LBS Research Online

D Bertsimas, R Cory-Wright and [J Pauphilet](#)

Mixed-Projection Conic Optimization: A New Paradigm for Modeling Rank Constraints
Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/1923/>

Bertsimas, D, Cory-Wright, R and [Pauphilet, J](#)

(2021)

Mixed-Projection Conic Optimization: A New Paradigm for Modeling Rank Constraints.

Operations Research.

ISSN 0030-364X

(In Press)

DOI: <https://doi.org/10.1287/opre.2021.2182>

INFORMS (Institute for Operations Research and Management Sciences)

<https://pubsonline.informs.org/doi/10.1287/opre.20...>

Mixed-Projection Conic Optimization: A New Paradigm for Modeling Rank Constraints

Dimitris Bertsimas

Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, USA.
ORCID: 0000-0002-1985-1003
dbertsim@mit.edu

Ryan Cory-Wright

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, USA,
ORCID: 0000-0002-4485-0619
ryancw@mit.edu

Jean Pauphilet

London Business School, London, UK,
ORCID: 0000-0001-6352-0984
jpauphilet@london.edu

We propose a framework for modeling and solving low-rank optimization problems to certifiable optimality. We introduce symmetric projection matrices that satisfy $\mathbf{Y}^2 = \mathbf{Y}$, the matrix analog of binary variables that satisfy $z^2 = z$, to model rank constraints. By leveraging regularization and strong duality, we prove that this modeling paradigm yields tractable convex optimization problems over the non-convex set of orthogonal projection matrices. Furthermore, we design outer-approximation algorithms to solve low-rank problems to certifiable optimality, compute lower bounds via their semidefinite relaxations, and provide near optimal solutions through rounding and local search techniques. We implement these numerical ingredients and, for the first time, solve low-rank optimization problems to certifiable optimality. Our algorithms also supply certifiably near-optimal solutions for larger problem sizes and outperform existing heuristics, by deriving an alternative to the popular nuclear norm relaxation which generalizes the perspective relaxation from vectors to matrices. Using currently available spatial branch-and-bound codes, not tailored to projection matrices, we can scale our exact (resp. near-exact) algorithms to matrices with up to 30 (resp. 600) rows/columns. All in all, our framework, which we name *Mixed-Projection Conic Optimization*, solves low-rank problems to certifiable optimality in a tractable and unified fashion.

Key words: rank minimization, semidefinite optimization, global optimization, discrete optimization, outer-approximation, regularization, perspective relaxation, matrix completion, nuclear norm

Subject classifications: programming: non-linear: quadratic, algorithms, applications

Area of review: Optimization

1. Introduction

Many central problems in optimization, machine learning, and control theory are equivalent to optimizing a low-rank matrix over a convex set. For instance, low-rank constraints successfully model notions of minimal complexity, low dimensionality, or orthogonality in a system. However, while rank constraints offer unparalleled modeling flexibility, no generic code currently solves these problems to certifiable optimality at even moderate sizes. This state of affairs has led influential works on low-rank optimization

(Candès and Plan 2010, Recht et al. 2010) to characterize low-rank optimization as intractable and advocate convex relaxations or heuristics which do not enjoy assumption-free optimality guarantees.

The manner in which rank constrained optimization is regarded today is reminiscent of how mixed-integer conic optimization (MICO), which can model NP-complete problems, was originally considered. After decades of research effort, however, algorithms and software for MICO are now widely available (see, e.g., Bonami et al. 2008, Coey et al. 2020) and solve large instances of disparate non-convex problems such as best subset selection (Bertsimas and Van Parys 2020) or sparse portfolio selection (Frangioni and Gentile 2007, Zheng et al. 2014, Bertsimas and Cory-Wright 2018) to certifiable optimality. Unfortunately, rank constraints cannot be represented using mixed-integer convex optimization (Lubin et al. 2021, Lemma 4.1) and do not benefit from these advances.

In this work, we characterize the complexity of rank constrained optimization and propose a new, more general framework, which we term *Mixed-Projection Conic Optimization* (MPCO). Our proposal generalizes MICO, by replacing binary variables z which satisfy $z^2 = z$ with symmetric orthogonal projection matrices $\mathbf{Y}$ which satisfy $\mathbf{Y}^2 = \mathbf{Y}$, and offers the following advantages over existing state-of-the-art methods: First, it supplies certificates of (near) optimality for low-rank problems. Second, it demonstrates that some of the best ideas in MICO, such as decomposition methods, cutting-planes, relaxations, and random rounding schemes, admit straightforward extensions to MPCO. Finally, we implement a near-optimal rounding strategy and a globally optimal cutting-plane algorithm that improve upon the state-of-the-art for matrix completion and sensor location problems. We hope that MPCO gives rise to exciting new challenges for the optimization community to tackle.

1.1. Scope of the Framework

Formally, we consider the problem:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \langle \mathbf{C}, \mathbf{X} \rangle + \lambda \cdot \text{Rank}(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{A}\mathbf{X} = \mathbf{B}, \text{Rank}(\mathbf{X}) \leq k, \mathbf{X} \in \mathcal{K}, \quad (1)$$

where λ (resp. k) prices (bounds) the rank of $\mathbf{X}$, $(\mathbf{A}, \mathbf{B}) \in \mathbb{R}^{\ell \times n} \times \mathbb{R}^{\ell \times m}$ defines an affine subspace, and $\mathcal{K}$ is a proper cone in the sense of Boyd and Vandenberghe (2004), i.e., closed, convex, solid and pointed. Observe that Problem (1) offers significant modeling flexibility, as it allows arbitrary conic constraints on $\mathbf{X}$. As a result, linear, convex quadratic, semidefinite, exponential, and power constraints and objectives can be captured by letting $\mathcal{K}$ be an appropriate product of the non-negative orthant and the second order, semidefinite, exponential, and power cones.

We now introduce our notation and present some central problems from the optimization and machine learning literature which admit low-rank formulations and fall within our framework.

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 $[n]$ denote the set of running indices $\{1, \dots, n\}$. 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.

We also use an assortment of matrix operators. We let $\sigma_i(\mathbf{X})$ denote the i th largest singular value of a matrix $\mathbf{X}$, $\langle \cdot, \cdot \rangle$ denote the Euclidean inner product between two vectors or matrices of the same dimension, $\mathbf{X}^\dagger$ denote the Moore-Penrose pseudoinverse of a matrix $\mathbf{X}$, $\|\cdot\|_F$ denote the Frobenius norm of a matrix, $\|\cdot\|_\sigma$ denote the spectral norm of a matrix, and $\|\cdot\|_*$ denote the nuclear norm of a matrix; see Horn and Johnson (1985) for a general theory of matrix operators.

Finally, we use a wide variety of convex cones. We let S^n denote the $n \times n$ cone of symmetric matrices, and S_+^n denote the $n \times n$ positive semidefinite cone.

1.1.1. Low-Rank Matrix Completion Given a sub-sample $(A_{i,j} : (i,j) \in \mathcal{I} \subseteq [n] \times [m])$ of a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, the matrix completion problem is to recover the entire matrix, by assuming $\mathbf{A}$ is low rank and seeking a rank- k matrix $\mathbf{X}$ which approximately fits the observed values. This problem arises in recommender system applications at Netflix and Amazon and admits the formulation:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \frac{1}{2} \sum_{(i,j) \in \mathcal{I}} (X_{i,j} - A_{i,j})^2 \quad \text{s.t.} \quad \text{Rank}(\mathbf{X}) \leq k. \quad (2)$$

Since there are $(n+m)k$ degrees of freedom in a singular value decomposition of a rank- k matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$, Problem (2) is not well-defined unless $|\mathcal{I}| \geq (n+m)k$.

1.1.2. Minimum Dimension Euclidean Distance Embedding Given a set of pairwise distances $d_{i,j}$, the Euclidean Distance Embedding (EDM) problem is to determine the lowest dimensional space which the distances can be embedded in, such that the distances correspond to Euclidean distances. As discussed by Blekherman et al. (2012) Theorem 2.49, a set of distances $d_{i,j}$ can be embedded in a Euclidean space of dimension k if and only if there exists some Gram matrix $\mathbf{G} \succeq \mathbf{0}$ of rank k such that $d_{i,j}^2 = G_{i,i} + G_{j,j} - 2G_{i,j}$, on all pairs (i,j) where $d_{i,j}$ is supplied. Denoting $D_{i,j} = d_{i,j}^2$, we write these constraints in matrix form, $\mathbf{D} = \text{Diag}(\mathbf{G})\mathbf{e}^\top + \mathbf{e}\text{Diag}(\mathbf{G})^\top - 2\mathbf{G}$, where the equality is implicitly imposed only for pairs (i,j) where $d_{i,j}$ is supplied. This is equivalent to:

$$\min_{\mathbf{G} \in S_+^n} \text{Rank}(\mathbf{G}) \quad \text{s.t.} \quad \text{Diag}(\mathbf{G})\mathbf{e}^\top + \mathbf{e}\text{Diag}(\mathbf{G})^\top - 2\mathbf{G} = \mathbf{D}. \quad (3)$$

1.1.3. Quadratically Constrained Quadratic Optimization A quadratically constrained quadratic optimization problem (QCQO) seeks an $\mathbf{x} \in \mathbb{R}^n$ which solves:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \mathbf{x}^\top \mathbf{Q}_0 \mathbf{x} + \mathbf{q}_0^\top \mathbf{x} \quad \text{s.t.} \quad \mathbf{x}^\top \mathbf{Q}_i \mathbf{x} + \mathbf{q}_i^\top \mathbf{x} \leq r_i \quad \forall i \in [m], \quad (4)$$

where $\mathbf{Q}_0, \mathbf{Q}_i, \mathbf{q}_0, \mathbf{q}_i, r_i$ are given problem data. We assume that $\mathbf{Q}_0, \mathbf{Q}_i$ are symmetric matrices, but do not assume that they are positive semidefinite. Therefore, this problem is non-convex, and encompasses binary quadratic optimization (Goemans and Williamson 1995) and alternating current optimal power

flow problems (Lavai and Low 2011). The fundamental difficulty in Problem (4) is the potential non-convexity of the outer product $\mathbf{x}\mathbf{x}^\top$. However, we can isolate this non-convexity by introducing a rank-one matrix $\mathbf{X}$ to model the outer product $\mathbf{x}\mathbf{x}^\top$. This leads to the following reformulation:

$$\min_{\mathbf{x} \in \mathbb{R}^n, \mathbf{X} \in S^n} \langle \mathbf{Q}_0, \mathbf{X} \rangle + \langle \mathbf{q}_0, \mathbf{x} \rangle \quad \text{s.t.} \quad \langle \mathbf{Q}_i, \mathbf{X} \rangle + \langle \mathbf{q}_i, \mathbf{x} \rangle \leq r_i \quad \forall i \in [m], \quad \text{Rank} \begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{X} \end{pmatrix} = 1. \quad (5)$$

We have established that QCQOs are rank constrained problems. Notably however, the converse is also true: rank constrained problems with linear, second-order cone, or semidefinite constraints are QCQOs. Indeed, the constraint $\text{Rank}(\mathbf{X}) \leq k$ is equivalent to requiring that $\mathbf{X} = \mathbf{U}\mathbf{V}^\top : \mathbf{U} \in \mathbb{R}^{n \times k}, \mathbf{V} \in \mathbb{R}^{m \times k}$, i.e., imposing $m \times n$ non-convex quadratic equalities. As modern solvers such as **Gurobi** can now solve non-convex QCQOs to global optimality, this QCQO formulation can be used to solve low-rank problems, although it is not particularly scalable; we expand on this point in Section 4.1.2.

