

Solving Large-Scale Sparse PCA to Certifiable (Near) Optimality

Dimitris Bertsimas

DBERTSIM@MIT.EDU

*Massachusetts Institute of Technology
Cambridge, MA 02139, USA*

Ryan Cory-Wright

RYANCW@MIT.EDU

*Massachusetts Institute of Technology
Cambridge, MA 02139, USA*

Jean Pauphilet

JPAUPHILET@LONDON.EDU

*London Business School
London, UK*

Editor: TBD

Abstract

Sparse principal component analysis (PCA) is a popular dimensionality reduction technique for obtaining principal components which are linear combinations of a small subset of the original features. Existing approaches cannot supply certifiably optimal principal components with more than $p = 100$ s of variables. By reformulating sparse PCA as a convex mixed-integer semidefinite optimization problem, we design a cutting-plane method which solves the problem to certifiable optimality at the scale of selecting $k = 5$ covariates from $p = 300$ variables, and provides small bound gaps at a larger scale. We also propose a convex relaxation and greedy rounding scheme that provides bound gaps of $1 - 2\%$ in practice within minutes for $p = 100$ s or hours for $p = 1,000$ s and is therefore a viable alternative to the exact method at scale. Using real-world financial and medical datasets, we illustrate our approach's ability to derive interpretable principal components tractably at scale.

Keywords: Sparse PCA, Sparse Eigenvalues, Semidefinite Optimization

1. Introduction

In the era of big data, interpretable methods for compressing a high-dimensional dataset into a lower dimensional set which shares the same essential characteristics are imperative. Since the work of Hotelling (1933), principal component analysis (PCA) has been one of the most popular approaches for completing this task. Formally, given centered data $\mathbf{A} \in \mathbb{R}^{n \times p}$ and its normalized empirical covariance matrix $\mathbf{\Sigma} := \frac{\mathbf{A}\mathbf{A}^\top}{n-1} \in \mathbb{R}^{p \times p}$, PCA selects one or more leading eigenvectors of $\mathbf{\Sigma}$ and subsequently projects \mathbf{A} onto these eigenvectors. This can be achieved in $O(p^3)$ time by taking a singular value decomposition $\mathbf{\Sigma} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$.

A common criticism¹ of PCA is that the columns of \mathbf{U} are not interpretable, since each eigenvector is a linear combination of all p original features. This causes difficulties because:

- In medical applications such as cancer detection, PCs generated during exploratory data analysis need to supply interpretable modes of variation (Hsu et al., 2014).
- In scientific applications such as protein folding, each original co-ordinate axis has a physical interpretation, and the reduced set of co-ordinate axes should too.
- In finance applications such as investing capital across index funds, each non-zero entry in each eigenvector used to reduce the feature space incurs a transaction cost.
- If $p \gg n$, PCA suffers from a curse of dimensionality and becomes physically meaningless (Amini and Wainwright, 2008).

One common method for obtaining interpretable principal components is to stipulate that they are sparse, i.e., maximize variance while containing at most k non-zero entries. This approach leads to the following non-convex mixed-integer quadratically constrained problem (see d’Aspremont et al., 2005):

$$\max_{\mathbf{x} \in \mathbb{R}^p} \mathbf{x}^\top \Sigma \mathbf{x} \quad \text{s.t.} \quad \mathbf{x}^\top \mathbf{x} = 1, \|\mathbf{x}\|_0 \leq k, \quad (1)$$

where the constraint $\|\mathbf{x}\|_0 \leq k$ forces variance to be explained in a compelling fashion.

1.1 Background and Literature Review

Owing to sparse PCA’s fundamental importance in a variety of applications including best subset selection (d’Aspremont et al., 2008), natural language processing (Zhang et al., 2012), compressed sensing (Candes and Tao, 2007), and clustering (Luss and d’Aspremont, 2010), three distinct classes of methods for addressing Problem (1) have arisen. Namely, (a) heuristic methods which obtain high-quality sparse PCs in an efficient fashion but do not supply guarantees on the quality of the solution, (b) convex relaxations which obtain certifiably near-optimal solutions by solving a convex relaxation and rounding, and (c) exact methods which obtain certifiably optimal solutions, albeit in exponential time.

Heuristic Approaches: The importance of identifying a small number of interpretable principal components has been well-documented in the literature since the work of Hotelling (1933) (see also Jeffers, 1967), giving rise to many distinct heuristic approaches for obtaining high-quality solutions to Problem (1). Two interesting such approaches are to rotate dense principal components to promote sparsity (Kaiser, 1958; Richman, 1986; Jolliffe, 1995), or

1. A second criticism of PCA is that, as set up here, it uses the sample correlation or covariance matrix. This is a drawback, because sample covariance matrices are poorly conditioned estimators which over-disperses the sample eigenvalues, particularly in high-dimensional settings. In practice, this can be rectified by, e.g., using a shrinkage estimator (see, e.g., Ledoit and Wolf, 2004). We do not do so here for simplicity, but we recommend doing so if using the techniques developed in this paper in practice.

apply an ℓ_1 penalty term as a convex surrogate to the cardinality constraint (Jolliffe et al., 2003; Zou et al., 2006). Unfortunately, the former approach does not provide performance guarantees, while the latter approach still results in a non-convex optimization problem.

More recently, motivated by the need to rapidly obtain high-quality sparse principal components at scale, a wide variety of first-order heuristic methods have emerged. The first such *modern* heuristic was developed by Journée et al. (2010), and involves combining the power method with thresholding and re-normalization steps. By pursuing similar ideas, several related methods have since been developed (see Witten et al., 2009; Hein and Bühler, 2010; Richtárik et al., 2020; Luss and Teboulle, 2013; Yuan and Zhang, 2013). Unfortunately, while these methods are often very effective in practice, they sometimes badly fail to recover an optimal sparse principal component, and a practitioner using a heuristic method typically has no way of knowing when this has occurred. Indeed, Berk and Bertsimas (2019) recently compared 7 heuristic methods, including most of those reviewed here, on 14 instances of sparse PCA, and found that none of the heuristic methods successfully recovered an optimal solution in all 14 cases (i.e., no heuristic was right all the time).

Convex Relaxations: Motivated by the shortcomings of heuristic approaches on high-dimensional datasets, and the successful application of semi-definite optimization in obtaining high-quality approximation bounds in other applications (see Goemans and Williamson, 1995; Wolkowicz et al., 2012), a variety of convex relaxations have been proposed for sparse PCA. The first such convex relaxation was proposed by d’Aspremont et al. (2005), who reformulated sparse PCA as the rank-constrained mixed-integer semidefinite optimization problem (MISDO):

$$\max_{\mathbf{X} \succeq \mathbf{0}} \langle \boldsymbol{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, \|\mathbf{X}\|_0 \leq k^2, \text{Rank}(\mathbf{X}) = 1, \quad (2)$$

where \mathbf{X} models the outer product $\mathbf{x}\mathbf{x}^\top$. Note that, for a rank-one matrix \mathbf{X} , the constraint $\|\mathbf{X}\|_0 \leq k^2$ in (2) is equivalent to the constraint $\|\mathbf{x}\|_0 \leq k$ in (1), since a vector \mathbf{x} is k -sparse if its outer product $\mathbf{x}\mathbf{x}^\top$ is k^2 -sparse. After performing this reformulation, d’Aspremont et al. (2005) relaxed both the cardinality and rank constraints and instead solved

$$\max_{\mathbf{X} \succeq \mathbf{0}} \langle \boldsymbol{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, \|\mathbf{X}\|_1 \leq k, \quad (3)$$

which supplies a valid upper bound on Problem (1)’s objective.

The semidefinite approach has since been refined in a number of follow-up works. Among others, d’Aspremont et al. (2008), building upon the work of Ben-Tal and Nemirovski (2002), proposed a different semidefinite relaxation which supplies a sufficient condition for optimality via the primal-dual KKT conditions, and d’Aspremont et al. (2014) analyzed the quality of the semidefinite relaxation in order to obtain high-quality approximation bounds. A common theme in these approaches is that they require solving large-scale semidefinite optimization problems. This presents difficulties for practitioners because state-of-the-art

implementations of interior point methods such as `Mosek` require $O(p^6)$ memory to solve Problem (3), and therefore currently cannot solve instances of Problem (3) with $p \geq 300$ (see Bertsimas and Cory-Wright, 2020, for a recent comparison). Techniques other than interior point methods, e.g., ADMM or augmented Lagrangian methods as reviewed in Majumdar et al. (2019) could also be used to solve Problem (3), although they tend to require more runtime than IPMs to obtain a solution of a similar accuracy and be numerically unstable for problem sizes where IPMs run out of memory (Majumdar et al., 2019).

A number of works have also studied the statistical estimation properties of Problem (3), by assuming an underlying probabilistic model. Among others, Amini and Wainwright (2008) have demonstrated the asymptotic consistency of Problem (3) under a spiked covariance model once the number of samples used to generate the covariance matrix exceeds a certain threshold; see Vu and Lei (2012); Berthet and Rigollet (2013); Wang et al. (2016) for further results in this direction, Miolane (2018) for a recent survey.

In an complementary direction, Dey et al. (2018) has recently questioned the modeling paradigm of lifting \mathbf{x} to a higher dimensional space by instead considering the following (tighter) relaxation of sparse PCA in the original problem space

$$\max_{\mathbf{x} \in \mathbb{R}^p} \mathbf{x}^\top \Sigma \mathbf{x} \quad \text{s.t.} \quad \|\mathbf{x}\|_2 = 1, \|\mathbf{x}\|_1 \leq \sqrt{k}. \quad (4)$$

Interestingly, Problem (4)'s relaxation provides a $\left(1 + \sqrt{\frac{k}{k+1}}\right)^2$ -factor bound approximation of Problem (1)'s objective, while Problem (3)'s upper bound may be exponentially larger in the worst case (Amini and Wainwright, 2008). This additional tightness, however, comes at a price: Problem (4) is NP-hard to solve—indeed, providing a constant-factor guarantee on sparse PCA is NP-hard (Magdon-Ismail, 2017)—and thus (4) is best formulated as a MIO, while Problem (3) can be solved in polynomial time.

More recently, by building on the work of Kim and Kojima (2001), Bertsimas and Cory-Wright (2020) introduced a second-order cone relaxation of (2) which scales to $p = 1000s$, and matches the semidefinite bound after imposing a small number of cuts. Moreover, it typically supplies bound gaps of less than 5%. However, it does not supply an *exact* certificate of optimality, which is often desirable, for instance in medical applications.

A fundamental drawback of existing convex relaxation techniques is that they are not coupled with rounding schemes for obtaining high-quality feasible solutions. This is problematic, because optimizers are typically interested in obtaining high-quality solutions, rather than certificates. In this paper, we take a step in this direction, by deriving new convex relaxations that naturally give rise to greedy and random rounding schemes. The fundamental point of difference between our relaxations and existing relaxations is that we derive our relaxations by rewriting sparse PCA as a MISDO and dropping an integrality constraint, rather than using more ad-hoc techniques.

Exact Methods: Motivated by the successful application of mixed-integer optimization for solving statistical learning problems such as best subset selection (Bertsimas and Van Parys, 2020) and sparse classification (Bertsimas et al., 2017), several exact methods for solving sparse PCA to certifiable optimality have been proposed. The first branch-and-bound algorithm for solving Problem (1) was proposed by Moghaddam et al. (2006), by applying norm equivalence relations to obtain valid bounds. However, Moghaddam et al. (2006) did not couple their approach with high-quality initial solutions and tractable bounds to prune partial solutions. Consequently, they could not scale their approach beyond $p = 40$.

A more sophisticated branch-and-bound scheme was recently proposed by Berk and Bertsimas (2019), which couples tighter Gershgorin Circle Theorem bounds (Horn and Johnson, 1990, Chapter 6) with a fast heuristic due to Yuan and Zhang (2013) to solve problems up to $p = 250$. However, their method cannot scale beyond $p = 100$ s, because the bounds obtained are too weak to avoid enumerating a sizeable portion of the tree.

Recently, the authors developed a framework for reformulating convex mixed-integer optimization problems with logical constraints (see Bertsimas et al., 2019), and demonstrated that this framework allows a number of problems of practical relevance to be solved to certifiably optimality via a cutting-plane method. In this paper, we build upon this work by reformulating Problem (1) as a *convex* mixed-integer semidefinite optimization problem, and leverage this reformulation to design a cutting-plane method which solves sparse PCA to certifiable optimality. A key feature of our approach is that we need not solve any semidefinite subproblems. Rather, we use concepts from SDO to design a semidefinite-free approach which uses simple linear algebra techniques.

Concurrently to our initial submission, Li and Xie (2020) also attempted to reformulate sparse PCA as an MISDO, and proposed valid inequalities for strengthening their formulation and local search algorithms for obtaining high-quality solutions at scale. Our work differs in the following two ways. First, we propose strengthening the MISDO formulation using the Gershgorin circle theorem and demonstrate that this allows our MISDO formulation to scale to problems with $p = 100$ s of features, while they do not, to our knowledge, solve any MISDOs to certifiable optimality where $p > 13$. Second, we develop tractable second-order cone relaxations and greedy rounding schemes which allow practitioners to obtain certifiably near optimal sparse principal components even in the presence of $p = 1,000$ s of features. More remarkable than the differences between the works however is the similarities: more than 15 years after d’Aspremont et al. (2005)’s landmark paper first appeared, both works proposed reformulating sparse PCA as an MISDO less than a week apart. In our view, this demonstrates that the ideas contained in both works transcend sparse PCA, and can perhaps be applied to other problems in the optimization literature which have not yet been formulated as MISDOs.

