}$ using Algorithm 1, Theorem 2 guarantees that

$$707 (1 - \varepsilon) \mathbf{X}_{(k)}^\top \mathbf{X}_{(k)} \prec \mathbf{X}_{(k+1)}^\top \mathbf{X}_{(k+1)}. \quad (10)$$

708 Combining (9) and (10), we obtain

$$709 (1 - \varepsilon)^{k+1} \mathbf{X}_{(0)}^\top \mathbf{X}_{(0)} \prec \cdots \prec (1 - \varepsilon) \mathbf{X}_{(k)}^\top \mathbf{X}_{(k)} \prec \mathbf{X}_{(k+1)}^\top \mathbf{X}_{(k+1)},$$

710 finalizing the proof. \square

711 C MISSING DETAILS IN PROOF OF THEOREM 3

712 PROOF OF EQUATIONS (1) AND (2)

713 We have that

$$714 \begin{aligned} g(\mathbf{u}_K) &= \log \det (\mathbf{X}^\top \mathbf{U}_K \mathbf{X}) \\ 715 &= \log \det \left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} \right) \\ 716 &= \log \left(\frac{1}{n^d} \det (\mathbf{X}^\top \mathbf{X}) \right) \\ 717 &= \log \det (\mathbf{X}^\top \mathbf{X}) - d \log n, \end{aligned}$$

718 and, similarly,

$$719 \begin{aligned} g_s(\mathbf{u}_0) &= \log \det (\mathbf{X}_s^\top \mathbf{U}_0 \mathbf{X}_s) \\ 720 &= \log \det \left(\frac{1}{s} \mathbf{X}_s^\top \mathbf{X}_s \right) \\ 721 &= \log \det (\mathbf{X}_s^\top \mathbf{X}_s) - d \log s. \end{aligned}$$

722 PROOF OF INEQUALITY (4)

723 We first need to prove the following lemma.

724 **Lemma 10.** *Suppose we have two $d \times d$ symmetric positive semidefinite matrices \mathbf{A} and \mathbf{B} , with*
 725 *$\mathbf{A} \preceq \mathbf{B}$. Let the eigenvalues of \mathbf{A} be given by $\lambda_1(\mathbf{A}) \leq \cdots \leq \lambda_d(\mathbf{A})$, and the eigenvalues of \mathbf{B}*
 726 *be given by $\lambda_1(\mathbf{B}) \leq \cdots \leq \lambda_d(\mathbf{B})$. Then for all $i = 1, \dots, d$, we have*

$$727 \lambda_i(\mathbf{A}) \leq \lambda_i(\mathbf{B}).$$

728 *Proof.* Use Min-max Theorem, along with the fact that $\mathbf{x}^\top \mathbf{A} \mathbf{x} \leq \mathbf{x}^\top \mathbf{B} \mathbf{x}$ for all \mathbf{x} . \square

729 We can now prove Inequality (4).

730 **Lemma 11.** *Suppose we have two $d \times d$ symmetric positive semidefinite matrices \mathbf{A} and \mathbf{B} , with*
 731 *$\mathbf{A} \preceq \mathbf{B}$. Then*

$$732 \det(\mathbf{A}) \leq \det(\mathbf{B}).$$

756 *Proof.* We have that

$$\begin{aligned}
 757 & \log \det(\mathbf{A}) = \log \left(\prod_{i=1}^d \lambda_i(\mathbf{A}) \right) \\
 758 & = \sum_{i=1}^d \log \lambda_i(\mathbf{A}) \\
 759 & \leq \sum_{i=1}^d \log \lambda_i(\mathbf{B}) \\
 760 & = \log \left(\prod_{i=1}^d \lambda_i(\mathbf{B}) \right) \\
 761 & = \log \det(\mathbf{B}),
 \end{aligned}$$

762 where the inequality uses Lemma 10 and the fact that log is a monotonically increasing function.
763 Then

$$764 \det(\mathbf{A}) = e^{\log \det(\mathbf{A})} \leq e^{\log \det(\mathbf{B})} = \det(\mathbf{B}),$$

765 where the inequality is because the exponential function is monotonically increasing. \square

766 D FINAL OPTIMALITY GAP

767 We break this proof into two cases:

- 768 1. $u_i^* > 0$, for all $i = 1, \dots, s$, and
- 769 2. $u_i^* \geq 0$, for all $i = 1, \dots, s$.