1.2. Background and Literature Review

Our work arises at the intersection of three complementary areas of the low-rank optimization literature: (a) global optimization algorithms for non-convex quadratically constrained problems, (b) the interplay of convex relaxations and their dual side, randomized rounding methods, and (c) heuristics which provide high-quality solutions to non-convex problems in an efficient fashion.

1.2.1. Global Optimization Techniques

Branch-and-bound: A broad class of global optimization algorithms have been proposed for QCQOs, since McCormick (1976) observed that convex envelopes of non-convex regions supply globally valid lower bounds. This gives rise to a numerical strategy where one recursively partitions the QCQO's feasible region into subregions, constructs convex envelopes for each subregion and uses these envelopes to construct iteratively improving lower bounds. This approach is known as spatial branch-and-bound; see Lee and Zou (2014) for a scheme which decomposes a matrix into a sparse matrix plus a low-rank matrix, Kocuk et al. (2016) for a modern implementation in alternating current optimal power flow, and Bertsimas et al. (2017) for an exact branch-and-bound approach to low-rank factor analysis.

Branch-and-cut: In a complementary direction, several branch-and-cut methods (Audet et al. 2000, Linderoth 2005) have been proposed for solving non-convex QCQOs, by borrowing decomposition schemes from the mixed-integer nonlinear optimization literature (Duran and Grossmann 1986). While often efficient in practice, a common theme in these methods is that the more efficient decomposition schemes used for MINLOs cannot be applied out-of-the-box, because they may fail to converge to a globally optimal solution (see Grossmann 2002, for a counterexample). As a result, non-convex problems need to be preprocessed in an expensive fashion. This preprocessing step has inhibited the use of global optimization methods for low-rank problems; indeed, we are not aware of any works which apply branch-and-cut techniques to solve low-rank problems to certifiable optimality.

Complementarity: In an opposite direction, several authors have proposed applying general nonlinear optimization techniques to address low-rank problems, since Ding et al. (2014) observed that a low-rank constraint is equivalent to a complementarity constraint over the positive semidefinite cone, and thus can be addressed by general techniques for mathematical programs with equilibrium constraints (see Luo et al. 1996, for a general theory). Among others, Bai et al. (2016) invoked the complementarity observation to design a completely positive reformulation of low-rank SDOs, and Bi et al. (2020) developed a multi-stage convex relaxation of the complementarity constraint. We are not aware of any works which use these ideas to solve low-rank problems exactly where say $n \geq 10$.

Algebraic: By taking an algebraic view of rank constraints, several algebraic geometry techniques have been proposed for addressing low-rank SDOs. Among others, d’Aspremont (2003) proposed reformulating low-rank constraints as systems of polynomial equations which can be addressed via the sum-of-squares hierarchy (Lasserre 2001). More recently, Naldi (2018) proposed a semi-algebraic reformulation of rank-constrained SDOs, which can be optimized over via Gröbner basis computation (Cox et al. 2013). Unfortunately, algebraic approaches do not scale well in practice. Indeed, as observed by Recht et al. (2010), it seems unlikely that algebraic approaches can solve low-rank SDOs when $n > 10$.

1.2.2. Convex Relaxations and Random Rounding Methods for Low-Rank Problems

Convex relaxations: A number of authors have studied convex relaxations of low-rank problems, since Fazel (2002) observed that the nuclear norm of a matrix is the convex envelope of a rank constraint on the set of matrices with spectral norm at most M , i.e.,

$$\text{Conv}(\{\mathbf{X} \in \mathbb{R}^{n \times m} : \|\mathbf{X}\|_{\sigma} \leq M, \text{Rank}(\mathbf{X}) \leq k\}) = \{\mathbf{X} \in \mathbb{R}^{n \times m} : \|\mathbf{X}\|_{\sigma} \leq M, \|\mathbf{X}\|_* \leq kM\}. \quad (6)$$

Because the epigraph of a nuclear norm is semidefinite representable, this gives rise to semidefinite relaxations of low-rank problems which can be computed in polynomial time.

Rounding methods: A complementary line of work aims to supply certifiably near-optimal solutions to low-rank problems, by rounding their semidefinite relaxations. Initiated by Goemans and Williamson (1995) in the context of binary quadratic optimization, who established that randomly rounding an SDO relaxation supplies a 0.878-approximation, it has evolved into a successful framework for solving rank-one optimization problems; see Nemirovski et al. (1999) for a unified approach in the rank-one case. However, this line of work has a key drawback. Namely, existing rounding methods do not address rank- k problems such as matrix completion, due to the analytic difficulty of constructing a rounding mechanism which preserves both feasibility and near-optimality in the rank- k case.

1.2.3. Heuristic Methods Due to the computational difficulty of solving Problem (1) to certifiable optimality, and the analytic difficulty of deriving a high-quality randomized rounding procedure, a variety of heuristic methods have been proposed for solving Problem (1), originating with methods for solving low-rank linear matrix inequalities in the optimal control literature (Boyd et al. 1994).

Although slow and somewhat ad-hoc in their original implementations, heuristic methods were moved front-and-center by the works of Fazel (2002), Burer and Monteiro (2003, 2005). Fazel (2002) observed that low-rank positive semidefinite matrices lie on the boundary of the PSD cone, and used this observation to justify a “log-det” heuristic, where a rank minimization objective is replaced with the function $\log \det(\mathbf{X} + \delta \mathbb{I})$. Burer and Monteiro (2003, 2005) proposed implicitly modeling a rank constraint $\text{Rank}(\mathbf{X}) \leq k$ by applying the non-linear reformulation $\mathbf{X} = \mathbf{U}\mathbf{V}^\top$, where $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times k}$ and eliminating $\mathbf{X}$, to obtain a problem which is non-convex in $(\mathbf{U}, \mathbf{V})$. Although originally solved using augmented Lagrangian techniques, modern implementations of the Burer-Monterio heuristic typically use alternating minimization (Jain et al. 2013), to enhance the methods scalability and ensure convergence towards a second-order critical point.

The modern era of heuristics methods for low-rank matrix optimization was initiated by matrix completion and the Netflix competition (Bell and Koren 2007). The magnitude of the data made available, challenged the aforementioned methods and led to new techniques such as Wen et al. (2012), who proposed a nonlinear successive over-relaxation approach that scales to $10,000 \times 10,000$ matrices, Recht and Ré (2013), who proposed a stochastic gradient descent method which scales to $1,000,000 \times 1,000,000$ matrices; see Udell et al. (2016), Nguyen et al. (2019) for reviews of heuristic approaches.

1.3. Contributions and Structure

The key contributions of the paper are threefold: First, we propose using orthogonal projection matrices which satisfy $\mathbf{Y}^2 = \mathbf{Y}$, the matrix analogue of binary variables which satisfy $z^2 = z$, to model low-rank constraints via the non-linear equation $\mathbf{X} = \mathbf{Y}\mathbf{X}$. Under this lens, low-rank problems admit reformulations as optimization problems where some decision variables comprise a projection matrix. We term this family of problems *Mixed-Projection Conic Optimization* (MPCO), in reference to mixed-integer optimization. To our knowledge, our approach is the first mathematical framework which solves low-rank optimization problems to certifiable optimality. Second, by leveraging regularization and strong duality we rewrite low-rank optimization problems as saddle-point problems over the space of orthogonal projection matrices, and propose an outer-approximation method to solve the saddle-point problem to certifiable optimality. Third, by analyzing the saddle-point problem, we derive new convex relaxations and rounding schemes which provide certifiably near optimal solutions in polynomial time in theory and rapidly in practice. Using a generic spatial branch-and-bound code, we are able to solve low-rank optimization problems exactly for matrices with 30 rows and columns, and find near-exact solutions for matrices with up to 600 rows and columns. We believe branch-and-bound schemes tailored to the set of projection matrices and dedicated semi-definite codes for the relaxations have to potential to increase the numerical scalability of MPCO even further and constitute an exciting future research direction.

We note that the idea of using projection matrices to model low-rank constraints is not entirely new, as Peng and Xia (2005) have proposed reformulating k -means clustering as a semidefinite optimization

problem over the set of orthogonal projection matrices. However, our proposal has several key points of difference. Namely, (1) we consider optimizing over projection matrices directly, while Peng and Xia (2005) use projection matrices as a vehicle to derive semidefinite relaxations, and (2) we use projection matrices to solve low-rank optimization problems that do not admit mixed-integer reformulations, while k -means clustering certainly admits a mixed-integer reformulation (Grötschel and Wakabayashi 1989).

The rest of the paper is laid out as follows:

In Section 2, we show that projection matrices are a natural generalization of binary vectors to matrices. Inspired by a common tactic in cardinality constrained optimization, namely introducing binary variables to encode the support of the decision vector, we propose introducing a projection matrix to encode the image of the decision matrix and thereby model rank. We also investigate the complexity of low-rank optimization problems and show that rank minimization is as hard as the existential theory of the reals (i.e., deciding whether a semi-algebraic set is non-empty), and thus in PSPACE.

In Section 3, we derive the MPCO formulations of the aforementioned rank optimization problems. By introducing a constraint on the spectral norm of $\mathbf{X}$ or a penalty on its Frobenius norm - the matrix analogs of big- M constraints and perspective formulations (Günlük and Linderoth 2012) respectively, we leverage strong duality, reformulate Problem (1) as a saddle-point problem, and prove the resulting optimization problem admits a convex objective.