1.2 Contributions and Structure

The main contributions of the paper are twofold. First, we reformulate sparse PCA exactly as a mixed-integer semidefinite optimization problem; a reformulation which is, to the best of our knowledge, novel. Second, we leverage this MISDO formulation to design efficient algorithms for solving non-convex mixed-integer quadratic optimization problems, such as sparse PCA, to certifiable optimality or within 1 – 2% of optimality in practice at a larger scale than existing state-of-the-art methods. The structure and detailed contributions of the paper are as follows:

- In Section 2.1, we reformulate Problem (1) as a mixed-integer SDO. We propose a cutting-plane method which solves it to certifiable optimality in Section 2.2. Our algorithm decomposes the problem into a purely binary master problem and a semidefinite separation problem. Interestingly, we show in Section 2.3 that the separation problems can be solved efficiently via a leading eigenvalue computation and does not require any SDO solver. Finally, the Gershgorin Circle theorem has been empirically successful for deriving upper-bounds on the objective value of (1) (Berk and Bertsimas, 2019). We theoretically analyze the quality of such bounds in Section 2.4 and show in Section 2.5 that tighter bounds derived from Brauer’s ovals of Cassini theorem can also be imposed via mixed-integer second-order cone constraints.
- In Section 3, we analyze the semidefinite reformulation’s convex relaxation, and introduce a greedy rounding scheme (Section 3.1) which supplies high-quality solutions to Problem (1) in polynomial time, together with a sub-optimality gap (see numerical experiments in Section 4). To further improve the quality of rounded solution and the optimality gap, we introduce strengthening inequalities (Section 3.2). While solving the strengthened formulation exactly would result in an intractable MISDO problem, solving its relaxation and rounding the solution appears as an efficient strategy to return high-quality solutions with a numerical certificate of near-optimality.
- In Section 4, we apply the cutting-plane and random rounding methods to derive optimal and near optimal sparse principal components for problems in the UCI dataset. We also compare our method’s performance against the method of Berk and Bertsimas (2019), and find that our exact cutting-plane method performs comparably, while our relax+round approach successfully scales to problems an order of magnitude larger and often returns solutions which outperform the exact method at sizes which the exact method cannot currently scale to. A key feature of our numerical success is that we sidestep the computational difficulties in solving SDOs at scale by proposing semidefinite-free methods for solving the convex relaxations, i.e., solving second-order cone relaxations.

Notation: We let nonbold face characters such as b denote scalars, lowercase bold faced characters such as \mathbf{x} denote vectors, uppercase bold faced characters such as \mathbf{X} denote

matrices, and calligraphic uppercase characters such as \mathcal{Z} denote sets. We let $[p]$ denote the set of running indices $\{1, \dots, p\}$. We let \mathbf{e} denote a vector of all 1's, $\mathbf{0}$ denote a vector of all 0's, and \mathbb{I} denote the identity matrix, with dimension implied by the context.

We also use an assortment of matrix operators. We let $\langle \cdot, \cdot \rangle$ denote the Euclidean inner product between two matrices, $\| \cdot \|_F$ denote the Frobenius norm of a matrix, $\| \cdot \|_\sigma$ denote the spectral norm of a matrix, $\| \cdot \|_*$ denote the nuclear norm of a matrix, \mathbf{X}^\dagger denote the Moore-Penrose pseudoinverse of a matrix \mathbf{X} and S_+^p denote the $p \times p$ positive semidefinite cone; see Horn and Johnson (1990) for a general theory of matrix operators.

2. An Exact Mixed-Integer Semidefinite Optimization Algorithm

In Section 2.1, we reformulate Problem (1) as a convex mixed-integer semidefinite optimization problem. From this formulation, we propose an outer-approximation scheme (Section 2.2) which, as we show in Section 2.3, does not require solving any semidefinite problems. We improve convergence of the algorithm by deriving quality upper-bounds on Problem's (1) objective value in Section 2.4 and 2.5.

2.1 A Mixed-Integer Semidefinite Reformulation

Starting from the rank-constrained SDO formulation (2), we introduce binary variables z_i to model whether $X_{i,j}$ is non-zero, via the logical constraint $X_{i,j} = 0$ if $z_i = 0$; note that we need not require that $X_{i,j} = 0$ if $z_j = 0$, since \mathbf{X} is a symmetric matrix. By enforcing the logical constraint via $-M_{i,j}z_i \leq X_{i,j} \leq M_{i,j}z_i$ for sufficiently large $M_{i,j} > 0$, Problem (2) becomes

$$\begin{aligned} \max_{z \in \{0,1\}^p: \mathbf{e}^\top z \leq k} \quad & \max_{\mathbf{X} \in S_+^p} \quad \langle \Sigma, \mathbf{X} \rangle \\ \text{s.t.} \quad & \text{tr}(\mathbf{X}) = 1, \quad -M_{i,j}z_i \leq X_{i,j} \leq M_{i,j}z_i \quad \forall i, j \in [p], \quad \text{Rank}(\mathbf{X}) = 1. \end{aligned}$$

To obtain a MISDO reformulation, we omit the rank constraint. In general, omitting a rank constraint generates a relaxation and induces some loss of optimality. Remarkably, this omission is without loss of optimality in this case. Indeed, the objective is convex and therefore some rank-one extreme matrices \mathbf{X} is optimal. We formalize this observation in the following theorem; note that a similar result—although in the context of computing Restricted Isometry constants and with a different proof—exists (Gally and Pfetsch, 2016):

Theorem 1 *Problem (1) attains the same optimal objective value as the problem:*

$$\begin{aligned}
 & \max_{z \in \{0,1\}^p: e^\top z \leq k} \max_{\mathbf{X} \in S_+^p} \langle \boldsymbol{\Sigma}, \mathbf{X} \rangle \\
 & \text{s.t.} \quad \text{tr}(\mathbf{X}) = 1 \\
 & \quad |X_{i,j}| \leq M_{i,j} z_i \quad \forall i, j \in [p], \\
 & \quad \sum_{j=1}^p |X_{i,j}| \leq \sqrt{k} z_i \quad \forall i \in [p],
 \end{aligned} \tag{5}$$

where $M_{i,i} = 1$, and $M_{i,j} = \frac{1}{2}$ if $j \neq i$.

Remark 2 *Observe that if $k \leq \sqrt{n}$ and we set $M_{i,j} = 1 \forall i, j \in [p]$ in Problem (5) and omit the valid inequality $\sum_{j=1}^p |X_{i,j}| \leq \sqrt{k} z_i$ then the optimal value of the continuous relaxation is trivially $\lambda_{\max}(\boldsymbol{\Sigma})$. Indeed, letting \mathbf{x} be a leading eigenvector of the unconstrained problem (where $\|\mathbf{x}\|_2 = 1$), we can set $z_i = |x_i| \geq |x_i| |x_j|$, where the inequality holds since $\|\mathbf{x}\|_2 = 1$, and $X_{i,j} = x_i x_j$, meaning (a) $\sum_i z_i = \|\mathbf{x}\|_1 \leq k \leq \sqrt{n}$ by norm equivalence and (b) $|X_{i,j}| \leq z_i$ and thus (\mathbf{X}, \mathbf{z}) solves this continuous relaxation. Therefore, setting $M_{i,j} = \frac{1}{2}$ if $j \neq i$ and/or imposing the valid inequality $\sum_{j=1}^p |X_{i,j}| \leq \sqrt{k} z_i$ is necessary for obtaining non-trivial relaxations whenever k is small.*

Proof It suffices to demonstrate that for any feasible solution to (1) we can construct a feasible solution to (5) with an equal or greater payoff, and vice versa.

- Let $\mathbf{x} \in \mathbb{R}^p$ be a feasible solution to (1). Then, since $\|\mathbf{x}\|_1 \leq \sqrt{k}$, $(\mathbf{X} := \mathbf{x}\mathbf{x}^\top, \mathbf{z})$ is a feasible solution to (5) with equal cost, where $z_i = 1$ if $|x_i| > 0$, $z_i = 0$ otherwise.
- Let (\mathbf{X}, \mathbf{z}) be a feasible solution to Problem (5), and let $\mathbf{X} = \sum_{i=1}^p \sigma_i \mathbf{x}_i \mathbf{x}_i^\top$ be a Cholesky decomposition of \mathbf{X} , where $e^\top \boldsymbol{\sigma} = 1, \boldsymbol{\sigma} \geq \mathbf{0}$, and $\|\mathbf{x}_i\|_2 = 1 \forall i \in [p]$. Observe that $\|\mathbf{x}_i\|_0 \leq k \forall i \in [p]$, since we can perform the Cholesky decomposition on the submatrix of \mathbf{X} induced by \mathbf{z} , and “pad” out the remaining entries of each \mathbf{x}_i with 0s to obtain the decomposition of \mathbf{X} . Therefore, let us set $\hat{\mathbf{x}} := \arg \max_i [\mathbf{x}_i^\top \boldsymbol{\Sigma} \mathbf{x}_i]$. Then, $\hat{\mathbf{x}}$ is a feasible solution to (1) with an equal or greater payoff.

Finally, we let $M_{i,i} = 1, M_{i,j} = \frac{1}{2}$ if $i \neq j$, as the 2×2 minors imply $X_{i,j}^2 \leq X_{i,i} X_{j,j} \leq \frac{1}{4}$ whenever $i \neq j$ (c.f. Gally and Pfetsch, 2016, Lemma 1). ■

Theorem 1 reformulates Problem (1) as a mixed-integer SDO. Therefore, we can solve Problem (5) using general branch-and-cut techniques for semidefinite optimization problems (see Gally et al., 2018; Kobayashi and Takano, 2020). However, this approach is not scalable, as it comprises solving a large number of semidefinite subproblems and the community does not know how to efficiently warm-start interior point methods (IPMs) for SDOs.

Alternatively, we propose a saddle-point reformulation of Problem (5) which avoids the computational difficulty of solving a large number of SDOs by exploiting problem structure,

as we will show in Section 2.3. The following result reformulates Problem (5) as a max-min saddle-point problem amenable to outer-approximation:

Theorem 3 *Problem (5) attains the same optimal value as the following problem:*

$$\begin{aligned} & \max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z}) \tag{6} \\ \text{where } f(\mathbf{z}) := & \min_{\lambda \in \mathbb{R}, \boldsymbol{\alpha} \in \mathbb{R}^{p \times p}, \boldsymbol{\beta} \in \mathbb{R}^p} \lambda + \sum_{i=1}^p z_i \left(\sum_{j=1}^p M_{i,j} \max(0, |\alpha_{i,j}| - \beta_i) + \sqrt{k} \beta_i \right) \tag{7} \\ & \text{s.t. } \lambda \mathbb{I} + \boldsymbol{\alpha} \succeq \boldsymbol{\Sigma}. \end{aligned}$$

Remark 4 *The above theorem demonstrates that $f(\mathbf{z})$ is concave in \mathbf{z} , by rewriting it as the infimum of functions which are linear in \mathbf{z} (Boyd and Vandenberghe, 2004).*

Proof Let us introduce auxiliary variables $U_{i,j}$ to model the absolute value of $X_{i,j}$ and rewrite the inner optimization problem of (5) as

$$\begin{aligned} f(\mathbf{z}) := & \max_{\mathbf{X} \succeq \mathbf{0}, \mathbf{U}} \langle \boldsymbol{\Sigma}, \mathbf{X} \rangle \\ \text{s.t. } & \text{tr}(\mathbf{X}) = 1, \quad [\lambda] \\ & U_{i,j} \leq M_{i,j} z_i \quad \forall i, j \in [p], \quad [\sigma_{i,j}] \\ & |X_{i,j}| \leq U_{i,j} \quad \forall i, j \in [p], \quad [\alpha_{i,j}] \\ & \sum_{j=1}^p U_{i,j} \leq \sqrt{k} z_i \quad \forall i \in [p], \quad [\beta_i] \end{aligned} \tag{8}$$

where we associate dual constraint multipliers with primal constraints in square brackets. For \mathbf{z} such that $\mathbf{e}^\top \mathbf{z} \geq 1$, the maximization problem induced by $f(\mathbf{z})$ satisfies Slater's condition (see, e.g., Boyd and Vandenberghe, 2004, Chapter 5.2.3), strong duality applies and leads to

$$\begin{aligned} f(\mathbf{z}) = & \min_{\lambda, \boldsymbol{\sigma}, \boldsymbol{\alpha}, \boldsymbol{\beta} \geq \mathbf{0}} \lambda + \sum_{i,j} \sigma_{i,j} M_{i,j} z_i + \sum_{i=1}^p \beta_i \sqrt{k} z_i \\ \text{s.t. } & \lambda \mathbb{I} + \boldsymbol{\alpha} \succeq \boldsymbol{\Sigma}, |\alpha_{i,j}| \leq \sigma_{i,j} + \beta_i. \end{aligned}$$

We eliminate $\boldsymbol{\sigma}$ from the dual problem above by optimizing over $\sigma_{i,j}$ and setting $\sigma_{i,j}^* = \max(0, |\alpha_{i,j}| - \beta_i)$.