770 We do this because the proof of Inequality (6) relies on scaling $\sqrt{\mathbf{U}^*}$ so that the first s diagonal
771 terms are ≥ 1 , and the last $n - s$ diagonal terms are ≤ 1 . Having any zeros in the first s terms
772 prevents such a scaling.

773 We consider Case 1 first. First, follow the proof outline from Section 6. Then, what remains to be
774 shown is that Inequality (6) holds.

775 CASE 1: $u_i^* > 0$, FOR ALL $i = 1, \dots, s$

776 The following lemma will be useful for proving Inequality (6). It shows the change in leverage
777 scores when one row of \mathbf{X} is scaled.

778 **Lemma 12** (Ordozgoiti et al. (2022), Lemma 3.2). *Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ have rank d . Let $\alpha \geq 0$. Define
779 $\Sigma^{(i)}$ as the diagonal matrix satisfying $\Sigma_{ii}^{(i)} = \alpha$, $\Sigma_{jj}^{(i)} = 1$, if $j \neq i$. Then*

$$780 \ell_j \left(\Sigma^{(i)} \mathbf{X} \right) = \ell_j(\mathbf{X}) - \frac{(\alpha^2 - 1) \left(\mathbf{x}_j^\top (\mathbf{X}^\top \mathbf{X})^\top \mathbf{x}_i \right)^2}{1 + (\alpha^2 - 1) \ell_i(\mathbf{X})},$$

781 and, in particular,

$$782 \ell_i \left(\Sigma^{(i)} \mathbf{X} \right) = \frac{\alpha^2 \ell_i(\mathbf{X})}{1 + (\alpha^2 - 1) \ell_i(\mathbf{X})}.$$

783 We now ready to prove the inequality.

784 **Proposition 13.** *Use Algorithm 1 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$, with $\varepsilon \in (0, 1)$. Suppose that an optimal
785 solution \mathbf{u}^* of (D) satisfies $u_i^* > 0$ for all $i = 1, \dots, s$. Define $\mathbf{Y} := \sqrt{\mathbf{U}^*} \mathbf{X}$. Then*

$$786 \sum_{i=s+1}^n \ell_i(\mathbf{Y}) \leq \sum_{i=s+1}^n \ell_i(\mathbf{X}).$$

810 *Proof.* The structure of this proof is based on the proof of Theorem 3.2 in Ordozgoiti et al. (2022).
 811 We will show that the sum of the last $n-s$ leverage scores of \mathbf{X} do not increase when we premultiply
 812 it by $\sqrt{\mathbf{U}^*}$. Without loss of generality, we assume that $\min_{i=1,\dots,s} \sqrt{u_i^*} \geq \max_{i=s+1,\dots,n} \sqrt{u_i^*}$.
 813 This is because, for any k , we can write

$$814 \quad u_k^* \mathbf{x}_k \mathbf{x}_k^\top = \sum_{i=1}^{\beta} \frac{u_k^*}{\beta} \mathbf{x}_k \mathbf{x}_k^\top,$$

815 where $\beta \in \mathbf{N}$, such that $\frac{u_k^*}{\beta}$ is sufficiently small.
 816