We propose numerical algorithms to solve these MPCO problem to provable (near) optimality in Section 4, by extending some of the most successful techniques from MICO. First, we propose an outer-approximation scheme for solving Problem (1) exactly. Then, we obtain valid lower-bounds from solving its convex relaxations and propose an alternating minimization algorithm to do so. In addition, we prove that a singular value decomposition (SVD) followed by greedily rounding the eigenvalues provides certifiably near-optimal solutions in polynomial time. Finally, we propose a local-search strategy to improve the quality of the greedily rounded solution.

In Section 5, we implement and numerically evaluate our proposed algorithms. On examples from matrix completion and sensor location, we demonstrate that methods proposed in this paper solve instances of Problem (1) to certifiable optimality in minutes for n in the tens. To our knowledge, our work is the first to demonstrate that moderately sized rank constrained problems can be solved to provable optimality in a tractable fashion. For n in the hundreds, our proposal scales and provides in minutes solutions of higher quality than existing heuristics, such as nuclear norm minimization.

2. From Cardinality to Rank: A Unifying Nonlinear Perspective

Low rank constraints $\text{Rank}(\mathbf{X}) \leq k$ are a natural generalization of cardinality constraints $\|\mathbf{x}\|_0 \leq k$ from vectors to matrices. Indeed, if $\mathbf{X}$ is a diagonal matrix then $\text{Rank}(\mathbf{X}) \leq k$ if and only if $\|\mathbf{X}\|_0 \leq k$, and more generally $\text{Rank}(\mathbf{X}) \leq k$ if and only if $\|\sigma(\mathbf{X})\|_0 \leq k$, where $\sigma(\mathbf{X})$ is the vector of singular values of $\mathbf{X}$. However, while cardinality and rank constraints are intimately linked, they are addressed using

different algorithms. Namely, we can solve cardinality constrained problems with 100,000s of variables to optimality (Bertsimas and Van Parys 2020), while low-rank problems are dramatically harder and have not yet been solved to certifiable optimality for $n > 10$ (Naldi 2018).

In our opinion, the difference between the community’s understanding of cardinality and rank constraints has arisen because of two algorithmic barriers. The first barrier is that rank constraints belong to a harder complexity class. The second barrier arises because cardinality constraints can be represented using binary variables, while rank constraints cannot (Lubin et al. 2021, Corollary 4.1). This presents a challenge for researchers, who have developed scalable methods for cardinality constraints by exploiting advances in mixed-integer conic optimization (MICO), but cannot use these advances to address rank constraints. In this section, we question these barriers by characterizing the complexity of low-rank problems and proposing a new framework for modeling rank that generalizes (MICO).

2.1. On the Complexity of Rank-Constrained Optimization

Existing studies of Problem (1) typically claim that it is intractable, and support this claim by proving that it is NP-hard, by reduction from an NP-complete problem such as Boolean linear programming (see, e.g., Vandenbergh and Boyd 1996, Section 7.3). In our opinion, this argument needs to be revisited, for two separate reasons. First, NP-hardness is a worst-case analysis statement. In practice, NP-hard optimization problems are often tractable. For instance, sparse regression can usually be solved to certifiable optimality with 100,000s of features in minutes (Bertsimas and Van Parys 2020). Second, there is no evidence that Problem (1) is even in NP. Indeed, Problem (1) cannot be represented using mixed-integer convex optimization (Lubin et al. 2021, Corollary 4.1), while all 21 of Karp’s NP-complete problems can, and the best-known algorithms for Problem (1) run in EXPTIME (Chistov and Grigoriev 1984, Naldi 2018).

In this section, we provide a more complete characterization of Problem (1)’s complexity than is currently available in the literature. First, we demonstrate that it belongs to a different class than NP. In particular, it is *existential theory of the reals*-hard ($\exists\mathbb{R}$ -hard; see Renegar (1992) for a general theory), i.e., as hard as any polynomial optimization problem, which implies that, if $\text{NP} \subsetneq \exists\mathbb{R}$, Problem (1) is strictly harder than NP-complete problems. Second, we prove that Problem (1) is actually in $\exists\mathbb{R}$.

We now demonstrate that Problem (1) is existential theory of the reals complete (i.e., $\exists\mathbb{R}$ -complete). We begin by reminding the reader of the definition of the $\exists\mathbb{R}$ complexity class (c.f. Schaefer 2013):

DEFINITION 1. A decision problem belongs to the existential theory of the reals complexity class if it reduces to deciding whether a statement “ $(\exists x_1, \dots, x_n) \phi(x_1, \dots, x_n)$ ” is true or false, where $\phi(\cdot)$ is a quantifier-free Boolean formula involving polynomials equalities and inequalities, for instance, deciding the emptiness of a semi-algebraic set. We say a problem is $\exists\mathbb{R}$ -hard if any problem in $\exists\mathbb{R}$ reduces to it.

Note that 3-SAT $\in \exists\mathbb{R}$, so $\text{NP} \subseteq \exists\mathbb{R}$, and any statement in $\exists\mathbb{R}$ can be decided in PSPACE (Canny 1988), so $\exists\mathbb{R} \subseteq \text{PSPACE}$. To establish that Problem (1) is $\exists\mathbb{R}$ hard, we require the following proposition, which is essentially a restatement of (Schaefer 2013, Theorem 3.1) in the language of optimization.

PROPOSITION 1. *Let $G := (V, E)$ be a graph, and $\ell(e)$ be the length of edge e . Then, deciding if G can be embedded in $\mathbb{R}^2$ is $\exists\mathbb{R}$ complete, even when all edges have unit length.*

By reducing Proposition 1’s planar embedding problem to a Euclidean Distance Embedding problem, we obtain the following result (proof deferred to Appendix A.1):

THEOREM 1. *Problem (1) is $\exists\mathbb{R}$ -hard.*

Theorem 1 demonstrates that Problem (1) is, from a traditional complexity theory perspective, at least as hard as any problem in $\exists\mathbb{R}$. However, its complexity status remains unresolved. Indeed, while Candès and Plan (2010) have observed that Problem (1) is in EXPTIME, it seems likely that $\exists\mathbb{R} \subset \text{EXPTIME}$. We now address this matter, by proving that if $\mathcal{K}$ represents the semidefinite cone then Problem (1) is in $\exists\mathbb{R}$, and hence $\exists\mathbb{R}$ -complete; note that the examples listed in Section 1.1 can all be rewritten as low-rank SDOs, so this result applies to all of the aforementioned examples (proof of theorem deferred to Appendix A.2).

THEOREM 2. *Let $\mathcal{K} = S_+^n$ denote the $n \times n$ positive semidefinite cone. Then, Problem (1) is in $\exists\mathbb{R}$, and hence $\exists\mathbb{R}$ -complete.*

REMARK 1. Since $\exists\mathbb{R} \subseteq \text{PSPACE} \subseteq \text{EXPTIME}$, this upper bound improves upon the EXPTIME bound on Problem (1)’s complexity stated by Recht et al. (2010), Candès and Plan (2010) among others. Moreover, it seems unlikely to us that this bound can be further improved without settling fundamental questions in complexity theory (e.g., characterizing NP vs. $\exists\mathbb{R}$ vs. PSPACE vs. EXPTIME).

REMARK 2. Imposing an additional integrality constraint dramatically changes the complexity of Problem (1). Indeed, under the constraint $\mathbf{X} \in \mathbb{Z}^{n \times n}$, we can use a reduction from Hilbert’s 10th problem to show that we *cannot* decide in finite time whether Problem (1)’s optimal objective is 0, even if we know the objective is binary; see Appendix A.3 for a proof.

2.2. Projection Matrices for Modeling Rank

As previously discussed, rank constraints can be seen as a generalization to the matrix case of cardinality constraints. For a vector $\mathbf{x} \in \mathbb{R}^n$, the cardinality constraint $\|\mathbf{x}\|_0 \leq k$ ensures that at most k coordinates of $\mathbf{x}$ are non-zero, and can be modeled by introducing a vector of binary variables since

$$\|\mathbf{x}\|_0 \leq k \iff \exists \mathbf{z} \in \{0, 1\}^n : \mathbf{e}^\top \mathbf{z} \leq k, \mathbf{x} = \mathbf{z} \circ \mathbf{x}, \quad (7)$$

where $\mathbf{z} \circ \mathbf{x}$ denotes the component-wise product of $\mathbf{z}$ and $\mathbf{x}$. Actually, non-linear constraints of the form “ $\mathbf{x} = \mathbf{z} \circ \mathbf{x}$ ” where $\mathbf{z}$ is binary and $\mathbf{x}$ is continuous occur in a variety of mixed-integer optimization problems, far beyond cardinality constrained optimization. Bertsimas et al. (2019) observed that such non-linear constraints “ $\mathbf{x} = \mathbf{z} \circ \mathbf{x}$ ” actually lead to tractable optimization problems, provided that the overall objective is appropriately regularized. In particular, big- M constraints (Glover 1975) and

perspective reformulations (Günlük and Linderoth 2012) can be seen as appropriate regularizers. By building upon this observation and the work of several other authors (Fischetti et al. 2017, Bertsimas and Van Parys 2020), they successfully solve cardinality constrained problems at scale via a combination of branch-and-cut, randomized rounding and heuristic methods.

Unfortunately, rank constraints cannot be modeled using mixed-integer convex optimization (Lubin et al. 2021, Corollary 4.1) and therefore MICO techniques cannot be applied “out-of-the-box” to address rank constraints. Therefore, we now propose a new framework to model rank in optimization problems. Instead of a binary vector $\mathbf{z}$ to encode the support of $\mathbf{x}$, we introduce a projection matrix $\mathbf{Y}$ to capture the column space of $\mathbf{X}$ and obtain a similar non-linear reformulation.

DEFINITION 2. A matrix $\mathbf{Y} \in \mathbb{R}^{n \times n}$ is called a projection matrix if it satisfies the equality $\mathbf{Y}^2 = \mathbf{Y}$. In addition, if $\mathbf{Y}$ is symmetric, $\mathbf{Y}$ is called an orthogonal projection matrix.

As symmetric matrices, orthogonal projection matrices are diagonalizable and their eigenvalues satisfy $\lambda_i^2 = \lambda_i$, i.e., are binary. As a result, the Moore-Penrose pseudoinverse of an orthogonal projection $\mathbf{Y}$ is $\mathbf{Y}$ itself ($\mathbf{Y} = \mathbf{Y}^\dagger$). In addition, since its eigenvalues are binary, the trace of $\mathbf{Y}$ equals the number of non-zero eigenvalues, i.e., $\text{Rank}(\mathbf{Y}) = \text{tr}(\mathbf{Y})$. We are now in a position to link projection matrices and rank constraints.