Note that for $\mathbf{z} = \mathbf{0}$, the primal subproblem is infeasible and the dual subproblem has objective $-\infty$, but this can safely be ignored since $\mathbf{z} = \mathbf{0}$ is certainly suboptimal. \blacksquare

2.2 A Cutting-Plane Method

Theorem 3 shows that evaluating $f(\hat{\mathbf{z}})$ yields the globally valid overestimator:

$$f(\mathbf{z}) \leq f(\hat{\mathbf{z}}) + \mathbf{g}_{\hat{\mathbf{z}}}^\top (\mathbf{z} - \hat{\mathbf{z}}),$$

where $\mathbf{g}_{\hat{\mathbf{z}}}$ is a supergradient of f at $\hat{\mathbf{z}}$, at no additional cost. In particular, we have

$$\mathbf{g}_{\hat{\mathbf{z}},i} = \left(\sum_{j=1}^p M_{i,j} \max(0, |\alpha_{i,j}^*(\hat{\mathbf{z}})| - \beta_i(\hat{\mathbf{z}})) + \sqrt{k} \beta_i(\hat{\mathbf{z}}) \right),$$

where $\boldsymbol{\alpha}^*(\hat{\mathbf{z}})$, $\boldsymbol{\beta}^*(\hat{\mathbf{z}})$ constitutes an optimal choice of $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ for a fixed $\hat{\mathbf{z}}$. This observation leads to an efficient strategy for maximizing $f(\mathbf{z})$: iteratively maximizing and refining a piecewise linear upper estimator of $f(\mathbf{z})$. This strategy is called outer-approximation (OA), and was originally proposed by Duran and Grossmann (1986). OA works by iteratively constructing estimators of the following form at each iteration t :

$$f^t(\mathbf{z}) = \min_{1 \leq i \leq t} \left\{ f(\mathbf{z}_i) + \mathbf{g}_{\mathbf{z}_i}^\top (\mathbf{z} - \mathbf{z}_i) \right\}. \quad (9)$$

After constructing each overestimator, we maximize $f^t(\mathbf{z})$ over $\{0,1\}^p$ to obtain \mathbf{z}_t , and evaluate $f(\cdot)$ and its supergradient at \mathbf{z}_t . This procedure yields a non-increasing sequence of overestimators $\{f^t(\mathbf{z}_t)\}_{t=1}^T$ which converge to the optimal value of $f(\mathbf{z})$ within a finite number of iterations $T \leq \binom{p}{1} + \dots + \binom{p}{k}$, since $\{0,1\}^p$ is a finite set and OA never visits a point twice. Additionally, we can avoid solving a different MILO at each OA iteration by integrating the entire algorithm within a single branch-and-bound tree, as proposed by Quesada and Grossmann (1992), using `lazy constraint callbacks`. Lazy constraint callbacks are now standard components of modern MILO solvers such as `Gurobi` or `CPLEX` and substantially speed-up OA. We formalize this procedure in Algorithm 1; note that $\partial f(\mathbf{z}_{t+1})$ denotes the set of supergradients of f at \mathbf{z}_{t+1} .

Algorithm 1 An outer-approximation method for Problem (1)

Require: Initial solution \mathbf{z}_1

$t \leftarrow 1$

repeat

 Compute $\mathbf{z}_{t+1}, \theta_{t+1}$ solution of

$$\max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k, \theta} \theta \quad \text{s.t. } \theta \leq f(\mathbf{z}_i) + \mathbf{g}_{\mathbf{z}_i}^\top (\mathbf{z} - \mathbf{z}_i) \quad \forall i \in [t],$$

 Compute $f(\mathbf{z}_{t+1})$ and $\mathbf{g}_{\mathbf{z}_{t+1}} \in \partial f(\mathbf{z}_{t+1})$ by solving (7)

$t \leftarrow t + 1$

until $f(\mathbf{z}_t) - \theta_t \leq \varepsilon$

return \mathbf{z}_t

2.3 A Semidefinite-free Subproblem Strategy

Our derivation and analysis of Algorithm 1 indicates that we can solve Problem (1) to certifiable optimality by solving a (potentially large) number of semidefinite subproblems (7), which might be prohibitive in practice. Therefore, we now derive a computationally efficient subproblem strategy which crucially does not require solving *any* semidefinite programs. Formally, we have the following result:

Theorem 5 *For any $\mathbf{z} \in \{0, 1\}^p : \mathbf{e}^\top \mathbf{z} \leq k$, optimal dual variables in (7) are*

$$\lambda = \lambda_{\max}(\boldsymbol{\Sigma}_{1,1}), \quad \hat{\boldsymbol{\alpha}} = \begin{pmatrix} \hat{\boldsymbol{\alpha}}_{1,1} & \hat{\boldsymbol{\alpha}}_{1,2} \\ \hat{\boldsymbol{\alpha}}_{1,2}^\top & \hat{\boldsymbol{\alpha}}_{2,2} \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{2,2} - \lambda \mathbb{I} + \boldsymbol{\Sigma}_{1,2}^\top (\lambda \mathbb{I} - \boldsymbol{\Sigma}_{1,1})^\dagger \boldsymbol{\Sigma}_{1,2} \end{pmatrix}, \quad (10)$$

$$\beta_i = (1 - z_i) \left(|\hat{\alpha}_{i,1}|, |\hat{\alpha}_{i,2}|, \dots, |\hat{\alpha}_{i,i}|, |\hat{\alpha}_{i,i}|, \dots, |\hat{\alpha}_{i,p}| \right)_{[\lceil \sqrt{k} \rceil]} \quad \forall i \in [p], \quad (11)$$

where $\lambda_{\max}(\cdot)$ denotes the leading eigenvalue of a matrix, $\hat{\boldsymbol{\alpha}} = \begin{pmatrix} \hat{\boldsymbol{\alpha}}_{1,1} & \hat{\boldsymbol{\alpha}}_{1,2} \\ \hat{\boldsymbol{\alpha}}_{1,2}^\top & \hat{\boldsymbol{\alpha}}_{2,2} \end{pmatrix}$ is a permutation such that $\hat{\boldsymbol{\alpha}}_{1,1}$ (resp. $\hat{\boldsymbol{\alpha}}_{2,2}$) denotes the entries of $\hat{\boldsymbol{\alpha}}$ where $z_i = z_j = 1$ ($z_i = z_j = 0$); $\boldsymbol{\Sigma}$ is similar, and $(\mathbf{x})_{[k]}$ denotes the k th largest element of \mathbf{x} .

Remark 6 *By Theorem 5, Problem (7) can be solved by computing the leading eigenvalue of $\boldsymbol{\Sigma}_{1,1}$ and solving a linear system. This justifies our claim that we need not solve any SDOs in our algorithmic strategy.*

Proof We appeal to strong duality and complementary slackness. Observe that, for any $\mathbf{z} \in \{0, 1\}^p$, $f(\mathbf{z})$ is the optimal value of a maximization problem over a closed convex compact set. Therefore, there exists some optimal primal solution \mathbf{X}^* without loss of generality. Moreover, since the primal has non-empty relative interior with respect to the non-affine constraints, it satisfies the Slater constraint qualification and strong duality holds (see, e.g., Boyd and Vandenberghe, 2004, Chapter 5.2.3). Therefore, by complementary slackness (see, e.g., Boyd and Vandenberghe, 2004, Chapter 5.5.2), there must exist some dual-optimal solution $(\lambda, \hat{\boldsymbol{\alpha}}, \boldsymbol{\beta})$ which obeys complementarity with \mathbf{X}^* . Moreover, $|X_{i,j}| \leq M_{i,j}$ is implied by $\text{tr}(\mathbf{X}) = 1, \mathbf{X} \succeq \mathbf{0}$, while $\sum_{j=1}^p |X_{i,j}| \leq z_i \sqrt{k}$ is implied by $|X_{i,j}| \leq M_{i,j} z_i$ and $\mathbf{e}^\top \mathbf{z} \leq k$. Therefore, by complementary slackness, we can take the constraints $|X_{i,j}| \leq M_{i,j} z_i, \sum_{j=1}^p |X_{i,j}| \leq z_i \sqrt{k}$ to be inactive when $z_i = 1$ without loss of generality, which implies that $\hat{\alpha}_{i,j}^*, \beta_i^* = 0$ if $z_i = 1$ in some dual-optimal solution. Moreover, we also have $\hat{\alpha}_{i,j}^* = 0$ if $z_j = 1$, since $\hat{\boldsymbol{\alpha}}$ obeys the dual feasibility constraint $\lambda \mathbb{I} + \hat{\boldsymbol{\alpha}} \succeq \boldsymbol{\Sigma}$, and therefore is itself symmetric.

Next, observe that, by strong duality, $\lambda = \lambda_{\max}(\boldsymbol{\Sigma}_{1,1})$ in this dual-optimal solution, since $\boldsymbol{\alpha}$ only takes non-zero values if $z_i = z_j = 0$ and does not contribute to the objective, and $\boldsymbol{\beta}$ is similar.

Next, observe that, by strong duality and complementary slackness, any dual feasible $(\lambda, \hat{\alpha}, \beta)$ satisfying the above conditions is dual-optimal. Therefore, we need to find an $\hat{\alpha}_{2,2}$ such that

$$\begin{pmatrix} \lambda \mathbb{I} - \Sigma_{1,1} & -\Sigma_{1,2} \\ -\Sigma_{2,1} & \lambda \mathbb{I} + \hat{\alpha}_{2,2} - \Sigma_{2,2} \end{pmatrix} \succeq \mathbf{0}.$$

By the generalized Schur complement lemma (see Boyd et al., 1994, Equation 2.41), this is PSD if and only if

1. $\lambda \mathbb{I} - \Sigma_{1,1} \succeq \mathbf{0}$,
2. $(\mathbb{I} - (\lambda \mathbb{I} - \Sigma_{1,1})(\lambda \mathbb{I} - \Sigma_{1,1})^\dagger) \Sigma_{1,2} = \mathbf{0}$, and
3. $\lambda \mathbb{I} + \hat{\alpha}_{2,2} - \Sigma_{2,2} \succeq \Sigma_{1,2}^\top (\lambda \mathbb{I} - \Sigma_{1,1})^\dagger \Sigma_{1,2}$.

The first two conditions hold because, as argued above, λ is optimal and therefore feasible, and the conditions are independent of $\hat{\alpha}_{2,2}$. Therefore, it suffices to pick $\hat{\alpha}_{2,2}$ in order that the third condition holds. We achieve this by setting $\hat{\alpha}_{2,2}$ so the PSD constraint in condition (3) holds with equality.

Finally, let us optimize for β to obtain stronger cuts (when $z_i = 0$ we can pick any feasible β_i , but optimizing to set $\partial f(\mathbf{z})_i$ to be as small as possible gives stronger cuts). This is equivalent to solving the univariate minimization problem for each β_i :

$$\min_{\beta_i} \left(\sum_{j=1}^p M_{i,j} \max(0, |\alpha_{i,j}| - \beta_i) + \sqrt{k} \beta_i \right).$$

Moreover, it is a standard result from max- k optimization (see, e.g., Zakeri et al., 2014; Todd, 2018) that this is achieved by setting β_i to be the $\lceil \sqrt{k} \rceil$ largest element of $\{\alpha_{i,j}\}_{j \in [p]} \cup \{\alpha_{i,i}\}$ in absolute magnitude, where we include $\alpha_{i,i}$ twice since $M_{i,i} = 1$ while $M_{i,j} = 1/2$ if $j \neq i$. ■

2.4 Strengthening the Master Problem via the Gershgorin Circle Theorem

To accelerate Algorithm 1, we strengthen the master problem by imposing bounds from the circle theorem. Formally, we have the following result, which can be deduced from (Horn and Johnson, 1990, Theorem 6.1.1):

Theorem 7 *For any vector $\mathbf{z} \in \{0, 1\}^p$ we have the following upper bound on $f(\mathbf{z})$*

$$f(\mathbf{z}) \leq \max_{j \in [p]: z_j = 1} \sum_{i \in [p]} z_i |\Sigma_{i,j}|. \tag{12}$$

Observe that this bound cannot be used to *directly* strengthen Algorithm 1’s master problem, since the bound is not convex in \mathbf{z} . Nonetheless, it can be successfully applied if we (a) impose a big-M assumption on Problem (1)’s optimal objective and (b) introduce p additional binary variables $\mathbf{s} \in \{0, 1\}^p : \mathbf{e}^\top \mathbf{s} = 1$ which model whether the i th Gershgorin disc is active; recall that each eigenvalue is contained in the union of the discs. Formally, we impose the following valid inequalities in the master problem:

$$\exists \mathbf{s} \in \{0, 1\}^p : \theta \leq \sum_{i \in [p]} z_i |\Sigma_{i,j}| + M(1 - s_j) \quad \forall j \in [p], \mathbf{e}^\top \mathbf{s} = 1, \mathbf{s} \leq \mathbf{z}, \quad (13)$$

where θ is the epigraph variable maximized in the master problem stated in Algorithm 1, and M is an upper bound on the sum of the k largest absolute entries in any column of Σ . Note that we set $\mathbf{s} \leq \mathbf{z}$ since if $z_i = 0$ the i th column of Σ does not feature in the relevant submatrix of Σ . In the above inequalities, a valid M is given by any bound on the optimal objective. Since Theorem 7 supplies one such bound for any given \mathbf{z} , we can compute

$$M := \max_{j \in [p]} \max_{\mathbf{z} \in \{0, 1\}^p : \mathbf{e}^\top \mathbf{z} \leq k} \sum_{i \in [p]} z_i |\Sigma_{i,j}|, \quad (14)$$

which can be done in $O(p^2)$ time.