817 First, we scale $\sqrt{\mathbf{U}^*}$ so that the first s diagonal terms are ≥ 1 , and the last $n-s$ diagonal terms are
 818 ≤ 1 . Let $\alpha = 1/\min_{i=1,\dots,s} \sqrt{u_i^*}$. We define

$$819 \quad \mathbf{V} := \alpha \sqrt{\mathbf{U}^*}.$$

820 We exploit the fact that

$$\begin{aligned} 821 \quad \ell_i(\mathbf{V}\mathbf{X}) &= \ell_i(\alpha \sqrt{\mathbf{U}^*} \mathbf{X}) \\ 822 \quad &= \ell_i(\alpha \mathbf{Y}) \\ 823 \quad &= \left((\alpha \mathbf{Y}) \left((\alpha \mathbf{Y})^\top (\alpha \mathbf{Y}) \right)^{-1} (\alpha \mathbf{Y})^\top \right)_{ii} \\ 824 \quad &= \left(\alpha^2 \frac{1}{\alpha^2} \mathbf{Y} (\mathbf{Y}^\top \mathbf{Y})^{-1} \mathbf{Y}^\top \right)_{ii} \\ 825 \quad &= \ell_i(\mathbf{Y}). \end{aligned}$$

826 To see how \mathbf{V} affects the leverage scores of \mathbf{X} , we will consider scaling each row separately. We
 827 define $\mathbf{V}^{(i)}$ as the diagonal matrix satisfying $\mathbf{V}_{ii}^{(i)} = \mathbf{V}_{ii} = \alpha \sqrt{u_i^*}$ and $\mathbf{V}_{jj}^{(i)} = 1$, if $j \neq i$. Now,
 828 we consider the leverage scores of $\mathbf{V}^{(i)}\mathbf{X}$. From Lemma 12, we have

- 829 1. If $i \leq s$, then $\alpha \sqrt{u_i^*} \geq 1$. Hence $\ell_i(\mathbf{V}^{(i)}\mathbf{X}) \geq \ell_i(\mathbf{X})$, and $\ell_j(\mathbf{V}^{(i)}\mathbf{X}) \leq \ell_j(\mathbf{X})$, if
 830 $j \neq i$.
- 831 2. If $i \geq s$, then $\alpha \sqrt{u_i^*} \leq 1$. Hence $\ell_i(\mathbf{V}^{(i)}\mathbf{X}) \leq \ell_i(\mathbf{X})$, and $\ell_j(\mathbf{V}^{(i)}\mathbf{X}) \geq \ell_j(\mathbf{X})$, if
 832 $j \neq i$.

833 Now, we consider scaling the first s rows. From the previous discussion, we can conclude that the
 834 leverage scores of the last $n-s$ rows do not increase. Therefore, the sum of the first s leverage scores
 835 does not decrease. (This is because the sum of the leverage scores remains constant, since scaling
 836 rows by non-zero constants does not affect the rank.) Next, consider scaling the last $n-s$ rows.
 837 From the discussion above, all the leverage scores of the first s rows cannot decrease. That is, the
 838 sum of the first s leverage scores again does not decrease. We conclude that the sum of the last $n-s$
 839 leverage scores cannot increase. (This is because the sum of the leverage scores remains constant.
 840 Suppose $\text{rank}(\mathbf{V}\mathbf{X}) < d$. Then $\text{rank}(\mathbf{X}^\top \mathbf{U}^* \mathbf{X}) = \text{rank}\left(\left(\mathbf{V}\mathbf{X}\right)^\top \left(\mathbf{V}\mathbf{X}\right)\right) = \text{rank}(\mathbf{V}\mathbf{X}) < d$,
 841 which is a contradiction, since $\mathbf{X}^\top \mathbf{U}^* \mathbf{X}$ is invertible.) Hence

$$842 \quad \sum_{i=1}^n \ell_i(\mathbf{Y}) = \sum_{i=1}^n \ell_i(\mathbf{V}\mathbf{X}) \leq \sum_{i=1}^n \ell_i(\mathbf{X}).$$

843 \square

844 CASE 2: $u_i^* \geq 0$, FOR ALL $i = 1, \dots, s$

845 This time, we will examine a δ -feasible solution \mathbf{u} for (D), which has $u_i > 0$ for all $i = 1, \dots, n$.
 846 Such a solution can be achieved by using the Frank-Wolfe algorithm (Frank et al., 1956; Wolfe,
 847 1970) with Khachiyan's initialisation (Khachiyan, 1996).