PROPOSITION 2. For any $\mathbf{X} \in \mathbb{R}^{n \times m}$, $\text{Rank}(\mathbf{X}) \leq k \iff \exists \mathbf{Y} \in \mathcal{Y}_n : \text{tr}(\mathbf{Y}) \leq k, \mathbf{X} = \mathbf{Y}\mathbf{X}$, where $\mathcal{Y}_n := \{\mathbf{P} \in \mathbb{R}^{n \times n} : \mathbf{P}^2 = \mathbf{P}\}$ is the set of $n \times n$ orthogonal projection matrices.

Proof of Proposition 2 We prove the two implications successively.

- Let $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, with $\mathbf{U} \in \mathbb{R}^{n \times k}$, $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$, $\mathbf{V} \in \mathbb{R}^{m \times k}$, be a singular value decomposition of $\mathbf{X}$ and define $\mathbf{Y} = \mathbf{U}(\mathbf{U}^\top \mathbf{U})^{-1} \mathbf{U}^\top = \mathbf{U}\mathbf{U}^\top$. By construction, $\mathbf{X} = \mathbf{Y}\mathbf{X}$, since $\mathbf{U}^\top \mathbf{U} = \mathbb{I}$. Moreover, $\text{tr}(\mathbf{Y}) = \text{rank}(\mathbf{Y}) = \text{rank}(\mathbf{X}) \leq k$.
- Since $\mathbf{X} = \mathbf{Y}\mathbf{X}$, $\text{rank}(\mathbf{X}) \leq \text{rank}(\mathbf{Y}) = \text{tr}(\mathbf{Y}) \leq k$. □

REMARK 3. In Proposition 2, the rank constraint is expressed via a trace constraint on $\mathbf{Y}$, the orthogonal projection onto the image or column space of $\mathbf{X}$. Alternatively, one could model the rank constraint via a matrix $\mathbf{Y}' \in \mathcal{Y}_m$ such that $\text{tr}(\mathbf{Y}') \leq k$ and $\mathbf{X} = \mathbf{X}\mathbf{Y}'$. In this case, $\mathbf{Y}'$ encodes the projection onto the row space of $\mathbf{X}$. In practice, one could introduce both $\mathbf{Y}$ and $\mathbf{Y}'$ and obtain tighter formulations, at the price of introducing additional notation. We explore this idea in Appendix B.

Proposition 2 suggests that projection matrices are to rank constraints what binary variables are to cardinality constraints. Indeed, similarities between the two are evident: binary variables z are idempotent scalars which solve $z^2 = z$, while projection matrices $\mathbf{Y}$ are idempotent matrices which solve $\mathbf{Y}^2 = \mathbf{Y}$. Also, if $\mathbf{X}$ and $\mathbf{Y}$ are diagonal, Proposition 2 recovers cardinality constrained optimization.

Over the past decades, extensive efforts have been devoted to improving the scalability of mixed-integer optimization. We believe that similar achievements can be obtained for rank constrained problems by adapting techniques from MICO to MPCO. In this direction, Table 1 establishes a dictionary

linking cardinality and rank constraints, and demonstrates that many of the techniques developed for binary convex optimization admit generalizations to MPCO, including the main results from our recent work (Bertsimas et al. 2019). Note that we have not yet established most of the connections claimed in Table 1; this is the focus of the next two sections of the paper.

Table 1 Analogy between mixed-integer conic and mixed-projection conic optimization.

Framework	Bertsimas et al. (2019)	This paper
Parsimony concept	cardinality	rank
Non-convex outer set	binaries	orthogonal projection matrices
Strongly convex regularizer	ℓ_2^2 penalty	Frobenius norm squared
Boundedness regularizer	ℓ_∞ norm	spectral norm
Non-linear formulation	$\mathbf{x} = \mathbf{x} \circ \mathbf{z}; \mathbf{z} \in \{0, 1\}^n$	$\mathbf{X} = \mathbf{Y} \mathbf{X}, \mathbf{Y} \in \mathcal{Y}_n$
Big-M formulation	$-M\mathbf{z} \leq \mathbf{x} \leq M\mathbf{z}$	$\begin{pmatrix} M\mathbf{Y} & \mathbf{X} \\ \mathbf{X}^\top & M\mathbb{I} \end{pmatrix} \succeq \mathbf{0}$
Perspective formulation	$\begin{pmatrix} \theta_i & x_i \\ x_i & z_i \end{pmatrix} \succeq \mathbf{0}$	$\begin{pmatrix} \boldsymbol{\theta} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{Y} \end{pmatrix} \succeq \mathbf{0}$
Convex relaxation complexity	linear/second-order cone	semidefinite
Greedy rounding mechanism	coordinate-wise	singular value decomposition

3. Regularization and a Saddle-Point Reformulation

In this section, we prove that (9) can be reformulated as a saddle-point mixed-projection problem by leveraging regularization terms analogous to the big- M and ridge regularization techniques from MICO, and derive their semidefinite relaxations, as summarized in Table 1.

Throughout this paper, we let $\mathcal{Y}_n := \{\mathbf{P} \in S^n : \mathbf{P}^2 = \mathbf{P}\}$ denote the set of $n \times n$ orthogonal projection matrices and $\mathcal{Y}_n^k := \{\mathbf{P} \in S^n : \mathbf{P}^2 = \mathbf{P}, \text{tr}(\mathbf{P}) \leq k\}$ denote projection matrices with rank at most k . Although $\mathcal{Y}_n$ and $\mathcal{Y}_n^k$ do not commonly appear in the optimization literature, their convex hulls are well-studied, as we now remind the reader, by restating (Overton and Womersley 1992, Theorem 3):

LEMMA 1. *Let $\mathcal{Y}_n$ denote the $n \times n$ orthogonal projection matrices and $\mathcal{Y}_n^k$ denote the low-rank orthogonal projection matrices. Then, $\text{Conv}(\mathcal{Y}_n) = \{\mathbf{P} : \mathbf{0} \preceq \mathbf{P} \preceq \mathbb{I}\}$ and $\text{Conv}(\mathcal{Y}_n^k) = \{\mathbf{P} : \mathbf{0} \preceq \mathbf{P} \preceq \mathbb{I}, \text{tr}(\mathbf{P}) \leq k\}$. Moreover, the extreme points of $\text{Conv}(\mathcal{Y}_n)$ are $\mathcal{Y}_n$, and the extreme points of $\text{Conv}(\mathcal{Y}_n^k)$ are $\mathcal{Y}_n^k$.*

3.1. A Regularization Assumption

By invoking Proposition 2, we rewrite Problem (1) as the following mixed-projection conic problem:

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \langle \mathbf{C}, \mathbf{X} \rangle + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} = \mathbf{Y}\mathbf{X}, \mathbf{X} \in \mathcal{K}. \quad (8)$$

Observe that Problem (8) has a two-stage structure which involves first selecting a low-rank projection matrix $\mathbf{Y}$ and second selecting a matrix $\mathbf{X}$ under the constraint $\mathbf{X} = \mathbf{Y}\mathbf{X}$. Moreover, selecting an optimal $\mathbf{X}$ given $\mathbf{Y}$ is *easy*, because it involves solving a conic optimization problem under the linear

constraint $\mathbf{X} = \mathbf{Y}\mathbf{X}$, while selecting an optimal $\mathbf{Y}$ is *hard*, because $\mathcal{Y}_n^k$ is a non-convex set. Therefore, our modeling framework isolates the hardness of Problem (8) in $\mathcal{Y}_n^k$.

To cope with the non-linear constraints $\mathbf{X} = \mathbf{Y}\mathbf{X}$ in a tractable fashion, we augment the objective function in (8) with a regularization term. Namely, we consider

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \langle \mathbf{C}, \mathbf{X} \rangle + \Omega(\mathbf{X}) + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} = \mathbf{Y}\mathbf{X}, \mathbf{X} \in \mathcal{K}. \quad (9)$$

where the regularization term $\Omega(\mathbf{X})$ satisfies the following assumption:

ASSUMPTION 1. *In Problem (9), the regularization term $\Omega(\mathbf{X})$ is one of:*

- *A spectral norm penalty, $\Omega(\mathbf{X}) = 0$ if $\|\mathbf{X}\|_\sigma \leq M$ and $\Omega(\mathbf{X}) = +\infty$ otherwise.*
- *A Frobenius norm penalty, $\Omega(\mathbf{X}) = \frac{1}{2\gamma} \|\mathbf{X}\|_F^2$.*

As we demonstrate in Section 3.2, Assumption 1 is crucial for developing efficient low-rank algorithms, for the regularizer drives the convexity (see Theorem 3) and smoothness (see Lemma 3) of the problem, and also make computationally cheap to evaluate subgradients readily accessible (Table 2). The idea of leveraging regularization to optimize possibly non-smooth functions by obtaining computationally useful subgradients is classical (Nesterov 2005, 2007) and therefore its effectiveness should not be surprising.

The two regularizers are matrix analogues of the popular big-M constraints (constraints on the ℓ_∞ norm of the continuous variables) and ridge regularization (penalty on the ℓ_2^2 norm) for vectors. In mixed-integer optimization, such regularization terms can efficiently cope with non-linear constraints between continuous and binary variables (Bertsimas et al. 2019) and motivate our current approach. Practically speaking, regularization can be a natural component of the original problem (8), otherwise we advocate for introducing it artificially, for it leads to tractable algorithms with moderate impact on the resulting solution. For instance, if M is large enough so that the optimal solution to Problem (8), $\mathbf{X}^*$, satisfies $\|\mathbf{X}^*\|_\sigma \leq M$, Problems (9) and (8) are equivalent. In Section 3.5, we develop a disciplined technique for computing such an M . With the Frobenius norm penalty, the gap between Problem (9)'s and (8)'s objective is at most $\frac{1}{2\gamma} \|\mathbf{X}^*\|_F^2$, which can certainly be bounded whenever $\text{tr}(\mathbf{X})$ is bounded, as often occurs in practice.

For ease of notation, we let

$$g(\mathbf{X}) = \langle \mathbf{C}, \mathbf{X} \rangle + \begin{cases} 0, & \text{if } \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K}, \\ +\infty, & \text{otherwise,} \end{cases}$$

denote the unregularized second-stage cost for a given $\mathbf{X}$. Therefore, Problem (9) can be written as:

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k} f(\mathbf{Y}) + \lambda \cdot \text{tr}(\mathbf{Y}), \quad (10)$$

$$\text{where } f(\mathbf{Y}) := \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \Omega(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{X} = \mathbf{Y}\mathbf{X} \quad (11)$$

yields a best choice of $\mathbf{X}$ given $\mathbf{Y}$. As we establish in this section, this turns out to be a computationally useful reformulation, for f is convex in $\mathbf{Y}$ (see Theorem 3) and Lipschitz continuous (see Lemma 3), and therefore the non-convexity in the problem has been isolated within the set $\mathcal{Y}_n^k$.