To further improve Algorithm 1, we also make use of the Gershgorin circle theorem before generating each outer-approximation cut. Namely, at a given node in a branch-and-bound tree, there are indices i where z_i has been fixed to 1, indices i where z_i has been fixed to 0, and indices i where z_i has not yet been fixed. Accordingly, we compute the worst-case Gershgorin bound—by taking the worst-case bound over each index j such that z_j has not yet been fixed to 0, i.e.,

$$\max_{j: z_j \neq 0} \left\{ \max_{\mathbf{s} \in \{0, 1\}^p : \mathbf{e}^\top \mathbf{s} \leq k} \left\{ \sum_{i \in [p]} s_i |\Sigma_{i,j}| \text{ s.t. } s_i = 0 \text{ if } z_i = 0, s_i = 1 \text{ if } z_i = 1 \right\} \right\}.$$

If this bound is larger than our incumbent solution then we generate an outer-approximation cut, otherwise the entire subtree rooted at this node does not contain an optimal solution and we use instruct the solver to avoid exploring this node via a `callback`.

Our numerical results in Section 4 echo the empirical findings of Berk and Bertsimas (2019) and indicate that Algorithm 1 performs substantially better when the Gershgorin bound is supplied in the master problem. Therefore, it is interesting to theoretically investigate the tightness, or at least the quality, of Gershgorin’s bound. We supply some results in this direction in the following proposition:

Proposition 8 *Suppose that Σ is a scaled diagonally dominant matrix as defined by Boman et al. (2005), i.e., there exists some vector $\mathbf{d} > 0$ such that*

$$d_i \Sigma_{i,i} \geq \sum_{j \in [p]: j \neq i} d_j |\Sigma_{i,j}| \quad \forall i \in [p].$$

Then, letting $\rho := \max_{i,j \in [p]} \{\frac{d_i}{d_j}\}$, the Gershgorin circle theorem provides a $(1 + \rho)$ -factor approximation, i.e.,

$$f(\mathbf{z}) \leq \max_{j \in [p]} \left\{ \sum_{i \in [p]} z_i |\Sigma_{i,j}| \right\} \leq (1 + \rho) f(\mathbf{z}) \quad \forall \mathbf{z} \in \{0, 1\}^p. \quad (15)$$

Remark 9 Observe that, for a fixed \mathbf{z} , the ratio $\rho := \max_{i,j \in [p]} \{\frac{d_i}{d_j}\}$ need only be computed over indices i, j such that $z_i, z_j = 1$. Moreover, for a partially specified \mathbf{z} —which might arise at an intermediate node in a branch-and-bound tree generated by Algorithm 1—the ratio ρ need only be computed over indices i where z_i is unspecified or set to 1. This suggests that the quality of the Gershgorin bound improves upon branching.

Remark 10 In particular, if $\Sigma \in S_+^n$ is a diagonal matrix, then Equation (13)'s bound is tight - which follows from the fact that the spectrum of Σ and the discs coincide if and only if Σ is diagonal (see, e.g. Horn and Johnson, 1990, Chapter 6). Alternatively, if Σ is a diagonally dominant matrix then $\rho = 1$ and the Gershgorin circle theorem provides a 2-factor approximation.

Proof Scaled diagonally dominant matrices have scaled diagonally dominant principal minors—this is trivially true because

$$d_i \Sigma_{i,i} \geq \sum_{j \in [p]: j \neq i} d_j |\Sigma_{i,j}| \quad \forall i \in [p] \implies d_i \Sigma_{i,i} \geq \sum_{j \in [p]: j \neq i} d_j z_j |\Sigma_{i,j}| \quad \forall i \in [p] : z_i = 1$$

for the same vector $\mathbf{d} > \mathbf{0}$ and therefore the following chain of inequalities holds

$$\begin{aligned} f(\mathbf{z}) &\leq \max_{j \in [p]} \left\{ \sum_{i \in [p]} z_i |\Sigma_{i,j}| \right\} = \max_{j \in [p]} \left\{ z_j \Sigma_{j,j} + \sum_{i \in [p]: i \neq j} z_i |\Sigma_{i,j}| \right\} \\ &\leq \max_{j \in [p]} \left\{ z_j \Sigma_{j,j} + \sum_{i \in [p]: i \neq j} \rho \frac{d_i}{d_j} z_i |\Sigma_{i,j}| \right\} \leq (1 + \rho) \max_{j \in [p]} \{z_j \Sigma_{j,j}\} \leq (1 + \rho) f(\mathbf{z}) \quad \forall \mathbf{z} \in \{0, 1\}^p, \end{aligned}$$

where the second inequality follows because $\rho \geq \frac{d_i}{d_j}$, the third inequality follows from the scaled diagonal dominance of the principal submatrices of Σ , and the fourth inequality holds because the leading eigenvalue of a PSD matrix is at least as large as each diagonal. ■

To make clear the extent our numerical success depends upon Theorem 7, our results in Section 4 present implementations of Algorithm 1 both with and without the bound.

2.5 Beyond Gershgorin: Further Strengthening via Brauer's Ovals of Cassini

Given the relevance of Gershgorin's bound, we propose, in this section, a stronger —yet more expensive to implement— upper bound, based on an generalization of the Gershgorin Circle theorem, namely Brauer's ovals of Cassini.

First, we derive a new upper-bound on $f(\mathbf{z})$ that is at least as strong as the one presented in Theorem 7 and often strictly stronger (Horn and Johnson, 1990, Chapter 6):

Theorem 11 *For any vector $\mathbf{z} \in \{0, 1\}^p$, we have the following upper bound on $f(\mathbf{z})$:*

$$f(\mathbf{z}) \leq \max_{i,j \in [p]: i > j, z_i = z_j = 1} \left\{ \frac{\Sigma_{i,i} + \Sigma_{j,j}}{2} + \frac{\sqrt{(\Sigma_{i,i} - \Sigma_{j,j})^2 + 4R_i(\mathbf{z})R_j(\mathbf{z})}}{2} \right\}, \quad (16)$$

where $R_i(\mathbf{z}) := \sum_{j \in [p]: j \neq i} z_j |\Sigma_{i,j}|$ is the absolute sum of off-diagonal entries in the i th column of the submatrix of Σ induced by \mathbf{z} .

Proof Let us first recall that, per Brauer (1952)'s original result, all eigenvalues of a matrix $\Sigma \in S_+^p$ are contained in the union of the following $p(p-1)/2$ ovals of Cassini:

$$\bigcup_{i \in [p], j \in [p]: i < j} \{ \lambda \in \mathbb{R}_+ : |\lambda - \Sigma_{i,i}| |\lambda - \Sigma_{j,j}| \leq R_i R_j \},$$

where $R_i := \sum_{j \in [p]: j \neq i} |\Sigma_{i,j}|$ is the absolute sum of off-diagonal entries in the i th column of Σ . Next, let us observe that, if λ is a dominant eigenvalue of a PSD matrix Σ then $\lambda \geq \Sigma_{i,i} \forall i$ and, in the (i, j) th oval, the bound reduces to

$$\lambda^2 - \lambda(\Sigma_{i,i} + \Sigma_{j,j}) + \Sigma_{i,i}\Sigma_{j,j} - R_i R_j \leq 0, \quad (17)$$

which, by the quadratic formula, implies an upper bound is $\frac{\Sigma_{i,i} + \Sigma_{j,j}}{2} + \frac{\sqrt{(\Sigma_{i,i} - \Sigma_{j,j})^2 + 4R_i R_j}}{2}$. The result follows because if $z_i = 0$ the i th row of Σ cannot be used to bound $f(\mathbf{z})$. ■

Theorem 11's inequality can be enforced numerically as mixed-integer second order cone constraints. Indeed, the square root term in (16) can be modeled using second-order cone, and the bilinear terms only involve binary variables and can be linearized. Completing the square in Equation (17), (16) is equivalent to the following system of $p(p-1)/2$ mixed-integer second-order cone inequalities:

$$\left(\theta - \frac{1}{2}(\Sigma_{i,i} + \Sigma_{j,j}) \right)^2 \leq \sum_{s,t \in [p]: s \neq i, t \neq j} W_{s,t} |\Sigma_{i,s} \Sigma_{j,t}| - \frac{3}{4} \Sigma_{i,i} \Sigma_{j,j} + M(1 - s_{i,j}) \quad \forall i, j \in [p] : i < j,$$

$$\sum_{i,j \in [p]: i < j} s_{i,j} = 1, s_{i,j} \leq \min(z_i, z_j) \quad i, j \in [p] : i < j, \quad s_{i,j} \in \{0, 1\} \quad i, j \in [p] : i < j.$$

where $W_{i,j} = z_i z_j$ is a product of binary variables which can be modeled using, e.g., the McCormick inequalities $\max(0, z_i + z_j - 1) \leq W_{i,j} \leq \min(z_i, z_j)$, and M is an upper bound

on the right-hand-side of the inequality for any $i, j : i \neq j$, which can be computed in $O(p^3)$ time in much the same manner as a big- M constant was computed in the previous section. Note that we do not make use of these inequalities directly in our numerical experiments, due to their high computational cost. However, an interesting extension would be to introduce the binary variables dynamically, via branch-and-cut-and-price (Barnhart et al., 1998).

Since the bound derived from the ovals of Cassini (Theorem 11) is at least as strong as the Gershgorin circle's one (Theorem 7), it satisfies the same approximation guarantee (Proposition 8). In particular, it is tight when Σ is diagonal and provides a 2-factor approximation for diagonally dominant matrices. Actually, we now prove a stronger result and demonstrate that Theorem 11 provides a 2-factor bound on $f(\mathbf{z})$ for doubly diagonally dominant matrices—a broader class of matrices than diagonally dominant matrices (see Li and Tsatsomeros, 1997, for a general theory):

Proposition 12 *Let $\Sigma \in S_+^p$ be a doubly diagonally dominant matrix, i.e.,*

$$\Sigma_{i,i}\Sigma_{j,j} \geq R_i R_j \quad \forall i, j \in [p] : i > j,$$

where $R_i := \sum_{j \in [p]: j \neq i} |\Sigma_{i,j}|$ is the sum of the off-diagonal entries in the i th column of Σ . Then, we have that

$$f(\mathbf{z}) \leq \max_{i,j \in [p]: i > j, z_i = z_j = 1} \left\{ \frac{\Sigma_{i,i} + \Sigma_{j,j}}{2} + \frac{\sqrt{(\Sigma_{i,i} - \Sigma_{j,j})^2 + 4R_i(\mathbf{z})R_j(\mathbf{z})}}{2} \right\} \leq 2f(\mathbf{z}). \quad (18)$$

Proof Observe that if $\Sigma_{i,i}\Sigma_{j,j} \geq R_i R_j$ then

$$\sqrt{(\Sigma_{i,i} - \Sigma_{j,j})^2 + 4R_i R_j} \leq \sqrt{(\Sigma_{i,i} - \Sigma_{j,j})^2 + 4\Sigma_{i,i}\Sigma_{j,j}} = \Sigma_{i,i} + \Sigma_{j,j}.$$

The result then follows in essentially the same fashion as Proposition 8. ■

3. Convex Relaxations and Rounding Methods

For large-scale instances, high-quality solutions can be obtained by solving a convex relaxation of Problem (5) and rounding the optimal solution. In Section 3.1, we propose relaxing $\mathbf{z} \in \{0, 1\}^p$ in (5) to $\mathbf{z} \in [0, 1]^p$ and applying a greedy rounding scheme. We further tighten this relaxation using second-order cones constraints in Section 3.2.

3.1 A Boolean Relaxation and a Greedy Rounding Method

We first consider a Boolean relaxation of (5), which we obtain² by relaxing $\mathbf{z} \in \{0, 1\}^p$ to $\mathbf{z} \in [0, 1]^p$. This gives $\max_{\mathbf{z} \in [0, 1]^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z})$, i.e.,

$$\max_{\mathbf{z} \in [0, 1]^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \succeq \mathbf{0}} \langle \mathbf{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, |X_{i,j}| \leq M_{i,j} z_i \quad \forall i, j \in [p]. \quad (19)$$

A useful strategy for obtaining a high-quality feasible solution is to solve (19) and set $z_i = 1$ for k indices corresponding to the largest z_j 's in (19) as proposed in the randomized case for general integer optimization problems by Raghavan and Tompson (1987). We formalize this in Algorithm 2. We remark that rounding strategies for sparse PCA have previously been proposed (see Fountoulakis et al., 2017; Dey et al., 2017; Chowdhury et al., 2020), however, the idea of rounding \mathbf{z} and then optimizing for \mathbf{X} appears to be new.

Algorithm 2 A greedy rounding method for Problem (1)

Require: Covariance matrix $\mathbf{\Sigma}$, sparsity parameter k

Compute \mathbf{z}^* solution of (19) or (20)

Construct $\mathbf{z} \in \{0, 1\}^p : \mathbf{e}^\top \mathbf{z} = k$ such that $z_i \geq z_j$ if $z_i^* \geq z_j^*$.

Compute \mathbf{X} solution of

$$\max_{\mathbf{X} \in S_+^p} \langle \mathbf{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, X_{i,j} = 0 \text{ if } z_i z_j = 0 \quad \forall i, j \in [p].$$

return \mathbf{z}, \mathbf{X} .