864 Consider the feasible solution $\tilde{\mathbf{u}}_s$ for (D_s) , given by

$$865 \tilde{\mathbf{u}}_s = \frac{1}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{u}_s,$$

866 where \mathbf{u}_s contains the first s entries of \mathbf{u} . We would like a bound similar to the one in Theorem 2,
867 with \mathbf{X} replaced with

$$870 \mathbf{Y} := \sqrt{\mathbf{U}} \mathbf{X}.$$

871 We require the following result.

872 **Corollary 14.** *Use Algorithm 1 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$, with $\varepsilon \in (0, 1)$. Let \mathbf{u} be a δ -feasible
873 solution to (D) . Then*

$$874 \sum_{i=s+1}^n \ell_i(\mathbf{Y}) \leq \sum_{i=s+1}^n \ell_i(\mathbf{X}).$$

875 *Proof.* Follow the proof of Proposition 13, but replace \mathbf{u}^* with \mathbf{u} , and redefine $\mathbf{Y} := \sqrt{\mathbf{U}} \mathbf{X}$. \square

880 Then, by the construction of Algorithm 1, we have

$$881 \sum_{i=s+1}^n \ell_i(\mathbf{Y}) \leq \sum_{i=s+1}^n \ell_i(\mathbf{X}) < \varepsilon,$$

882 for some $\varepsilon \in (0, 1)$. Therefore, we may apply Theorem 2 with \mathbf{Y} instead of \mathbf{X} , to obtain the bound

$$883 (1 - \varepsilon) \mathbf{X}^\top \mathbf{U} \mathbf{X} \prec \mathbf{X}_s^\top \mathbf{U}_s \mathbf{X}_s.$$

884 That is, the feasible solution $\tilde{\mathbf{u}}_s$ satisfies

$$885 (1 - \varepsilon) \mathbf{X}^\top \mathbf{U} \mathbf{X} \preceq \frac{1 - \varepsilon}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{X}^\top \mathbf{U}_s \mathbf{X} \prec \frac{1}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{X}_s^\top \mathbf{U}_s \mathbf{X}_s = \mathbf{X}_s^\top \tilde{\mathbf{U}}_s \mathbf{X}_s.$$

886 We are now ready to prove the final optimality gap.

887 **Theorem 15.** *Use Algorithm 1 to construct $\mathbf{X}_s = \mathbf{R}\mathbf{X}$, with $\varepsilon \in (0, 1)$. Let \mathbf{u} be a δ -feasible
888 solution to (D) . Then*

$$889 g^* - g_s^* < d \log \left(\frac{1 + \delta}{1 - \varepsilon} \right).$$

890 *Proof.* For this proof, we will exploit the fact that

$$891 g^* - g_s^* = (g^* - g(\mathbf{u})) + (g(\mathbf{u}) - g_s^*).$$

892 We will bound above $g(\mathbf{u}) - g_s^*$ first. Let $\tilde{\mathbf{u}}_s = \frac{1}{\mathbf{e}^\top \mathbf{u}_s} \mathbf{u}_s$, where \mathbf{u}_s contains the first s entries of
893 \mathbf{u} . Since g_s is concave, the optimal value g_s^* must be greater or equal to g_s at any feasible point for
894 (D_s) . Hence

$$895 g_s^* \geq g_s(\tilde{\mathbf{u}}_s) \\ 896 = \log \det \left(\mathbf{X}_s^\top \tilde{\mathbf{U}}_s \mathbf{X}_s \right) \\ 897 > \log \det \left((1 - \varepsilon) \mathbf{X}^\top \mathbf{U} \mathbf{X} \right) \\ 898 = d \log (1 - \varepsilon) + \log \det \left(\mathbf{X}^\top \mathbf{U} \mathbf{X} \right),$$