Observe that both regularizers are coercive (i.e., “blow up” to $+\infty$ as $\|\mathbf{X}\| \rightarrow \infty$), and therefore render all unbounded solutions infeasible and ensure the compactness of the level sets of $\mathbf{X} \mapsto g(\mathbf{X}) + \Omega(\mathbf{X})$. This alleviates two of the major issues with conic duality (Ben-Tal and Nemirovski 2001, Theorem 2.4.1). First, regularization ensures that optimal solutions to conic problems are attained (see Blekherman et al. 2012, Example 2.27, for a regularization-free counterexample). Second, regularization ensures that infeasibility of a conic system is *certifiable*¹, i.e., there is either a feasible solution or a certificate of infeasibility. In general, such a procedure is not possible because a conic system could be infeasible but asymptotically feasible, i.e.,

$$\nexists \mathbf{X} : \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K} \text{ but } \exists \{\mathbf{X}_t\}_{t=1}^{\infty} : \mathbf{X}_t \in \mathcal{K} \quad \forall t \text{ with } \|\mathbf{A}\mathbf{X}_t - \mathbf{B}\| \rightarrow 0.$$

Here, the regularization term ensures that the set of feasible $\mathbf{X}$ (with objective at most $\theta_0 \in \mathbb{R}$) is a closed convex compact set. Therefore, $f(\mathbf{Y})$ cannot generate an asymptotically feasible problem.

Finally, the two regularization functions in Assumption 1 satisfy a non-trivial property which turns out to be crucial in both proving that $f(\mathbf{Y})$ is convex and deriving our overall algorithmic strategy:

LEMMA 2. *Consider a regularization function $\Omega(\mathbf{X})$ satisfying Assumption 1. There, there exists a Fenchel conjugate Ω^* (see, e.g., Boyd and Vandenberghe 2004, Chap. 3.3.1) such that, for any projection matrix $\mathbf{Y} \in \mathcal{Y}_n$ and any matrix $\boldsymbol{\alpha}$, we have*

$$\min_{\mathbf{X}} \{\Omega(\mathbf{Y}\mathbf{X}) + \langle \boldsymbol{\alpha}, \mathbf{Y}\mathbf{X} \rangle\} = \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22}),$$

and Ω^* is linear in $\mathbf{Y}$ (see Table 2 for its explicit definition).

Proof of Lemma 2 We start with the Frobenius regularization case, $\Omega(\mathbf{X}) = \frac{1}{2\gamma}\|\mathbf{X}\|_F$ and $\min_{\mathbf{X}} \{\Omega(\mathbf{Y}\mathbf{X}) + \langle \boldsymbol{\alpha}, \mathbf{Y}\mathbf{X} \rangle\} = \frac{1}{2\gamma}\|\mathbf{Y}\mathbf{X}\|_F + \langle \boldsymbol{\alpha}, \mathbf{Y}\mathbf{X} \rangle$. Any solution to the minimization problem satisfies the first-order condition $\frac{1}{\gamma}\mathbf{Y}\mathbf{X} + \mathbf{Y}\boldsymbol{\alpha} = 0$. Hence, since $\mathbf{Y}^2 = \mathbf{Y}$, $\mathbf{X}^* = -\gamma\mathbf{Y}\boldsymbol{\alpha}$ satisfies the first-order condition and the optimal objective value is $-\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22}) = -\frac{\gamma}{2}\langle \boldsymbol{\alpha}, \mathbf{Y}\boldsymbol{\alpha} \rangle$.

The spectral case is technically more challenging and detailed proofs are deferred to Appendix B. In the rectangular case, Lemma B.2 with $\mathbf{Y}' = \mathbb{I}_m$ yields

$$\min_{\mathbf{X}} \{\Omega(\mathbf{Y}\mathbf{X}) + \langle \boldsymbol{\alpha}, \mathbf{Y}\mathbf{X} \rangle\} = \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\frac{M}{2}\langle \mathbf{Y}, \mathbf{V}_{11} \rangle + \frac{M}{2}\langle \mathbf{I}_m, \mathbf{V}_{22} \rangle \text{ s.t. } \begin{pmatrix} \mathbf{V}_{11} & \boldsymbol{\alpha} \\ \boldsymbol{\alpha}^\top & \mathbf{V}_{22} \end{pmatrix} \succeq \mathbf{0}.$$

In the symmetric case, Lemma B.3 states that

$$\min_{\mathbf{X}} \{\Omega(\mathbf{Y}\mathbf{X}) + \langle \boldsymbol{\alpha}, \mathbf{Y}\mathbf{X} \rangle\} = \max_{\mathbf{V}_{11}, \mathbf{V}_{22} \succeq \mathbf{0}} -M\langle \mathbf{Y}, \mathbf{V}_{11} + \mathbf{V}_{22} \rangle \text{ s.t. } \boldsymbol{\alpha} = \mathbf{V}_{11} - \mathbf{V}_{22}. \quad \square$$

¹ Unless the conic dual is also infeasible, this case is unimportant for our purposes, because it only arises when the original problem is itself infeasible for any $\mathbf{Y}$, which can be checked a priori.

Table 2 Regularization penalties and conjugates, as defined in Lemma 2.

Penalty	$\Omega(\mathbf{X})$	$\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$	$\frac{\partial}{\partial Y_{i,j}} \Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$
Spectral norm ($\mathbf{X}$ rectangular)	$\begin{cases} 0, & \text{if } \ \mathbf{X}\ _\sigma \leq M, \\ +\infty, & \text{o.w.,} \end{cases}$	$\frac{M}{2} \langle \mathbf{Y}, \mathbf{V}_{11} \rangle + \frac{M}{2} \langle \mathbf{I}_m, \mathbf{V}_{22} \rangle$ s.t. $\begin{pmatrix} \mathbf{V}_{11} & \boldsymbol{\alpha} \\ \boldsymbol{\alpha}^\top & \mathbf{V}_{22} \end{pmatrix} \succeq \mathbf{0},$	$\frac{M}{2} V_{11,i,j}.$
Spectral norm ($\mathbf{X}$ symmetric)	$\begin{cases} 0, & \text{if } \ \mathbf{X}\ _\sigma \leq M, \\ +\infty, & \text{o.w.,} \end{cases}$	$M \langle \mathbf{Y}, \mathbf{V}_{11} + \mathbf{V}_{22} \rangle$ s.t. $\boldsymbol{\alpha} = \mathbf{V}_{11} - \mathbf{V}_{22},$ $\mathbf{V}_{11}, \mathbf{V}_{2,2} \succeq \mathbf{0},$	$M(V_{11} + V_{22})_{i,j}.$
Frobenius norm	$\frac{1}{2\gamma} \ \mathbf{X}\ _F^2$	$\frac{\gamma}{2} \langle \boldsymbol{\alpha}, \mathbf{Y} \boldsymbol{\alpha} \rangle$	$\frac{\gamma}{2} \langle \boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j \rangle$

3.2. A Saddle-Point Reformulation

We now reformulate Problem (9) as a saddle-point problem. This reformulation is significant for two reasons. First, as shown in the proof of Theorem 3, it leverages the nonlinear constraint $\mathbf{X} = \mathbf{Y}\mathbf{X}$ by introducing a new matrix of variables $\mathbf{V} \in \mathbb{R}^{n \times m}$ such that $\mathbf{V} = \mathbf{Y}\mathbf{X}$, giving:

$$f(\mathbf{Y}) = \min_{\mathbf{V}, \mathbf{X}} \{g(\mathbf{V}) + \Omega(\mathbf{Y}\mathbf{X}) : \mathbf{V} = \mathbf{Y}\mathbf{X}\},$$

a substitution reminiscent of the Douglas-Rachford splitting technique for composite convex optimization problems (Douglas and Rachford 1956, Eckstein and Bertsekas 1992)—the proof of Theorem 3 shows that this substitution does not change the optimal objective value. Second, it proves that the regularizer $\Omega(\mathbf{X})$ drives the convexity and smoothness of $f(\mathbf{Y})$. To derive the problem's dual, we require:

ASSUMPTION 2. *For each subproblem (11) generated by $f(\mathbf{Y})$ where $\mathbf{Y} \in \mathcal{Y}_n^k$, either the optimization problem is infeasible, or strong duality holds.*

Assumption 2 holds under Slater's constraint qualification (Boyd and Vandenberghe 2004, Section 5.2.3). By invoking Assumption 2, the following theorem reformulates (10) as a saddle-point problem:

THEOREM 3. *Suppose that Assumption 2 holds and $\Omega(\cdot)$ is either the spectral or Frobenius regularizer. Then, the following two optimization problems are equivalent:*

$$f(\mathbf{Y}) := \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \Omega(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{X} = \mathbf{Y}\mathbf{X}, \quad (12)$$

$$= \max_{\boldsymbol{\alpha}, \mathbf{V}_{11}, \mathbf{V}_{22}} h(\boldsymbol{\alpha}) - \Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22}), \quad (13)$$

where $h(\boldsymbol{\alpha}) := \max_{\boldsymbol{\Pi} : C - \boldsymbol{\alpha} - \mathbf{A}^\top \boldsymbol{\Pi} \in \mathcal{K}^*} \langle \mathbf{b}, \boldsymbol{\Pi} \rangle$, $\mathcal{K}^* := \{\mathbf{W} : \langle \mathbf{W}, \mathbf{X} \rangle \geq 0 \quad \forall \mathbf{X} \in \mathcal{K}\}$ denotes the dual cone to $\mathcal{K}$, and $\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$ is defined in Table 2.

Proof of Theorem 3 Let us fix $\mathbf{Y} \in \mathcal{Y}_n^k$, and suppose that strong duality holds for the inner minimization problem which defines $f(\mathbf{Y})$. To progress, we introduce a matrix $\mathbf{V} \in \mathbb{R}^{n \times m}$ such that $\mathbf{V} = \mathbf{Y}\mathbf{X}$ and obtain the relaxation:

$$\min_{\mathbf{X}, \mathbf{V}} g(\mathbf{V}) + \Omega(\mathbf{Y}\mathbf{X}) \quad \text{s.t.} \quad \mathbf{V} = \mathbf{Y}\mathbf{X}. \quad (14)$$

Let us verify that this relaxation is a valid substitution, i.e., that Problems (12) and (14) have the same optimal objective, $f(\mathbf{Y})$. If $\mathbf{X}$ is feasible for (12), then $(\mathbf{V} = \mathbf{X}, \mathbf{X})$ is obviously feasible for (14) with same objective value. Similarly, let $(\mathbf{V}, \mathbf{X})$ be feasible for (14). $\mathbf{YV} = \mathbf{Y}^2\mathbf{X} = \mathbf{YX} = \mathbf{V}$ since $\mathbf{Y}^2 = \mathbf{Y}$. Hence, $\mathbf{V}$ is feasible for (12) with same objective value.