Remark 13 *Our numerical results in Section 4 reveal that explicitly imposing a PSD constraint on \mathbf{X} in the relaxation (19)—or the ones derived later in the following section—prevents our approximation algorithm from scaling to larger problem sizes than the exact Algorithm 1 can already solve. Therefore, to improve scalability, the semidefinite cone can be safely approximated via its second-order cone relaxation, $X_{i,j}^2 \leq X_{i,i} X_{j,j} \quad \forall i, j \in [p]$, plus a small number of cuts of the form $\langle \mathbf{X}, \mathbf{x}_t \mathbf{x}_t^\top \rangle \geq 0$ as presented in Bertsimas and Cory-Wright (2020).*

Remark 14 *Rather than relaxing and greedily rounding \mathbf{z} , one could consider a higher dimensional relax-and-round scheme where we let \mathbf{Z} model the outer product $\mathbf{z}\mathbf{z}^\top$ via $\mathbf{Z} \succeq \mathbf{z}\mathbf{z}^\top$, $\max(0, z_i + z_j - 1) \leq Z_{i,j} \leq \min(z_i, z_j) \quad \forall i, j \in [p]$, $Z_{i,i} = z_i$, and require that $\sum_{i,j \in [p]} Z_{i,j} \leq k^2$. Indeed, a natural “round” component of such a relax-and-round scheme is precisely Goemans-Williamson rounding (Goemans and Williamson, 1995; Bertsimas and Ye, 1998), which performs at least as well as greedy rounding in both theory and*

2. Note that we omit the $\sum_{j=1}^p |X_{i,j}| \leq \sqrt{k z_i}$ constraints when we develop our convex relaxations, since they are essentially dominated by the $\|\mathbf{X}\|_1 \leq k$ constraint we introduce in the next section; we introduced these inequalities to improve our semidefinite-free subproblem strategy for the exact method.

practice. Unfortunately, some preliminary numerical experiments indicated that Goemans-Williamson rounding is not actually much better than greedy rounding in practice, and is considerably more expensive to implement. Therefore, we defer the details of the Goemans-Williamson scheme to Appendix A, and do not consider it any further in this paper.

3.2 Valid Inequalities for Strengthening Convex Relaxations

We now propose valid inequalities which allow us to improve the quality of the convex relaxations discussed previously. Note that as convex relaxations and random rounding methods are two sides of the same coin (Barak et al., 2014), applying these valid inequalities also improves the quality of the randomly rounded solutions.

Theorem 15 *Let \mathcal{P}_{strong} denote the optimal objective value of the following problem:*

$$\begin{aligned} \max_{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \in S_+^p} \langle \Sigma, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, |X_{i,j}| \leq M_{i,j} z_i \quad \forall i, j \in [p], \\ \sum_{j \in [p]} X_{i,j}^2 \leq X_{i,i} z_i, \|\mathbf{X}\|_1 \leq k. \end{aligned} \quad (20)$$

Then, (20) is a stronger relaxation than (19), i.e., the following inequalities hold:

$$\max_{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z}) \geq \mathcal{P}_{strong} \geq \max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z}).$$

Moreover, suppose an optimal solution to (20) is of rank one. Then, the relaxation is tight:

$$\mathcal{P}_{strong} = \max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z}).$$

Proof The first inequality $\max_{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k} f(\mathbf{z}) \geq \mathcal{P}_{strong}$ is trivial. The second inequality holds because \mathcal{P}_{strong} is indeed a valid relaxation of Problem (1). Indeed, $\|\mathbf{X}\|_1 \leq k$ follows from the cardinality and big-M constraints. The semidefinite constraint $\mathbf{X} \succeq 0$ impose second-order cone constraints on the 2×2 minors of \mathbf{X} , $X_{i,j}^2 \leq z_i X_{i,i} X_{j,j}$, which can be aggregated into $\sum_{j \in [p]} X_{i,j}^2 \leq X_{i,i} z_i$ (see Bertsimas and Cory-Wright, 2020, for derivations).

Finally, suppose that an optimal solution to Problem (20) is of rank one, i.e., the optimal matrix \mathbf{X} can be decomposed as $\mathbf{X} = \mathbf{x}\mathbf{x}^\top$. Then, the SOCP inequalities imply that $\sum_{j \in [p]} x_i^2 x_j^2 \leq x_i^2 z_i$. However, $\sum_{j \in [p]} x_j^2 = \text{tr}(\mathbf{X}) = 1$, which implies that $x_i^2 \leq x_i^2 z_i$, i.e., $z_i = 1$ for any index i such that $|x_i| > 0$. Since $\mathbf{e}^\top \mathbf{z} \leq k$, this implies that $\|\mathbf{x}\|_0 \leq k$, i.e., \mathbf{X} also solves Problem (2). \blacksquare

As our numerical experiments will demonstrate and despite the simplicity of our rounding mechanism in Algorithm 2, the relaxation (20) provides high-quality solutions to the original sparse PCA problem (1), without introducing any additional variables. We remark that other inequalities, including the second-order cone inequalities proposed in Li and Xie (2020, Lemma 2 (ii)), could further improve the convex relaxation; we leave integrating these inequalities within our framework as future work.

4. Numerical Results

We now assess the numerical behavior of the algorithms proposed in Section 2 and 3. To bridge the gap between theory and practice, we present a `Julia` code which implements the described convex relaxation and greedy rounding procedure on GitHub³. The code requires a conic solver such as `Mosek` and several open source `Julia` packages to be installed.

4.1 Performance of Exact Methods

In this section, we apply Algorithm 1 to medium and large-scale sparse principal component analysis problems, with and without Gershgorin circle theorem bounds in the master problem. All experiments were implemented in `Julia` 1.3, using `Gurobi` 9.1 and `JuMP.jl` 0.21.6, and performed on a standard Macbook Pro laptop, with a 2.9GHz 6-Core Intel i9 CPU, using 16 GB DDR4 RAM. We compare our approach to the branch-and-bound algorithm⁴ developed by Berk and Bertsimas (2019) on the UCI `pitprops`, `wine`, `miniboone`, `communities`, `arrythmia` and `micromass` datasets, both in terms of runtime and the number of nodes expanded; we refer to Berk and Bertsimas (2019); Bertsimas and Cory-Wright (2020) for descriptions of these datasets. Note that we normalized all datasets before running the method (i.e., we compute the leading sparse principal components of correlation matrices). Additionally, we warm-start all methods with the solution from the method of Yuan and Zhang (2013), to maintain a fair comparison.

Table 1 reports the time for Algorithm 1 (with and without Gershgorin circle theorem bounds in the master problem) and the method of Berk and Bertsimas (2019) to identify the leading k -sparse principal component for $k \in \{5, 10, 20\}$, along with the number of nodes expanded, and the number of outer approximation cuts generated. We impose a relative optimality tolerance of 10^{-3} for all approaches, i.e., terminate each method when $(UB - LB)/UB \leq 10^{-3}$ where UB denotes the current objective bound and LB denotes the current incumbent objective value. Note that p denotes the dimensionality of the correlation matrix, and $k \leq p$ denotes the target sparsity.

Our main findings from these experiments are as follows:

- For smaller problems, the strength of Algorithm 1’s cuts allows it to outperform state-of-the-art methods such as the method of Berk and Bertsimas (2019). Moreover, for larger problem sizes, the adaptive branching strategy performs comparably to Algorithm 1. This suggests that the relative merits of both approaches are roughly even, and which method is preferable may depend on the problem data.

3. <https://github.com/ryancorywright/ScalableSPCA.jl>

4. The solve times for their method, as reported here, differ from those reported in Berk and Bertsimas (2019) due to a small typo in their implementation (line 110 of their `branchAndBound.jl` code should read “if $y[i] == -1 || y[i] == 1$ ”, not “if $y[i] == -1$ ” in order to correctly compute the Gershgorin circle theorem bound); correcting this is necessary to ensure that we obtain correct results from their method.

Table 1: Runtime in seconds per approach. We impose a time limit of 600s. If a solver fails to converge, we report the relative bound gap at termination in brackets.

Dataset	p	k	Alg. 1			Alg. 1+ Circle Theorem			Method of B.+B.	
			Time(s)	Nodes	Cuts	Time(s)	Nodes	Cuts	Time(s)	Nodes
Pitprops	13	5	0.30	1,608	1,176	0.06	38	8	1.49	22
		10	0.14	414	387	0.02	18	21	0.02	14
Wine	13	5	0.57	2,313	1,646	0.02	46	11	0.04	34
		10	0.17	376	311	0.03	54	58	0.02	12
Miniboone	50	5	0.01	0	11	0.01	0	3	0.04	2
		10	0.01	0	16	0.02	0	3	0.04	2
		20	0.03	0	26	0.01	0	3	1.30	5,480
Communities	101	5	(2.87%)	28,462	25,483	0.20	201	3	0.57	101
		10	(13.3%)	37,479	36,251	0.34	406	39	0.94	1,298
		20	(39.6%)	24,566	24,632	(12.1%)	42,120	37,383	(9.97%)	669,500
Arrhythmia	274	5	(18.1%)	22,771	20,722	6.07	135	1,233	4.17	1,469
		10	(32.6%)	19,500	19,314	(2.92%)	15,510	6,977	(0.83%)	471,680
		20	(74.4%)	33,773	12,374	(24.3%)	33,123	19,662	(18.45%)	311,400
Micromass	1300	5	(1.29%)	3,859	3,099	163.60	2,738	6	24.31	1,096
		10	(10.6%)	3,366	3,369	241.86	3,233	121	362.4	36,690
		20	(35.9%)	2,797	2,839	(35.9%)	2,676	2,115	(10.34%)	31,990

- Generating outer-approximation cuts and valid upper bounds from the Gershgorin circle theorem are both powerful ideas, but the greatest aggregate power appears to arise from intersecting these bounds, rather than using one bound alone.
- Once both k and p are sufficiently large (e.g. $p > 300$ and $k > 10$), no approach is able to solve the problem to provable optimality within 600s. This motivates our study of convex relaxations and randomized rounding methods in the next section.

4.2 Convex Relaxations and Randomized Rounding Methods

In this section, we apply Algorithm 2 to obtain high quality convex relaxations and feasible solutions for the datasets studied in the previous subsection, and compare the relaxation to a difference convex relaxation developed by d’Aspremont et al. (2008), in terms of the quality of the upper bound and the resulting greedily rounded solutions. All experiments were implemented using the same specifications as the previous section. Note that d’Aspremont

et al. (2008)’s upper bound⁵ which we compare against is:

$$\max_{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \succeq \mathbf{0}, \mathbf{P}_i \succeq \mathbf{0} \forall i \in [p]} \sum_{i \in [p]} \langle \mathbf{a}_i \mathbf{a}_i^\top, \mathbf{P}_i \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, \text{tr}(\mathbf{P}_i) = z_i, \mathbf{X} \succeq \mathbf{P}_i \forall i \in [p], \quad (21)$$

where $\mathbf{\Sigma} = \sum_{i=1}^p \mathbf{a}_i \mathbf{a}_i^\top$ is a Cholesky decomposition of $\mathbf{\Sigma}$, and we obtain feasible solutions from this relaxation by greedily rounding an optimal \mathbf{z} in the bound *à la* Algorithm 2. To allow for a fair comparison, we also consider augmenting this formulation with the inequalities derived in Section 3.2 to obtain the following stronger yet more expensive to solve relaxation:

$$\max_{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\substack{\mathbf{X} \succeq \mathbf{0}, \\ \mathbf{P}_i \succeq \mathbf{0} \forall i \in [p]}} \sum_{i \in [p]} \langle \mathbf{a}_i \mathbf{a}_i^\top, \mathbf{P}_i \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, \text{tr}(\mathbf{P}_i) = z_i, \mathbf{X} \succeq \mathbf{P}_i \forall i \in [p], \\ \sum_{j \in [p]} X_{i,j}^2 \leq X_{i,i} z_i, \|\mathbf{X}\|_1 \leq k. \quad (22)$$

We first apply these relaxations on datasets where Algorithm 1 terminates, hence the optimal solution is known and can be compared against. We report the quality of both methods with and without the additional inequalities discussed in Section 3.2, in Tables 2-3 respectively⁶.

Observe that applying Algorithm 2 without the additional inequalities (Table 2) yields rather poor relaxations and randomly rounded solutions. However, by intersecting our relaxations with the additional inequalities from Section 3.2 (Table 3), we obtain extremely high quality relaxations. Indeed, with the additional inequalities, Algorithm 2 using formulation (20) identifies the optimal solution in all instances (0% O. gap), and always supplies a bound gap of less than 2%. Moreover, in terms of obtaining high-quality solutions, the new inequalities allow Problem (20) to perform as well or better as Problem (21), despite optimizing over one semidefinite matrix, rather than $p + 1$ semidefinite matrices. This suggests that Problem (20) should be considered as a viable, more scalable and more accurate alternative to existing SDO relaxations such as Problem (21). For this reason, we shall only consider using Problem (20)’s formulation for the rest of the paper.

5. Strictly speaking, d’Aspremont et al. (2008) does not actually write down this formulation in their work. Indeed, their bound involves dual variables which cannot be used directly to generate feasible solutions via greedy rounding. However, the fact that this bound and (d’Aspremont et al., 2008, Problem (8)) are dual to each other follows directly from strong semidefinite duality, and therefore we refer to this formulation as being due to d’Aspremont et al. (2008) (it essentially is).

6. For the instances of (21) or (22) where $p > 13$ we used SCS version 2.1.1 (with default parameters) instead of Mosek, since Mosek required more memory than was available in our computing environment, and SCS takes an augmented Lagrangian approach which is less numerically stable but requires significantly less memory. That is, (21)’s formulation is too expensive to solve via IPMs on a laptop when $p = 50$.

Table 2: Quality of relaxation gap (upper bound vs. optimal solution-denoted R. gap), objective gap (rounded solution vs. optimal solution-denoted O. gap) and runtime in seconds per method.