899 where the inequality uses Lemma 11. Hence

$$900 g(\mathbf{u}) - g_s^* < \log \det \left(\mathbf{X}^\top \mathbf{U} \mathbf{X} \right) - \left(d \log (1 - \varepsilon) + \log \det \left(\mathbf{X}^\top \mathbf{U} \mathbf{X} \right) \right) \\ 901 = d \log \left(\frac{1}{1 - \varepsilon} \right).$$

Now, recall that \mathbf{u} is δ -feasible for (D). By Proposition 1, we have

$$g^* - g(\mathbf{u}) \leq d \log(1 + \delta).$$

Hence

$$\begin{aligned} g^* - g_s^* &= (g^* - g(\mathbf{u})) + (g(\mathbf{u}) - g_s^*) \\ &< d \log\left(\frac{1}{1 - \varepsilon}\right) + d \log(1 + \delta) \\ &= d \log\left(\frac{1 + \delta}{1 - \varepsilon}\right). \end{aligned}$$

□

E MORE NUMERICAL RESULTS

COMPARISON OF OUR ALGORITHM WITH THE FIXED POINT ALGORITHM

All computations in this section are performed on a personal laptop with a 64 bit Windows 11 Home operating system (Version 23H2), and a 3.30 GHz 11th Gen Intel Core i7-11370H processor with 40 GB of RAM. The algorithms are run using MATLAB (R2021a).

In this subsection, we compare our algorithm to the Fixed Point algorithm (Cohen et al., 2019). For consistency, we run our algorithm using $s = 10\%$ of n . We run both algorithms on Gaussian datasets of size $n = 1\,000\,000$, with dimension d ranging from 2 to 10. Table 1 summarizes the total computation time required for both algorithms, for varying values of accuracy parameter δ .

Table 1: Time summary for the two algorithms. Tests run on Gaussian datasets of size $n = 1\,000\,000$, and d ranging from 2 to 10. Each value represents mean and sample standard deviation of 100 runs.

n	d	δ	Time (s)		
			Our algorithm	Fixed Point algorithm	
1 000 000	2	10^{-2}	0.0914 ± 0.0220	1.4331 ± 0.0877	
		10^{-3}	0.0922 ± 0.0210	12.1627 ± 0.8764	
		10^{-4}	0.0873 ± 0.0215	63.1388 ± 6.2418	
		5	10^{-2}	0.1330 ± 0.0219	3.6558 ± 0.2074
	10^{-3}		0.1956 ± 0.0442	12.1783 ± 1.3380	
	10^{-4}		0.2125 ± 0.0390	47.2873 ± 11.6625	
	10	10^{-2}	0.2483 ± 0.0618	6.0298 ± 0.7460	
		10^{-3}	0.4177 ± 0.1025	21.8280 ± 3.9056	
		10^{-4}		0.4313 ± 0.1077	117.9867 ± 41.5669

Table 2: Time summary for the two algorithms. Tests run on Gaussian datasets of dimension $d = 100$, with varying n . Each value represents mean and sample standard deviation (when multiple runs are computed).

d	n	δ	# runs	Time (s)	
				Our algorithm	Fixed Point algorithm
100	10 000	10^{-2}	100	0.0927 ± 0.0402	0.7270 ± 0.1886
		10^{-3}	100	0.2578 ± 0.1991	5.6216 ± 1.5382
		10^{-4}	100	0.3013 ± 0.0833	18.6236 ± 4.8307
		10^{-5}	50	0.5162 ± 0.1292	82.4097 ± 26.7802
		100 000	10^{-3}	20	1.9359 ± 0.1371
	1 000 000		10^{-3}	1	101.0625

972 We now apply the same methods to Gaussian matrices of dimension $d = 100$, with varying number
973 of points n and accuracy parameter δ . Table 2 summarizes the total computation time required for
974 both algorithms.

975 In both Tables 1 and 2, we notice that as δ increases, the time taken for both methods generally
976 increases. However, the increase in time for our algorithm is minimal, while the increase in time for
977 the Fixed Point algorithm is significant. In the final row of Table 2, our algorithm takes less than 2
978 minutes, compared with the Fixed Point algorithm, which took about 80 minutes.