Now, let $\boldsymbol{\alpha}$ denote the dual variables associated with the coupling constraints $\mathbf{V} = \mathbf{YX}$. The minimization problem is then equivalent to its dual problem, which is given by:

$$f(\mathbf{Y}) = \max_{\boldsymbol{\alpha}} h(\boldsymbol{\alpha}) + \min_{\mathbf{X}} [\Omega(\mathbf{YX}) + \langle \boldsymbol{\alpha}, \mathbf{YX} \rangle],$$

where $h(\boldsymbol{\alpha}) := \inf_{\mathbf{V}} g(\mathbf{V}) - \langle \mathbf{V}, \boldsymbol{\alpha} \rangle$ is, up to a sign, the Fenchel conjugate of g . By a standard application of Fenchel duality, it follows that

$$h(\boldsymbol{\alpha}) = \max_{\boldsymbol{\Pi}} \langle \mathbf{b}, \boldsymbol{\Pi} \rangle + \begin{cases} 0, & \text{if } \mathbf{C} - \boldsymbol{\alpha} - \mathbf{A}^\top \boldsymbol{\Pi} \in \mathcal{K}^*, \\ +\infty, & \text{otherwise.} \end{cases}$$

Finally, from Lemma 2 we have $\min_{\mathbf{X}} \{\Omega(\mathbf{YX}) + \langle \boldsymbol{\alpha}, \mathbf{YX} \rangle\} = \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$, which concludes the proof.

Alternatively, under either penalty, if the inner minimization problem defining $f(\mathbf{Y})$ is infeasible, then its dual problem is unbounded by weak duality.² $\square$

REMARK 4. In the unregularized case, i.e., $\Omega(\mathbf{X}) = 0$, we can derive a similar reformulation:

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k} \max_{\boldsymbol{\alpha} \in \mathbb{R}^{n \times m}} h(\boldsymbol{\alpha}) + \lambda \cdot \text{tr}(\mathbf{Y}) \text{ s.t. } \mathbf{Y}\boldsymbol{\alpha} = \mathbf{0}. \quad (15)$$

Under this lens, regularization of the primal problem is equivalent to a relaxation in the dual formulation: the hard constraint $\mathbf{Y}\boldsymbol{\alpha} = \mathbf{0}$ is penalized by $-\Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$.

REMARK 5. By Theorem 3 and Lemma 2, $f(\mathbf{Y})$ is convex as the point-wise maximum of functions which are linear in $\mathbf{Y}$.

By Theorem 3, when we evaluate $f(\hat{\mathbf{Y}})$, one of two alternatives occur. The first is that we have $f(\hat{\mathbf{Y}}) < +\infty$ and there is some optimal $(\boldsymbol{\alpha}, \mathbf{V}_{11}, \mathbf{V}_{22})$. In this case, we construct the lower approximation

$$f(\mathbf{Y}) \geq f(\hat{\mathbf{Y}}) + \langle \mathbf{H}, \mathbf{Y} - \hat{\mathbf{Y}} \rangle,$$

where $H_{i,j} = \frac{\partial}{\partial Y_{i,j}} \Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22})$ (see Table 2 for closed-form expression of the partial derivatives, which follow readily from Danskin's theorem (see, e.g., Bertsekas 2016, Prop. B.22)). The second alternative is that $f(\hat{\mathbf{Y}}) = +\infty$, in which case, by the conic duality theorem (see Ben-Tal and Nemirovski 2001, Chapter 2) there exists a $(\boldsymbol{\alpha}, \boldsymbol{\Pi})$ such that

$$\mathbf{C} - \boldsymbol{\alpha} - \mathbf{A}^\top \boldsymbol{\Pi} \in \mathcal{K}^*, \text{ and } \langle \mathbf{b}, \boldsymbol{\Pi} \rangle > \langle -\mathbf{H}, \hat{\mathbf{Y}} \rangle. \quad (16)$$

² Weak duality implies that the dual problem is either unfeasible or unbounded. Since the feasible set of the maximization problem does not depend on $\mathbf{Y}$, it is always feasible, unless the original problem is itself infeasible. Therefore, we assume without loss of generality that it is unbounded.

Under this alternative, we can separate $\hat{\mathbf{Y}}$ from the set of feasible $\mathbf{Y}$'s by imposing the cut $0 \geq \langle \mathbf{b}, \mathbf{\Pi} \rangle + \langle \mathbf{H}, \mathbf{Y} \rangle$. Under either alternative, we obtain a globally valid first-order underestimator of the form

$$zf(\mathbf{Y}) \geq h + \langle \mathbf{H}, \mathbf{Y} - \hat{\mathbf{Y}} \rangle, \quad (17)$$

where z, h are defined as

$$z = \begin{cases} 1, & \text{if } f(\hat{\mathbf{Y}}) < +\infty, \\ 0, & \text{if } f(\hat{\mathbf{Y}}) = +\infty, \end{cases} \quad \text{and} \quad h = \begin{cases} f(\hat{\mathbf{Y}}), & \text{if } f(\hat{\mathbf{Y}}) < +\infty, \\ \langle \mathbf{b}, \mathbf{\Pi} \rangle + \langle \mathbf{H}, \hat{\mathbf{Y}} \rangle, & \text{if } f(\hat{\mathbf{Y}}) = +\infty. \end{cases} \quad (18)$$

This observation suggests that a valid numerical strategy for minimizing $f(\mathbf{Y})$ is to iteratively minimize and refine a piecewise linear underestimator of $f(\mathbf{Y})$ defined by the pointwise supremum of a finite number of underestimators of the form $zf(\mathbf{Y}) \geq h + \langle \mathbf{H}, \mathbf{Y} - \mathbf{Y} \rangle$. Indeed, as we will see in Section 4, this strategy gives rise to the global optimization algorithm known as outer-approximation.

Smoothness We now demonstrate that $f(\mathbf{Y})$ is smooth, in the sense of Lipschitz continuity, under a boundedness assumption on the size of the dual variables, which is a crucial property for ensuring the convergence of our global optimization methods and bounding the quality of our semidefinite relaxation and greedy rounding methods. Formally, the following result follows directly from Theorem 3.

LEMMA 3. *Let $\mathbf{Y}, \mathbf{Y}' \in \text{Conv}(\mathcal{Y}_n^k)$ be on the convex hull of the orthogonal projection matrices. Then*

$$f(\mathbf{Y}) - f(\mathbf{Y}') \leq \Omega^*(\boldsymbol{\alpha}^*(\mathbf{Y}), \mathbf{Y}' - \mathbf{Y}, \mathbf{V}_{11}^*(\mathbf{Y}), \mathbf{V}_{22}^*(\mathbf{Y})).$$

Moreover, suppose $\boldsymbol{\alpha}^(\mathbf{Y}), \mathbf{V}_{11}^*(\mathbf{Y}), \mathbf{V}_{22}^*(\mathbf{Y})$ can be bounded independently from $\mathbf{Y}$, i.e., $\|\boldsymbol{\alpha}^*(\mathbf{Y})\|_\sigma \leq L_1$, $\|\mathbf{V}_{11}^*(\mathbf{Y})\|_\sigma \leq L_2$, $\|\mathbf{V}_{22}^*(\mathbf{Y})\|_\sigma \leq L_2$. Then, under spectral regularization we have*

$$f(\mathbf{Y}) - f(\mathbf{Y}') \leq M \langle \mathbf{V}_{11}^*(\mathbf{Y}), \mathbf{Y}' - \mathbf{Y} \rangle \leq ML_2 \|\mathbf{Y}' - \mathbf{Y}\|_*, \quad (19)$$

and under Frobenius regularization we have

$$f(\mathbf{Y}) - f(\mathbf{Y}') \leq \frac{\gamma}{2} \langle \boldsymbol{\alpha}^{*\top}(\mathbf{Y}) \boldsymbol{\alpha}^*(\mathbf{Y}), \mathbf{Y}' - \mathbf{Y} \rangle \leq \frac{\gamma}{2} L_1^2 \|\mathbf{Y}' - \mathbf{Y}\|_*, \quad (20)$$

where the bounds involving L_1, L_2 follow from Holder's inequality³.

REMARK 6. Section 3.5 develops disciplined techniques for computing an M such that the constraint $\|\mathbf{X}\|_\sigma \leq M$ in the primal does not alter the optimal objective. The same technique, applied to the dual, yields explicit bounds on L_1 . Moreover, since there exists an optimal pair $(\mathbf{V}_{11}, \mathbf{V}_{22})$ which is an explicit functions of an optimal $\boldsymbol{\alpha}$, this translates into explicit bounds on L_2 .

3.3. Semidefinite Relaxations

To lower bound (10)'s objective, we invoke Lemma 1 to relax the non-convex constraint $\mathbf{Y} \in \mathcal{Y}_n^k$ to

$$\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k) = \{\mathbf{Y} \in S^n : \mathbf{0} \preceq \mathbf{Y} \preceq \mathbb{I}, \text{tr}(\mathbf{Y}) \leq k\}.$$

³ Namely, $|\langle \mathbf{X}, \mathbf{Y} \rangle| \leq \|\mathbf{X}\|_\sigma \|\mathbf{Y}\|_*$, since the $\|\cdot\|_\sigma$ and $\|\cdot\|_*$, as the matrix analogs of the ℓ_∞ and ℓ_1 norms, are dual.