Dataset	p	k	Alg. 2 with (19)			Alg. 2 with (21)		
			R. gap (%)	O. gap (%)	Time(s)	R. gap (%)	O. gap (%)	Time(s)
Pitprops	13	5	23.8%	0.00%	0.02	23.8%	16.1%	0.46
		10	1.10%	0.30%	0.03	1.10%	1.33%	0.46
Wine	13	5	36.8%	0.00%	0.02	36.8%	40.4%	0.433
		10	2.43%	0.26%	0.03	2.43%	15.0%	0.463
Miniboone	50	5	781.3%	235.6%	7.37	781.2%	34.7%	1,191.0
		10	340.6%	117.6%	7.50	340.6%	44.9%	1,102.6
		20	120.3%	38.08%	6.25	120.3%	31.9%	1,140.2

Table 3: Quality of relaxation gap (upper bound vs. optimal solution-denoted R. gap), objective gap (rounded solution vs. optimal solution-denoted O. gap) and runtime in seconds per method, with additional inequalities from Section 3.2.

Dataset	p	k	Alg. 2 with (20)			Alg. 2 with (22)		
			R. gap (%)	O. gap (%)	Time(s)	R. gap (%)	O. gap (%)	Time(s)
Pitprops	13	5	0.71%	0.00%	0.17	1.53%	0.00%	0.55
		10	0.12%	0.00%	0.27	1.10%	0.00%	3.27
Wine	13	5	1.56%	0.00%	0.24	2.98%	15.03%	0.95
		10	0.40%	0.00%	0.22	2.04%	0.00%	1.15
Miniboone	50	5	0.00%	0.00%	163.3	0.00%	0.01%	500.7
		10	0.00%	0.00%	148.5	0.00%	0.02%	489.9
		20	0.00%	0.00%	194.5	0.00%	0.00%	776.3

We remark however that the key drawback of applying these methods is that, as implemented in this section, they do not scale to sizes beyond which Algorithm 1 successfully solves. This is a drawback because Algorithm 1 supplies an exact certificate of optimality, while these methods do not. In the following set of experiments, we therefore investigate numerical techniques to improve the scalability of Algorithm 2.

4.3 Scalable Dual Bounds and Randomized Rounding Methods

To improve the scalability of Algorithm 2, we relax the PSD constraint on \mathbf{X} in (19) and (20). With these enhancements, we demonstrate that Algorithm 2 can be successfully scaled to generate high-quality bounds for $1000s \times 1000s$ matrices.

As discussed in Remark 13, we can replace the PSD constraint $\mathbf{X} \succeq \mathbf{0}$ by requiring that the $p(p-1)/2$ two by two minors of \mathbf{X} are non-negative: $X_{i,j}^2 \leq X_{i,i}X_{j,j}$. Second, we consider adding 20 linear inequalities of the form $\langle \mathbf{X}, \mathbf{x}_t \mathbf{x}_t^\top \rangle \geq 0$, for some vector \mathbf{x}_t (see Bertsimas and Cory-Wright, 2020, for a discussion). Table 4 reports the performance of Algorithm 2 (with the relaxation (20)) with these two approximations of the positive semidefinite cone, “Minors” and “Minors + 20 inequalities” respectively. Note that we report the entire duality gap (i.e. do not break the gap down into its relaxation and objective gap components) since, as reflected in Table 1, some of these instances are currently too large to solve to optimality.

Table 4: Quality of bound gap (rounded solution vs. upper bound) and runtime in seconds of Algorithm 2 with (20), outer-approximation of the PSD cone.

Dataset	p	k	Minors		Minors + 20 inequalities	
			Gap (%)	Time(s)	Gap (%)	Time(s)
Pitprops	13	5	1.51%	0.02	0.72%	0.36
		10	5.29%	0.02	1.12%	0.36
Wine	13	5	2.22%	0.02	1.59%	0.38
		10	3.81%	0.02	1.50%	0.37
Miniboone	50	5	0.00%	0.11	0.00%	0.11
		10	0.00%	0.12	0.00%	0.12
		20	0.00%	0.39	0.00%	0.39
Communities	101	5	0.07%	0.67	0.07%	14.8
		10	0.66%	0.68	0.66%	14.4
		20	3.32%	1.84	2.23%	33.5
Arrhythmia	274	5	3.37%	27.2	1.39%	203.6
		10	3.01%	25.6	1.33%	184.0
		20	8.87%	21.8	4.48%	426.8
Micromass	1300	5	0.04%	239.4	0.01%	4,639
		10	0.63%	232.6	0.32%	6,392
		20	13.1%	983.5	5.88%	16,350

Observe that if we impose constraints on the 2×2 minors only then we obtain a solution certifiably within 13% of optimality in seconds (resp. minutes) for $p = 100$ s (resp. $p = 1000$ s). Moreover, adding 20 linear inequalities, we obtain a solution within 6% of optimality in minutes (resp. hours) for $p = 100$ s (resp. $p = 1000$ s). Moreover, the bound gaps compare favorably to Algorithm 1 and the method of Berk and Bertsimas (2019) for instances which these methods could not solve to certifiable optimality. For instance, for the Arrhythmia dataset when $k = 20$ we obtain a bound gap of less than 9% in 20s, while the method of Berk and Bertsimas (2019) obtains a bound gap of 18.45% in 600s. This illustrates the

value of the proposed relax+round method on datasets which are currently too large to be optimized over exactly.

To conclude this section, we explore Algorithm 2’s ability to scale to even higher dimensional datasets in a high performance setting, by running the method on one Intel Xeon E5–2690 v4 2.6GHz CPU core using 600 GB RAM. Table 5 reports the methods scalability and performance on the Wilshire 5000, and **Arcene** UCI datasets. For the **Gisette** dataset, we report on the methods performance when we include the first 3,000 and 4,000 rows/columns (as well as all 5,000 rows/columns). Similarly, for the **Arcene** dataset we report on the method’s performance when we include the first 6,000, 7,000 or 8,000 rows/columns. We do not report results for the **Arcene** dataset for $p > 8,000$, as computing this requires more memory than was available (i.e. > 600 GB RAM). We do not report the method’s performance when we impose linear inequalities for the PSD cone, as solving the relaxation without them is already rather time consuming. Moreover, we do not impose the 2×2 minor constraints to save memory, do not impose $|X_{i,j}| \leq M_{i,j}z_i$ when $p \geq 4000$ to save even more memory, and report the overall bound gap, as improving upon the randomly rounded solution is challenging in a high-dimensional setting.

These results suggest that if we solve the SOC relaxation using a first-order method rather than an interior point method, our approach could successfully generate certifiably near-optimal PCs when $p = 10,000$ s, particularly if combined with a feature screening technique (see d’Aspremont et al., 2008; Atamtürk and Gomez, 2020).

4.4 Performance of Exact and Approximate Methods on Synthetic Data

We now compare the exact and approximate methods against existing state-of-the-art methods in a spiked covariance matrix setting. We use the experimental setup laid out in d’Aspremont et al. (2008, Section 7.1). We recover the leading principal component of a test matrix⁷ $\Sigma \in S_+^p$, where $p = 150$, $\Sigma = \frac{1}{n}\mathbf{U}^\top\mathbf{U} + \frac{\sigma}{\|\mathbf{v}\|_2}\mathbf{v}\mathbf{v}^\top$, $\mathbf{U} \in [0, 1]^{150 \times 150}$ is a noisy matrix with i.i.d. standard uniform entries, $\mathbf{v} \in \mathbb{R}^{150}$ is a vector of signals such that

$$v_i = \begin{cases} 1, & \text{if } i \leq 50, \\ \frac{1}{i-50}, & \text{if } 51 \leq i \leq 100, \\ 0, & \text{otherwise,} \end{cases} \quad (23)$$

and $\sigma = 2$ is the signal-to-noise ratio. The methods which we compare are:

- **Exact:** Algorithm 1 with Gershgorin inequalities and a time limit of 600s.
- **Approximate:** Algorithm 2 with Problem (20), the SOC outer-approximation of the PSD cone, no PSD cuts, and the additional SOC inequalities.

7. This statement of the test matrix is different to d’Aspremont et al. (2008, Section 7.1), who write $\Sigma = \mathbf{U}^\top\mathbf{U} + \sigma\mathbf{v}\mathbf{v}^\top$, rather than $\Sigma = \frac{1}{n}\mathbf{U}^\top\mathbf{U} + \frac{\sigma}{\|\mathbf{v}\|_2}\mathbf{v}\mathbf{v}^\top$. However, it agrees with their source code.

Table 5: Quality of bound gap (rounded solution vs. upper bound) and runtime in seconds.

Dataset	p	k	Algorithm 2 (SOC relax)+Inequalities	
			Bound gap (%)	Time(s)
Wilshire 5000	2130	5	0.38%	1,036
		10	0.24%	1,014
		20	0.36%	1,059
Gisette	3000	5	1.67%	2,249
		10	35.81%	2,562
		20	10.61%	3,424
Gisette	4000	5	1.55%	1,402
		10	54.4%	1,203
		20	11.84%	1,435
Gisette	5000	5	1.89%	2,169
		10	2.22%	2,455
		20	7.16%	2,190
Arcene	6000	5	0.01%	3,333
		10	0.06%	3,616
		20	0.14%	3,198
Arcene	7000	5	0.03%	4,160
		10	0.05%	4,594
		20	0.25%	4,730
Arcene	8000	5	0.02%	6,895
		10	0.17%	8,479
		20	0.21%	6,335

- **Greedy:** as proposed by Moghaddam et al. (2006) and laid out in (d’Aspremont et al., 2008, Algorithm 1), start with a solution \mathbf{z} of cardinality 1 and iteratively augment this solution vector with the index which gives the maximum variance contribution. Note that d’Aspremont et al. (2008) found this method outperformed the 3 other methods (approximate greedy, thresholding and sorting) they considered in their work.
- **Truncated Power Method:** as proposed by Yuan and Zhang (2013), alternate between applying the power method to the solution vector and truncating the vector to ensure that it is k -sparse. Note that Berk and Bertsimas (2019) found that this approach performed better than 5 other state-of-the-art methods across the real-world datasets studied in the previous section of this paper and often matched the performance of the method of Berk and Bertsimas (2019)—indeed, it functions as a warm-start for the later method.

- **Sorting:** sort the entries of $\Sigma_{i,i}$ by magnitude and set $z_i = 1$ for the k largest entries of Σ , as studied in d’Aspremont et al. (2008). This naive method serves as a benchmark for the value of optimization in the more sophisticated methods considered here.

Figure 1 depicts the ROC curve (true positive rate vs. false positive rate for recovering the support of \mathbf{v}) over 20 synthetic random instances, as we vary k for each instance. We observe that among all methods, the sorting method is the least accurate, with a substantially larger false detection rate for a given true positive rate than the remaining methods (AUC= 0.7028). The truncated power method and our exact method⁸ then offer a substantial improvement over sorting, with respective AUCs of 0.7482 and 0.7483. The greedy method then offers a modest improvement over them (AUC= 0.7561) and the approximate relax+round method is the most accurate (AUC= 0.7593).

In addition to support recovery, Figure 2 reports average runtime (left panel) and average optimality gap (right panel) over the same instances. Observe that among all methods, only the exact and the approximate relax+round methods provide optimality gaps, i.e., numerical certificates of near optimality. On this metric, relax+round supplies average bound gaps of 1% or less on all instances, while the exact method typically supplies bound gaps of 30% or more. This comparison illustrates the tightness of the valid inequalities from Section 3.2 that we included in the relaxation. Moreover, the relax+round method converges in less than one minute on all instances. All told, the relax+round method is the best performing method overall, although if k is set to be sufficiently close to 0 or p all methods behave comparably. In particular, the relax+round method should be preferred over the exact method, even though the exact method performs better at smaller problem sizes.

4.5 Summary and Guidelines From Experiments

In summary, our main findings from our numerical experiments are as follows:

- For small or medium scale problems where $p \leq 100$ or $k \leq 10$, exact methods such as Algorithm 1 or the method of Berk and Bertsimas (2019) reliably obtain certifiably optimal or near-optimal solutions in a short amount of time, and should therefore be preferred over other methods. However, for larger-scale sparse PCA problems, exact methods currently do not scale as well as approximate or heuristic methods.
- For larger-scale sparse PCA problems, our proposed combination of solving a second-order cone relaxation and rounding greedily reliably supplies certifiably near-optimal solutions in practice (if not in theory) in a relatively small amount of time. Moreover, it outperforms other state-of-the-art heuristics including the greedy method of Moghaddam et al. (2006); d’Aspremont et al. (2008) and the Truncated Power Method

8. Note that the exact method would dominate the remaining methods if given an unlimited runtime budget. Its poor performance reflects its inability to find the true optimal solution within 600 seconds.