979 We emphasize that the examples considered here are smaller than those we considered in Section 7,
980 and the accuracy parameter δ is much larger than our desired 10^{-9} . Regardless, we ran the Fixed
981 Point algorithm on the examples from Section 7. We found that for all three datasets, the algorithm
982 did not converge in two hours of runtime.
983

984 REAL WORLD DATA

985 All computations in this section are performed on a personal laptop with a 64 bit MacOS 13 operat-
986 ing system, and a 2.4 GHz Quad-Core Intel Core i5 processor with 8 GB of RAM. The algorithms
987 are run using MATLAB (R2021a).

988 We now test our algorithm on three smaller datasets from the UCI Machine Learning Repository.
989 Table 3 summarizes the size of each dataset, and time to compute the MVCE using the WA algorithm
990 (Wolfe, 1970; Atwood, 1973).
991

992 We then apply our algorithm on each dataset, varying sample size s from 1% to 10% of n . In Table
993 4, we summarize the optimality gaps. The deterministic sampling performed well on all datasets,
994 achieving very small optimality gaps when s is 10% of n . In Table 5, we summarize the computation
995 times. Although the deterministic sampling was the slowest, all sampling methods decreased the
996 total computation time required to calculate the MVCE.
997

998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

Table 3: Dataset description. This includes dimension and time to compute MVCE using the WA algorithm.

Dataset	n	d	Time (s)
Ethylene CO (Fonollosa, 2015)	4 208 261	19	85.92 ± 0.17
Ethylene CH4 (Fonollosa, 2015)	4 178 504	19	75.89 ± 0.38
Skin (Bhatt & Dhall, 2012)	245 057	4	0.54 ± 0.04

Table 4: Optimality gap summary. Sample size s takes values of 1%, 5%, and 10% of n . Each value represents mean and sample standard deviation of 5 runs.

Dataset	s	det	Optimality Gap	
			prob	unif
Ethylene CO (Fonollosa, 2015)	1%	4.66	1.01 ± 0.11	5.47 ± 0.65
	5%	0.02	0.43 ± 0.03	4.40 ± 1.03
	10%	5.67E-06	0.25 ± 0.04	3.24 ± 0.48
Ethylene CH4 (Fonollosa, 2015)	1%	1.59	1.48 ± 0.09	3.53 ± 0.42
	5%	0.05	0.66 ± 0.05	1.69 ± 0.15
	10%	0.01	0.37 ± 0.08	1.35 ± 0.11
Skin (Bhatt & Dhall, 2012)	1%	0.75	0.07 ± 0.02	0.25 ± 0.13
	5%	0.56	0.02 ± 0.01	0.05 ± 0.03
	10%	-3.55E-15	0.03 ± 0.01	0.04 ± 0.02

Table 5: Time summary. Sample size s takes values of 1%, 5%, and 10% of n . Each value represents mean and sample standard deviation of 5 runs.

Dataset	s	det	Time (s)	
			prob	unif
Ethylene CO (Fonollosa, 2015)	1%	6.00 ± 0.12	5.42 ± 0.74	4.01 ± 0.39
	5%	15.00 ± 0.56	9.94 ± 0.47	7.55 ± 0.05
	10%	15.71 ± 0.14	14.32 ± 1.26	11.66 ± 0.95
Ethylene CH4 (Fonollosa, 2015)	1%	5.91 ± 0.30	5.24 ± 0.36	3.35 ± 0.30
	5%	13.56 ± 0.58	10.00 ± 1.10	7.95 ± 1.08
	10%	15.18 ± 0.34	13.05 ± 0.81	10.26 ± 0.93
Skin (Bhatt & Dhall, 2012)	1%	0.04 ± 0.01	0.04 ± 0.01	0.05 ± 0.03
	5%	0.15 ± 0.03	0.13 ± 0.01	0.10 ± 0.02
	10%	0.15 ± 0.03	0.17 ± 0.03	0.13 ± 0.02