This yields the saddle-point problem

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \max_{\boldsymbol{\alpha}, \mathbf{V}_{11}, \mathbf{V}_{22} \in S^m} h(\boldsymbol{\alpha}) - \Omega^*(\boldsymbol{\alpha}, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22}) + \lambda \cdot \text{tr}(\mathbf{Y}). \quad (21)$$

Problem (21) can in turn be reformulated as an SDO. Indeed, under Assumption 2, we obtain a semidefinite formulation by taking Problem (21)'s dual with respect to $\boldsymbol{\alpha}$. Formally, we have the following results (proofs deferred to Appendix A.4 and A.5 respectively):

LEMMA 4. *Suppose that Assumption 2 holds. Then, strong duality holds between:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \max_{\boldsymbol{\alpha} \in \mathbb{R}^{n \times m}} h(\boldsymbol{\alpha}) - \frac{\gamma}{2} \langle \boldsymbol{\alpha}, \mathbf{Y} \boldsymbol{\alpha} \rangle + \lambda \cdot \text{tr}(\mathbf{Y}), \quad (22)$$

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}, \boldsymbol{\theta} \in S^n} g(\mathbf{X}) + \frac{1}{2\gamma} \text{tr}(\boldsymbol{\theta}) + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{\theta} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{Y} \end{pmatrix} \succeq \mathbf{0}. \quad (23)$$

LEMMA 5. *Suppose that Assumption 2 holds. Then, strong duality holds between:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \max_{\boldsymbol{\alpha} \in S^n, \mathbf{V}_{11}, \mathbf{V}_{22} \succeq \mathbf{0}} h(\boldsymbol{\alpha}) - M \langle \mathbf{Y}, \mathbf{V}_{11} + \mathbf{V}_{22} \rangle + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad \boldsymbol{\alpha} = \mathbf{V}_{11} - \mathbf{V}_{22}, \quad (24)$$

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \min_{\mathbf{X} \in S^n} g(\mathbf{X}) + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad -M\mathbf{Y} \preceq \mathbf{X} \preceq M\mathbf{Y}. \quad (25)$$

We now offer some remarks on these bi-dual problems:

- We can derive a more general version of Lemma 5 without the symmetry assumption on $\mathbf{X}$ in much the same manner, via the Schur complement lemma.
- Problem (23)'s formulation generalizes the perspective relaxation from vectors to matrices. This suggests that (23) is an efficient formulation for addressing rank constraints, as perspective formulations efficiently address cardinality constrained problems with conic quadratic (Günlük and Linderoth 2012) or power cone (Aktürk et al. 2009) objectives, indeed, they provide a theoretical basis for scalable algorithms for sparse regression (Bertsimas and Van Parys 2020, Hazimeh et al. 2020), sparse portfolio selection (Zheng et al. 2014, Bertsimas and Cory-Wright 2018) and network design (Fischetti et al. 2017) problems among others..

3.4. Convex Penalty Interpretations of Relaxations

In this section, we consider instances where rank is penalized in the objective only and interpret the above convex relaxations as penalty functions, in the tradition of Fazel (2002), Recht et al. (2010). In the presence of the Frobenius penalty, our first result generalizes the *reverse Huber penalty* of Pilanci et al. (2015), Dong et al. (2015) from cardinality to rank objectives (proof deferred to Appendix A.7).

LEMMA 6. *Suppose that Assumption 2 holds. Then, the following problems are equivalent:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}, \boldsymbol{\theta} \in S^n} g(\mathbf{X}) + \frac{1}{2\gamma} \text{tr}(\boldsymbol{\theta}) + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{\theta} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{Y} \end{pmatrix} \succeq \mathbf{0}, \quad (26)$$

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \sum_{i=1}^n \min \left(\sqrt{\frac{2\lambda}{\gamma}} \sigma_i(\mathbf{X}), \lambda + \frac{\sigma_i(\mathbf{X})^2}{2\gamma} \right). \quad (27)$$

REMARK 7.

$$\text{Since } \min_{0 \leq \theta \leq 1} \left[\lambda \theta + \frac{t^2}{\theta} \right] = \begin{cases} 2\sqrt{\lambda}|t|, & \text{if } |t| \leq \sqrt{\lambda}, \\ t^2 + \lambda, & \text{otherwise,} \end{cases}$$

the proof of Lemma 6 reveals that Problems (26)-(27) are equivalent to minimizing

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}, \boldsymbol{\theta} \in \mathbb{R}^n: \mathbf{0} \leq \boldsymbol{\theta} \leq \mathbf{e}} g(\mathbf{X}) + \sum_{i=1}^n \left(\lambda \theta_i + \frac{\sigma_i(\mathbf{X})^2}{2\gamma \theta_i} \right), \quad (28)$$

which applies the smooth penalty $t \rightarrow \lambda \theta + \frac{t^2}{2\gamma \theta} : 0 \leq \theta \leq 1$ to model the non-convex cost $t \rightarrow \lambda \|t\|_0 + \frac{t^2}{2\gamma}$ incurred by each singular value of $\mathbf{X}$. Indeed, this smooth penalty is precisely the convex envelope of the non-convex cost function (see, e.g., Günlük and Linderoth 2012). Compared to other penalties for low-rank problems (Fan and Li 2001, Zhang 2010), this generalized Huber penalty is convex, amenable to efficient alternating minimization procedures (see Section 4.2.2) and could be of independent interest to the statistical learning community.

Lemma 6 proposes an alternative to the nuclear norm penalty for approximately solving low-rank problems. This is significant, as many low-rank problems have constraints $\mathbf{X} \succeq \mathbf{0}$, $\text{tr}(\mathbf{X}) = k$ (e.g. sparse PCA (d’Aspremont et al. 2007), k -means clustering (Peng and Wei 2007)), and under these constraints a nuclear norm cannot encourage low-rank solutions (Zhang et al. 2013), while Lemma 6’s penalty can.

Our next results relate rank minimization problems with a spectral regularizer to the nuclear norm penalty, in both the square symmetric and the rectangular case (proofs deferred to Appendix A.6):

LEMMA 7. *Suppose that Assumption 2 holds. Then, the following problems are equivalent:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)} \min_{\mathbf{X} \in \mathbb{S}^n} g(\mathbf{X}) + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad -M\mathbf{Y} \preceq \mathbf{X} \preceq M\mathbf{Y}, \quad (29)$$

$$\min_{\mathbf{X} \in \mathbb{S}^n} g(\mathbf{X}) + \frac{\lambda}{M} \|\mathbf{X}\|_* \quad \text{s.t.} \quad \|\mathbf{X}\|_\sigma \leq M. \quad (30)$$

LEMMA 8. *Suppose that Assumption 2 holds. Then, the following problems are equivalent:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n), \mathbf{Y}' \in \text{Conv}(\mathcal{Y}_m)} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \frac{\lambda}{2} \text{tr}(\mathbf{Y}) + \frac{\lambda}{2} \text{tr}(\mathbf{Y}') \quad \text{s.t.} \quad \begin{pmatrix} M\mathbf{Y} & \mathbf{X} \\ \mathbf{X}^\top & M\mathbf{Y}' \end{pmatrix} \succeq \mathbf{0}, \quad (31)$$

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \frac{\lambda}{M} \|\mathbf{X}\|_* \quad \text{s.t.} \quad \|\mathbf{X}\|_\sigma \leq M. \quad (32)$$

3.5. Bounding the Spectral Penalty

We now present techniques for computing an M such that the optimal values of (8) and (9) agree; these are essentially a generalization of similar techniques for logically constrained MICOs (Bertsimas et al. 2016, Section 2.3). We first consider the positive semidefinite case, then develop the general case. Note that the M s obtained here are not, in general, the smallest possible—computing this quantity is NP-hard even for MICOs (Kleinert et al. 2020).

Positive semidefinite case. Let $\mathbf{X} \in S_+^n$. Then, since $\text{tr}(\mathbf{X}) = \sum_{i=1}^n \lambda_i(\mathbf{X}) \geq \lambda_1(\mathbf{X})$, the optimal value of the following problem gives a valid bound on M :

$$M_{tr} := \max_{\mathbf{X} \in S_+^n} \text{tr}(\mathbf{X}) \text{ s.t. } \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K}. \quad (33)$$

Alternatively, since the volume of $\mathbf{X}$ and its spectral radius are related, we can maximize

$$\mathbf{X}^* = \arg \max_{\mathbf{X} \in S_+^n, \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K}} \log \det(\mathbf{X}). \quad (34)$$

We have $\det(\mathbf{X}) \leq (\lambda_{\max}(\mathbf{X}))^n$, which implies we can set $M_{det} = \sqrt[n]{\det(\mathbf{X}^*)}$. Alternatively, we could solve both (33) and (34) and set $M = \min(M_{tr}, M_{det})$, which also gives a valid bound.

General case Let $\mathbf{X} \in \mathbb{R}^{n \times m}$. Then, computing a valid M is more expensive, because $\|\mathbf{X}\|_\sigma$, $\|\mathbf{X}\|_F$, and $\|\mathbf{X}\|_*$ are not mixed-integer convex representable when maximizing (Lubin et al. 2021, Corollary 4.1). To progress, let us evaluate $n \cdot m$ values $M_{i,j}$, each computed by solving two conic problems:

$$M_{i,j} := \max_{\mathbf{X} \in \mathbb{R}^{n \times m}, \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K}} |X_{i,j}| \quad (35)$$

Then, a valid M is given by $\sum_{i,j} M_{i,j}$, since $\sum_{i,j} M_{i,j} \geq \|\mathbf{X}\|_1 \geq \|\mathbf{X}\|_\sigma$ (Boyd and Vandenberghe 2004).

Unbounded interpretation. We remind the reader that interpreting the case where $M = +\infty$ requires caution. Indeed, when a feasible $\mathbf{X}$ and extreme ray $\mathbf{W}$ give an unbounded direction such that

$$\text{rank}(\mathbf{X} + \lambda \mathbf{W}) \leq k \quad \forall \lambda \geq 0$$

we have a certificate that no valid bound on M exists. Alternatively, when $\text{rank}(\mathbf{X} + \lambda \mathbf{W}) \geq k + 1$, we don't actually know whether a valid M exists, since the set $\{\mathbf{X} \in S_+^n : \mathbf{A}\mathbf{X} = \mathbf{B}, \mathbf{X} \in \mathcal{K}\}$ could be unbounded, even while its intersection with a low-rank set is bounded. This difficulty also arises in the sparsity-constrained (MICO) case however, and therefore should not be unexpected; it can be dealt with by cross-validating M/γ , which is usually acceptable since $\mathbf{A}, \mathbf{B}$ are usually estimated from data.

4. Efficient Algorithmic Approaches

In this section, we present an efficient numerical approach to solve Problem (1) and its convex relaxations. The backbone is an outer-approximation strategy, embedded within a non-convex QCQO branch-and-bound procedure to solve the problem exactly. We also propose rounding heuristics to find good feasible solutions, and semidefinite free methods for optimizing over (1)'s convex relaxations.

The primary motivations for developing an outer-approximation procedure and solving mixed-projection problem as saddle-point problems are twofold. First, we are not aware of any exact solvers which address mixed-projection problems with semidefinite constraints. Instead, a decomposition strategy like outer-approximation can be readily implemented using a conjunction of **Gurobi** (to solve non-convex quadratically constrained master problems) and **Mosek** (to solve conic subproblems). Second, decomposition schemes for mixed-integer semidefinite problems typically outperform one-shot strategies (Belotti et al. 2013), so we expect - and observe in Section 5.3 - a similar comparison for mixed-projection optimization, hence connecting the frameworks in both theory (see Table 1) and practice.