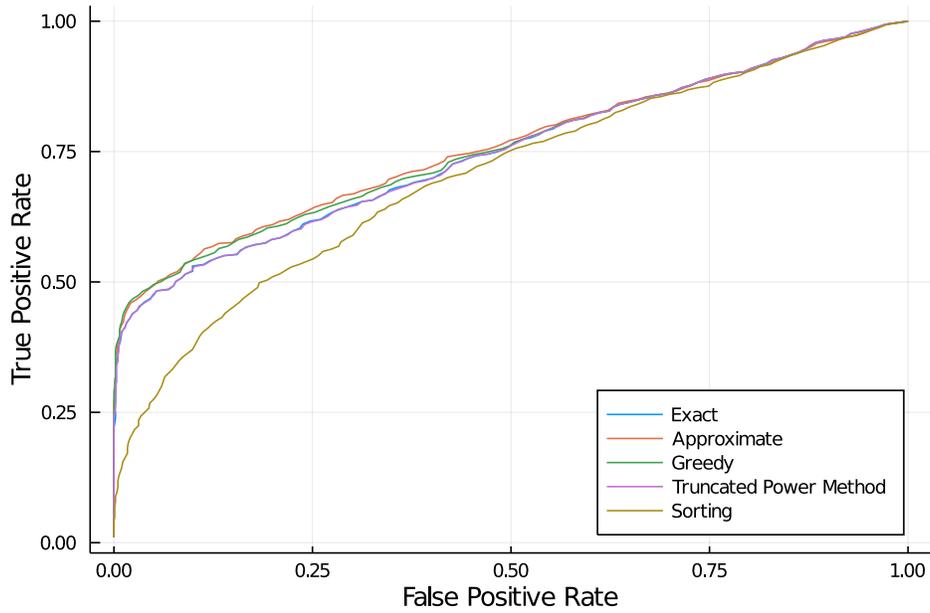


Figure 1: ROC curve over 20 synthetic instances where $p = 150$, $k_{\text{true}} = 100$ is unspecified.

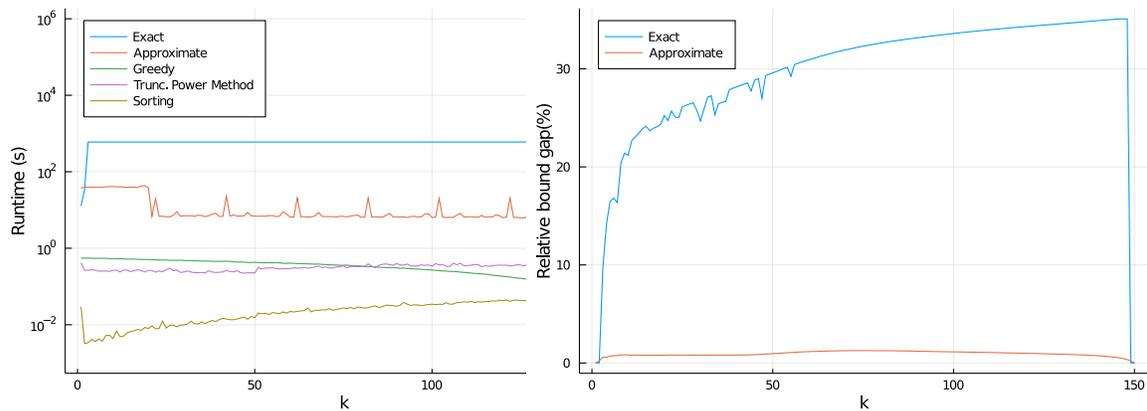


Figure 2: Average time to compute solution, and optimality gap over 20 synthetic instances where $p = 150$, $k_{\text{true}} = 100$ is unspecified.

of Yuan and Zhang (2013). Accordingly, it should be considered as a reliable and more accurate alternative for problems where $p = 1000$ s.

- In practice, for even larger-scale problem sizes, we recommend using a combination of these methods: a computationally cheaper method (with k set in the 1000s) as a feature screening method, to be followed by the approximate relax+round method (with k set in the 100s) and/or the exact method, if time permits.

5. Three Extensions and their Mixed-Integer Conic Formulations

We conclude by discussing three extensions of sparse PCA where our methodology applies.

5.1 Non-Negative Sparse PCA

One potential extension to this paper would be to develop a certifiably optimal algorithm for non-negative sparse PCA (see Zass and Shashua, 2007, for a discussion), i.e., develop a tractable reformulation of

$$\max_{\mathbf{x} \in \mathbb{R}^p} \langle \mathbf{x}\mathbf{x}^\top, \boldsymbol{\Sigma} \rangle \text{ s.t. } \mathbf{x}^\top \mathbf{x} = 1, \mathbf{x} \geq \mathbf{0}, \|\mathbf{x}\|_0 \leq k.$$

Unfortunately, we cannot develop a MISDO reformulation of non-negative sparse PCA *mutatis mutandis* Theorem 1. Indeed, while we can still set $\mathbf{X} = \mathbf{x}\mathbf{x}^\top$ and relax the rank-one constraint, if we do so then, by the non-negativity of \mathbf{x} , lifting \mathbf{x} yields:

$$\begin{aligned} \max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \in \mathcal{C}_n} \langle \boldsymbol{\Sigma}, \mathbf{X} \rangle \\ \text{s.t. } \text{tr}(\mathbf{X}) = 1, X_{i,j} = 0 \text{ if } z_i = 0, X_{i,j} = 0 \text{ if } z_j = 0 \forall i, j \in [p]. \end{aligned} \quad (24)$$

where $\mathcal{C}_n := \{\mathbf{X} : \exists \mathbf{U} \geq \mathbf{0}, \mathbf{X} = \mathbf{U}^\top \mathbf{U}\}$ denotes the completely positive cone, which is NP-hard to separate over and cannot currently be optimized over tractably (Dong and Anstreicher, 2013). Nonetheless, we can develop relatively tractable mixed-integer conic upper and lower bounds for non-negative sparse PCA. Indeed, we can obtain a fairly tight upper bound by replacing the completely positive cone with the larger doubly non-negative cone $\mathcal{D}_n := \{\mathbf{X} \in S_+^p : \mathbf{X} \geq \mathbf{0}\}$, which is a high-quality outer-approximation of \mathcal{C}_n , indeed exact when $k \leq 4$ (Burer et al., 2009).

Unfortunately, this relaxation is strictly different in general, since the extreme rays of the doubly non-negative cone are not necessarily rank-one when $k \geq 5$ (Burer et al., 2009). Nonetheless, to obtain feasible solutions which supply lower bounds, we could inner approximate the completely positive cone with the cone of non-negative scaled diagonally dominant matrices (see Ahmadi and Majumdar, 2019; Bostanabad et al., 2020).

5.2 Sparse PCA on Rectangular Matrices

A second extension would be to extend our methodology to the non-square case:

$$\max_{\mathbf{x} \in \mathbb{R}^m, \mathbf{y} \in \mathbb{R}^n} \mathbf{x}^\top \mathbf{A} \mathbf{y} \text{ s.t. } \|\mathbf{x}\|_2 = 1, \|\mathbf{y}\|_2 = 1, \|\mathbf{x}\|_0 \leq k, \|\mathbf{y}\|_0 \leq k. \quad (25)$$

Observe that computing the spectral norm of a matrix \mathbf{A} is equivalent to:

$$\max_{\mathbf{X} \in \mathbb{R}^{n \times m}} \langle \mathbf{A}, \mathbf{X} \rangle \text{ s.t. } \begin{pmatrix} \mathbf{U} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{V} \end{pmatrix} \succeq \mathbf{0}, \text{tr}(\mathbf{U}) + \text{tr}(\mathbf{V}) = 2, \quad (26)$$

where, in an optimal solution, \mathbf{U} stands for $\mathbf{x}\mathbf{x}^\top$, \mathbf{V} stands for $\mathbf{y}\mathbf{y}^\top$ and \mathbf{X} stands for $\mathbf{x}\mathbf{y}^\top$ —this can be seen by taking the dual of (Recht et al., 2010, Equation 2.4).

Therefore, by using the same argument as in the positive semidefinite case, we can rewrite sparse PCA on rectangular matrices as the following MISDO:

$$\begin{aligned} \max_{\mathbf{w} \in \{0,1\}^m, \mathbf{z} \in \{0,1\}^n} \max_{\mathbf{X} \in \mathbb{R}^{n \times m}} \langle \mathbf{A}, \mathbf{X} \rangle \\ \text{s.t. } \begin{pmatrix} \mathbf{U} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{V} \end{pmatrix} \succeq \mathbf{0}, \text{tr}(\mathbf{U}) + \text{tr}(\mathbf{V}) = 2, \\ U_{i,j} = 0 \text{ if } w_i = 0 \ \forall i, j \in [m], \\ V_{i,j} = 0 \text{ if } z_i = 0 \ \forall i, j \in [n], \mathbf{e}^\top \mathbf{w} \leq k, \mathbf{e}^\top \mathbf{z} \leq k. \end{aligned} \quad (27)$$

5.3 Sparse PCA with Multiple Principal Components

A third extension where our methodology is applicable is the problem of obtaining multiple principal components simultaneously, rather than deflating Σ after obtaining each principal component. As there are multiple definitions of this problem, we now discuss the extent to which our framework encompasses each case.

Common Support: Perhaps the simplest extension of sparse PCA to a multi-component setting arises when all r principal components have common support. By retaining the vector of binary variables \mathbf{z} and employing the Ky-Fan theorem (c.f. Wolkowicz et al., 2012, Theorem 2.3.8) to cope with multiple principal components, we obtain the following formulation in much the same manner as previously:

$$\max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \in S_+^p} \langle \mathbf{X}, \Sigma \rangle \text{ s.t. } \mathbf{0} \preceq \mathbf{X} \preceq \mathbb{I}, \text{tr}(\mathbf{X}) = r, X_{i,j} = 0 \text{ if } z_i = 0 \ \forall i \in [p]. \quad (28)$$

Notably, the logical constraint $X_{i,j} = 0$ if $z_i = 0$, which formed the basis of our subproblem strategy, still successfully models the sparsity constraint. This suggests that (a) one can derive an equivalent subproblem strategy under common support, and (b) a cutting-plane method for common support should scale equally well as with a single component.

Disjoint Support: In a sparse PCA problem with disjoint support (Vu and Lei, 2012), simultaneously computing the first r principal components is equivalent to solving:

$$\begin{aligned} \max_{\mathbf{z} \in \{0,1\}^{p \times r}: \mathbf{e}^\top \mathbf{z}_t \leq k \ \forall t \in [r], \mathbf{z} \mathbf{e} \leq \mathbf{e}} \max_{\mathbf{W} \in \mathbb{R}^{p \times r}} \langle \mathbf{W}\mathbf{W}^\top, \Sigma \rangle \\ \mathbf{W}^\top \mathbf{W} = \mathbb{I}_r, W_{i,j} = 0 \text{ if } z_{i,t} = 0 \ \forall i \in [p], t \in [r], \end{aligned} \quad (29)$$

where $z_{i,t}$ is a binary variable denoting whether feature i is a member of the t th principal component. By applying the technique used to derive Theorem 1 *mutatis mutandis*, and invoking the Ky-Fan theorem (c.f. Wolkowicz et al., 2012, Theorem 2.3.8) to cope with the rank- r constraint, we obtain:

$$\begin{aligned} \max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \in S^p} \langle \mathbf{X}, \boldsymbol{\Sigma} \rangle \\ \mathbf{0} \preceq \mathbf{X} \preceq \mathbb{I}, \operatorname{tr}(\mathbf{X}) = r, X_{i,j} = 0 \text{ if } Y_{i,j} = 0 \forall i \in [p], \end{aligned} \quad (30)$$

where $Y_{i,j} = \sum_{t=1}^r z_{i,t} z_{j,t}$ is a binary matrix denoting whether features i and j are members of the same principal component; this problem can be addressed by a cutting-plane method in much the same manner as when $r = 1$.

Acknowledgments

We are grateful to the three anonymous referees and the associate editor for many valuable comments which improved the paper.

References

- A. A. Ahmadi and A. Majumdar. DSOS and SDSOS optimization: More tractable alternatives to sum of squares and semidefinite optimization. *SIAM Journal on Applied Algebra and Geometry*, 3(2):193–230, 2019.
- A. A. Amini and M. J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. In *2008 IEEE international symposium on information theory*, pages 2454–2458. IEEE, 2008.
- A. Atamtürk and A. Gomez. Safe screening rules for l0-regression. *Optimization Online*, 2020.
- X. Bao, N. V. Sahinidis, and M. Tawarmalani. Semidefinite relaxations for quadratically constrained quadratic programming: A review and comparisons. *Mathematical Programming*, 129(1):129, 2011.
- B. Barak, J. A. Kelner, and D. Steurer. Rounding sum-of-squares relaxations. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 31–40, 2014.
- C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. Savelsbergh, and P. H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46(3):316–329, 1998.
- A. Ben-Tal and A. Nemirovski. On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty. *SIAM Journal on Optimization*, 12(3):811–833, 2002.
- L. Berk and D. Bertsimas. Certifiably optimal sparse principal component analysis. *Mathematical Programming Computation*, 11(3):381–420, 2019.
- Q. Berthet and P. Rigollet. Optimal detection of sparse principal components in high dimension. *Annals of Statistics*, 41(4):1780–1815, 2013.
- D. Bertsimas and R. Cory-Wright. On polyhedral and second-order-cone decompositions of semidefinite optimization problems. *Operations Research Letters*, 48(1):78–85, 2020.

- D. Bertsimas and B. Van Parys. Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *Annals of Statistics*, 48(1):300–323, 2020.
- D. Bertsimas and Y. Ye. Semidefinite relaxations, multivariate normal distributions, and order statistics. In *Handbook of Combinatorial Optimization*, pages 1473–1491. Springer, 1998.
- D. Bertsimas, J. Pauphilet, and B. Van Parys. Sparse classification and phase transitions: A discrete optimization perspective. *arXiv preprint arXiv:1710.01352*, 2017.
- D. Bertsimas, R. Cory-Wright, and J. Pauphilet. A unified approach to mixed-integer optimization problems with logical constraints. *arXiv preprint arXiv:1907.02109*, 2019.
- E. G. Boman, D. Chen, O. Parekh, and S. Toledo. On factor width and symmetric h-matrices. *Linear Algebra and its Applications*, 405:239–248, 2005.
- M. S. Bostanabad, J. Gouveia, and T. K. Pong. Inner approximating the completely positive cone via the cone of scaled diagonally dominant matrices. *Journal on Global Optimization*, 76:383–405, 2020.
- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, Cambridge, 2004.
- S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. Studies in Applied Mathematics, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1994.
- A. Brauer. Limits for the characteristic roots of a matrix iv. *Duke Mathematical Journal*, 19:75–91, 1952.
- S. Burer, K. M. Anstreicher, and M. Dür. The difference between 5×5 doubly nonnegative and completely positive matrices. *Linear Algebra and its Applications*, 431(9):1539–1552, 2009.
- E. Candes and T. Tao. The Dantzig selector: Statistical estimation when p is much larger than n. *Annals of Statistics*, 35(6):2313–2351, 2007.
- S. O. Chan, D. Papailiopoulos, and A. Rubinfeld. On the approximability of sparse PCA. In *Conference on Learning Theory*, pages 623–646, 2016.
- Y. Chen, Y. Ye, and M. Wang. Approximation hardness for a class of sparse optimization problems. *Journal of Machine Learning Research*, 20(38):1–27, 2019.
- A. Chowdhury, P. Drineas, D. P. Woodruff, and S. Zhou. Approximation algorithms for sparse principal component analysis. *arXiv preprint arXiv:2006.12748*, 2020.
- A. d’Aspremont, L. E. Ghaoui, M. I. Jordan, and G. R. Lanckriet. A direct formulation for sparse pca using semidefinite programming. In *Advances in neural information processing systems*, pages 41–48, 2005.
- A. d’Aspremont, F. Bach, and L. E. Ghaoui. Optimal solutions for sparse principal component analysis. *Journal of Machine Learning Research*, 9(Jul):1269–1294, 2008.
- S. S. Dey, R. Mazumder, M. Molinaro, and G. Wang. Sparse principal component analysis and its l_1 -relaxation. *arXiv preprint arXiv:1712.00800*, 2017.
- S. S. Dey, R. Mazumder, and G. Wang. A convex integer programming approach for optimal sparse PCA. *arXiv preprint arXiv:1810.09062*, 2018.

- H. Dong and K. Anstreicher. Separating doubly nonnegative and completely positive matrices. *Mathematical Programming*, 137(1-2):131–153, 2013.
- H. Dong, K. Chen, and J. Linderoth. Regularization vs. relaxation: A conic optimization perspective of statistical variable selection. *arXiv:1510.06083*, 2015.
- M. A. Duran and I. E. Grossmann. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Mathematical Programming*, 36(3):307–339, 1986.
- A. d’Aspremont and S. Boyd. Relaxations and randomized methods for nonconvex QCQPs. *EE392o Class Notes, Stanford University*, 1:1–16, 2003.
- A. d’Aspremont, F. Bach, and L. El Ghaoui. Approximation bounds for sparse principal component analysis. *Mathematical Programming*, 148(1-2):89–110, 2014.
- K. Fountoulakis, A. Kundu, E.-M. Kontopoulou, and P. Drineas. A randomized rounding algorithm for sparse PCA. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 11(3):1–26, 2017.
- T. Gally and M. E. Pfetsch. Computing restricted isometry constants via mixed-integer semidefinite programming. *Optimization Online*, 2016.
- T. Gally, M. E. Pfetsch, and S. Ulbrich. A framework for solving mixed-integer semidefinite programs. *Optimization Methods & Software*, 33(3):594–632, 2018.
- M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- M. Hein and T. Bühler. An inverse power method for nonlinear eigenproblems with applications in 1-spectral clustering and sparse PCA. In *Advances in neural information processing systems*, pages 847–855, 2010.
- R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 1990.
- H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- Y.-L. Hsu, P.-Y. Huang, and D.-T. Chen. Sparse principal component analysis in cancer research. *Translational cancer research*, 3(3):182, 2014.
- J. N. Jeffers. Two case studies in the application of principal component analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 16(3):225–236, 1967.
- I. T. Jolliffe. Rotation of principal components: choice of normalization constraints. *Journal of Applied Statistics*, 22(1):29–35, 1995.
- I. T. Jolliffe, N. T. Trendafilov, and M. Uddin. A modified principal component technique based on the lasso. *Journal of computational and Graphical Statistics*, 12(3):531–547, 2003.
- M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *Journal of Machine Learning Research*, 11(2), 2010.
- H. F. Kaiser. The varimax criterion for analytic rotation in factor analysis. *Psychometrika*, 23(3):187–200, 1958.

- S. Kim and M. Kojima. Second order cone programming relaxation of nonconvex quadratic optimization problems. *Optimization methods and software*, 15(3-4):201–224, 2001.
- K. Kobayashi and Y. Takano. A branch-and-cut algorithm for solving mixed-integer semidefinite optimization problems. *Computational Optimization & Applications*, 75(2):493–513, 2020.
- O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411, 2004.
- B. Li and M. Tsatsomeros. Doubly diagonally dominant matrices. *Linear Algebra and Its Applications*, 261(1-3):221–235, 1997.
- Y. Li and W. Xie. Exact and approximation algorithms for sparse PCA. *arXiv preprint arXiv:2008.12438*, 2020.
- R. Luss and A. d’Aspremont. Clustering and feature selection using sparse principal component analysis. *Optimization & Engineering*, 11(1):145–157, 2010.
- R. Luss and M. Teboulle. Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint. *SIAM Review*, 55(1):65–98, 2013.
- M. Magdon-Ismail. NP-hardness and inapproximability of sparse PCA. *Information Processing Letters*, 126:35–38, 2017.
- A. Majumdar, G. Hall, and A. A. Ahmadi. A survey of recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *arXiv preprint arXiv:1908.05209*, 2019.
- L. Miolane. Phase transitions in spiked matrix estimation: information-theoretic analysis. *arXiv preprint arXiv:1806.04343*, 2018.
- B. Moghaddam, Y. Weiss, and S. Avidan. Spectral bounds for sparse pca: Exact and greedy algorithms. In *Advances in neural information processing systems*, pages 915–922, 2006.
- I. Quesada and I. E. Grossmann. An LP/NLP based branch and bound algorithm for convex MINLP optimization problems. *Computers & Chemical Engineering*, 16(10-11):937–947, 1992.
- P. Raghavan and C. D. Tompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
- P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 755–764, 2010.
- B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- M. B. Richman. Rotation of principal components. *Journal of climatology*, 6(3):293–335, 1986.
- P. Richtárik, M. Jahani, S. D. Ahipaşaoğlu, and M. Takáč. Alternating maximization: Unifying framework for 8 sparse PCA formulations and efficient parallel codes. *Optimization and Engineering*, pages 1–27, 2020.
- M. J. Todd. On max-k-sums. *Mathematical Programming*, 171(1):489–517, 2018.

- V. Vu and J. Lei. Minimax rates of estimation for sparse PCA in high dimensions. In *Artificial intelligence and statistics*, pages 1278–1286, 2012.
- T. Wang, Q. Berthet, and R. J. Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *Annals of Statistics*, 44(5):1896–1930, 2016.
- D. M. Witten, R. Tibshirani, and T. Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.
- H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of semidefinite programming: theory, algorithms, and applications*, volume 27. Springer Science & Business Media, 2012.
- X.-T. Yuan and T. Zhang. Truncated power method for sparse eigenvalue problems. *Journal of Machine Learning Research*, 14(Apr):899–925, 2013.
- G. Zakeri, D. Craigie, A. Philpott, and M. Todd. Optimization of demand response through peak shaving. *Oper. Res. Letters*, 42(1):97–101, 2014.
- R. Zass and A. Shashua. Nonnegative sparse PCA. In *Advances in neural information processing systems*, pages 1561–1568, 2007.
- Y. Zhang, A. d’Aspremont, and L. El Ghaoui. Sparse PCA: Convex relaxations, algorithms and applications. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, pages 915–940. Springer, 2012.
- H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.

Appendix A. A Doubly Non-Negative Relaxation and a Goemans-Williamson Rounding Scheme

The MISDO formulation (5) we derived in Section 2 features big- M constraints of the form $|X_{i,j}| \leq M_{i,j}z_i$. We did not include the equally valid inequalities $|X_{i,j}| \leq M_{i,j}z_j$, because they are redundant with the fact that \mathbf{X} is symmetric. Actually, (5) is equivalent to

$$\max_{\mathbf{z} \in \{0,1\}^p: \mathbf{e}^\top \mathbf{z} \leq k} \max_{\mathbf{X} \in S_+^p} \langle \mathbf{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, |X_{i,j}| \leq M_{i,j}z_i z_j, \forall i, j \in [p]. \quad (31)$$

The formulation above features products of binary variables $z_i z_j$. Therefore, unlike several other problems involving cardinality constraints such as compressed sensing, relaxations of sparse PCA benefit from invoking an optimization hierarchy (see d’Aspremont and Boyd, 2003, Section 2.4.1, for a counterexample specific to compressed sensing). In particular, let us model the outer product $\mathbf{z}\mathbf{z}^\top$ by introducing a matrix \mathbf{Z} and imposing the semidefinite constraint $\mathbf{Z} \succeq \mathbf{z}\mathbf{z}^\top$. We tighten the formulation by requiring that $Z_{i,i} = z_i$ and imposing the linear inequalities $\max(z_i + z_j - 1, 0) \leq Z_{i,j} \leq \min(z_i, z_j)$. Hence, we obtain:

$$\max_{\substack{\mathbf{z} \in [0,1]^p: \mathbf{e}^\top \mathbf{z} \leq k, \\ \mathbf{Z} \in \mathbb{R}_+^{p \times p}}} \max_{\mathbf{X} \succeq \mathbf{0}} \langle \mathbf{\Sigma}, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, |X_{i,j}| \leq M_{i,j}Z_{i,j}, \langle \mathbf{E}, \mathbf{Z} \rangle \leq k^2, Z_{i,i} = z_i, \quad (32)$$

$$\max(z_i + z_j - 1, 0) \leq Z_{i,j} \leq \min(z_i, z_j), \begin{pmatrix} 1 & \mathbf{z}^\top \\ \mathbf{z} & \mathbf{Z} \end{pmatrix} \succeq \mathbf{0}.$$

Problem (32) is a doubly non-negative relaxation, as we have intersected the Shor and RLT relaxations. This is noteworthy, because doubly non-negative relaxations dominate most other popular relaxations with $O(p^2)$ variables (Bao et al., 2011, Theorem 1).

Relaxation (32) is amenable to a *Goemans-Williamson* rounding scheme (Goemans and Williamson, 1995). Namely, let $(\mathbf{z}^*, \mathbf{Z}^*)$ denote optimal choices of (\mathbf{z}, \mathbf{Z}) in Problem (32), $\hat{\mathbf{z}}$ be normally distributed random vector such that $\hat{\mathbf{z}} \sim \mathcal{N}(\mathbf{z}^*, \mathbf{Z}^* - \mathbf{z}^* \mathbf{z}^{*\top})$, and $\bar{\mathbf{z}}$ be a rounding of the vector such that $\bar{z}_i = 1$ for the k largest entries of \hat{z}_i ; this is, up to feasibility on $\hat{\mathbf{z}}$, equivalent to the hyperplane rounding scheme of Goemans and Williamson (1995) (see Bertsimas and Ye, 1998, for a proof). We formalize this procedure in Algorithm 3. As Algorithm 3 returns one of multiple possible $\bar{\mathbf{z}}$'s, a computationally useful strategy is to run the random rounding component several times and return the best solution.

Algorithm 3 A Goemans-Williamson rounding method for Problem (1)

Require: Covariance matrix Σ , sparsity parameter k

Compute $\mathbf{z}^*, \mathbf{Z}^*$ solution of (32)

Compute $\hat{\mathbf{z}} \sim \mathcal{N}(\mathbf{z}^*, \mathbf{Z}^* - \mathbf{z}^* \mathbf{z}^{*\top})$

Construct $\bar{\mathbf{z}} \in \{0, 1\}^p : \mathbf{e}^\top \bar{\mathbf{z}} = k$ such that $\bar{z}_i \geq \bar{z}_j$ if $\hat{z}_i \geq \hat{z}_j$.

Compute \mathbf{X} solution of

$$\max_{\mathbf{X} \in S_+^p} \langle \Sigma, \mathbf{X} \rangle \text{ s.t. } \text{tr}(\mathbf{X}) = 1, X_{i,j} = 0 \text{ if } \bar{z}_i \bar{z}_j = 0 \forall i, j \in [p].$$

return \mathbf{z}, \mathbf{X} .

A very interesting question is whether it is possible to produce a constant factor guarantee on the quality of Algorithm 3's rounding, as Goemans and Williamson (1995) successfully did for binary quadratic optimization. Unfortunately, despite our best effort, this does not appear to be possible as the quality of the rounding depends on the value of the optimal dual variables, which are hard to control in this setting. This should not be too surprising for two distinct reasons. Namely, (a) sparse regression, which reduces to sparse PCA (see d'Aspremont et al., 2008, Section 6.1) is strongly NP-hard (Chen et al., 2019), and (b) sparse PCA is hard to approximate within a constant factor under the Small Set Expansion (SSE) hypothesis (Chan et al., 2016), meaning that producing a constant factor guarantee would contradict the SSE hypothesis of Raghavendra and Steurer (2010).

We close this appendix by noting that a similar in spirit (although different in both derivation and implementation) combination of taking a semidefinite relaxation of $\mathbf{z} \in \{0, 1\}^p$ and rounding *à la* Goemans-Williamson has been proposed for sparse regression problems (Dong et al., 2015).