4.1. A Globally Optimal Cutting-Plane Method

The analysis in the previous section reveals that evaluating $f(\mathbf{Y})$ yields a globally valid first-order underestimator of $f(\cdot)$. Therefore, a numerically efficient strategy for minimizing $f(\mathbf{Y})$ is to iteratively minimize and refine a piecewise linear underestimator of $f(\mathbf{Y})$. This strategy is known as outer-approximation (OA), and was originally proposed by Duran and Grossmann (1986). OA iteratively constructs underestimators of the following form at each iterate $t + 1$:

$$f_{t+1}(\mathbf{Y}) = \max_{1 \leq i \leq t} \{f(\mathbf{Y}_i) + \langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}_i \rangle\}. \quad (36)$$

By iteratively minimizing $f_{t+1}(\mathbf{Y})$ and imposing the resulting cuts when constructing the next underestimator, we obtain a non-decreasing sequence of underestimators $f_t(\mathbf{Y}_t)$ and non-increasing sequence of overestimators $\min_{i \in [t]} f(\mathbf{Y}_i)$ which converge to an ϵ -optimal solution within a finite number of iterations; see also Section 3.2 for details on cut generation. Indeed, since $\text{Conv}(\mathcal{Y}_n^k)$ is a compact set and $f(\cdot)$ is an L -Lipschitz continuous function in $\mathbf{Y}$, OA never visits a ball of radius $\frac{\epsilon}{L}$ twice.

We now formalize this numerical procedure in Algorithm 1, and state its convergence properties (proof of convergence deferred to Appendix A.8):

Algorithm 1 An outer-approximation method for Problem (10)

Require: Initial solution $\mathbf{Y}_1$

$t \leftarrow 1$

repeat

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

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k, \theta} \theta + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t. } z_i \theta \geq h_i + \langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}_i \rangle \quad \forall i \in [t].$$

 Compute $f(\mathbf{Y}_{t+1}), \mathbf{H}_{t+1}, z_{t+1}, d_{t+1}$

until $f(\mathbf{Y}_t) - \theta_t \leq \epsilon$

return $\mathbf{Y}_t$

THEOREM 4. *Suppose that Assumptions 1-2 hold, and that there exists some Lipschitz constant L such that for any feasible $\mathbf{Y}, \mathbf{Y}' \in \text{Conv}(\mathcal{Y}_n^k)$ we have: $|f(\mathbf{Y}) - f(\mathbf{Y}')| \leq L\|\mathbf{Y} - \mathbf{Y}'\|_F$, and for any feasibility cut $\langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}_i \rangle + h_i \leq 0$ we have $|\langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}' \rangle| \leq L\|\mathbf{Y} - \mathbf{Y}'\|_F$. Let $\mathbf{Y}_t \in \mathcal{Y}_n^k$ be a feasible solution returned by the t^{th} iterate of Algorithm 1, where*

$$t \geq \left(\frac{Lk}{\epsilon} + 1 \right)^{n^2}.$$

Then, $\mathbf{Y}_t$ is an ϵ -optimal and ϵ -feasible solution to Problem (9). Moreover, suppose that we set $\epsilon \rightarrow 0$. Then, any limit point of $\{\mathbf{Y}_t\}_{t=1}^\infty$ solves (9).

4.1.1. Optimizing Over Orthogonal Projection Matrices To successfully implement Algorithm 1, we need to repeatedly solve optimization problems of the form

$$\min_{\mathbf{Y} \in \mathcal{Y}_n^k, \theta} \theta + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t. } z_i \theta \geq h_i + \langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}_i \rangle \quad \forall i \in [t], \quad (37)$$

which requires a tractable representation of $\mathcal{Y}_n^k$. Fortunately, Gurobi 9.0 contains a globally optimal spatial branch-and-bound method for general QCQOs which recursively partitions the feasible region into boxes and invokes the ubiquitous McCormick inequalities to obtain valid upper and lower bounds on each box—see Achterberg and Towle (2020) for a discussion of Gurobi’s bilinear solver, Belotti et al. (2013) for a general theory of spatial branch-and-bound. Therefore, we represent $\mathbf{Y}$ by introducing a matrix $\mathbf{U} \in \mathbb{R}^{n \times k}$ and requiring that $\mathbf{Y} = \mathbf{U}\mathbf{U}^\top$ and $\mathbf{U}^\top \mathbf{U} = \mathbb{I}$. This allows Algorithm 1 to be implemented by iteratively solving a sequence of QCQOs and conic optimization problems. Moreover, to decrease the amount of branching required in each iteration of Algorithm 1, we impose an outer-approximation of the valid constraint $\mathbf{Y} \succeq \mathbf{U}\mathbf{U}^\top$. Specifically, we strengthen the formulation by imposing second-order cone relaxations of the PSD constraint. First, we require that the 2×2 minors in $\mathbf{Y}$ are non-negative, i.e., $Y_{i,j}^2 \leq Y_{i,i}Y_{j,j} \quad \forall i, j \in [n]$, as proposed in Ahmadi and Majumdar (2019), Bertsimas and Cory-Wright (2020). Second, we require that the on-diagonal entries of $\mathbf{Y} \succeq \mathbf{U}\mathbf{U}^\top$ are non-negative i.e., $Y_{i,i} \geq \sum_{t=1}^k U_{i,t}^2 \quad \forall i \in [n]$. Finally, we follow Atamtürk and Gomez (2019, Proposition 5) in taking a second-order cone approximation of the 2×2 minors in $\mathbf{Y} \succeq \mathbf{U}\mathbf{U}^\top$ i.e., $0 \geq \|\mathbf{U}_i \pm \mathbf{U}_j\|_2^2 \pm 2Y_{i,j} - Y_{i,i} - Y_{j,j} \quad \forall i, j \in [n]$. All told, we have⁴:

$$\begin{aligned} \min_{\mathbf{Y} \in S^n, \mathbf{U} \in \mathbb{R}^{n \times k}, \theta} \quad & \theta + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t. } z_i \theta \geq h_i + \langle \mathbf{H}_i, \mathbf{Y} - \mathbf{Y}_i \rangle \quad \forall i \in [t], \\ & \mathbf{Y} = \mathbf{U}\mathbf{U}^\top, \mathbf{U}^\top \mathbf{U} = \mathbb{I}, Y_{i,i}Y_{j,j} \geq Y_{i,j}^2 \quad \forall i, j \in [n], Y_{i,i} \geq \sum_{t=1}^k U_{i,t}^2 \quad \forall i \in [n], \text{tr}(\mathbf{Y}) = k, \\ & 0 \geq \|\mathbf{U}_i + \mathbf{U}_j\|_2^2 - 2Y_{i,j} - Y_{i,i} - Y_{j,j}, \quad 0 \geq \|\mathbf{U}_i - \mathbf{U}_j\|_2^2 + 2Y_{i,j} - Y_{i,i} - Y_{j,j} \quad \forall i, j \in [n]. \end{aligned} \quad (38)$$

Finally, for a given $\mathbf{Y}, \mathbf{U}$, we strengthen this formulation by imposing second-order cone cuts of the form $\langle \mathbf{Y} - \mathbf{U}\mathbf{U}^\top, \mathbf{u}\mathbf{u}^\top \rangle \geq 0$, where $\mathbf{u}$ is the most negative eigenvector of $\mathbf{Y} - \mathbf{U}\mathbf{U}^\top$, as proposed by Sherali and Fraticelli (2002).

As described, a linear optimization problem over the set of orthogonal projection matrices is solved at each iteration, hence building a new branch-and-bound tree each time. We refer to this implementation as a “multi-tree” method. Although inefficient if implemented naively, multi-tree methods benefit from gradually tightening the numerical tolerance of the solver as the number of cuts increases.

To improve the efficiency of Algorithm 1, one can integrate the entire procedure within a single branch-and-cut tree using lazy callbacks, as originally proposed in the context of MICO by Quesada and Grossmann (1992). Henceforth, we refer to this implementation as a “single-tree” method. However, the

⁴ It should be noted that this formulation is rather complicated because non-convex QCQO solvers such as Gurobi currently do not model PSD constraints. If they did, we would supplant the second-order cone constraints with $\mathbf{Y} \succeq \mathbf{U}\mathbf{U}^\top$ and thereby obtain a simpler master problem.

benefit from using multi-tree over single-tree is not straightforward for it depends on how the method is engineered. We benchmark both implementations in Section 5.3.

4.1.2. A Simple Benchmark We now lay out a simple approach for solving low-rank problems exactly, which we will compare against in our numerical experiments. Rather than introducing an orthogonal projection matrix $\mathbf{Y}$, we let $\mathbf{X} = \mathbf{U}\mathbf{V}^\top$ where $\mathbf{U} \in \mathbb{R}^{n \times k}$ and $\mathbf{V} \in \mathbb{R}^{m \times k}$, and $\mathbf{U}, \mathbf{V}$ are both bounded in absolute value by big-M constraints of the form $|U_{i,j}| \leq 1, |V_{i,j}| \leq M$. Assuming that the objective and constraints are QCQO representable, as occurs for all of the examples mentioned in the introduction, this formulation can then be optimized over using Gurobi’s piecewise linear reformulation technique for general QCQOs. Formally, a rank constraint $\text{Rank}(\mathbf{X}) \leq k$ leads to:

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}, \mathbf{U} \in \mathbb{R}^{n \times k}, \mathbf{V} \in \mathbb{R}^{m \times k}} \quad & \Omega(\mathbf{X}) + g(\mathbf{X}) \\ \text{s.t.} \quad & \mathbf{X} \geq \mathbf{U}\mathbf{V}^\top - \epsilon \mathbf{E}, \mathbf{X} \leq \mathbf{U}\mathbf{V}^\top + \epsilon \mathbf{E}, \|\mathbf{U}_i\|_2 \leq 1 \quad \forall i \in [n], \|\mathbf{V}\|_\infty \leq M, \end{aligned}$$

where $\mathbf{E}$ is a matrix of all ones. Note however that, as we observe in Section 5, this approach is significantly less efficient than the previously described cutting-plane approaches.

4.2. Lower bounds via Semidefinite Relaxations

To certify optimality, high-quality lower bounds are of interest and can be obtained by relaxing the non-convex constraint $\mathbf{Y} \in \mathcal{Y}_n^k$ to $\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)$ to obtain a semidefinite relaxation as discussed in Lemma 1. In addition to a valid lower bound on (10)’s objective, the optimal solution to the relaxation $\mathbf{Y}^*$ is a natural candidate for a random rounding strategy, for stronger convex relaxations lead to superior random rounding strategies. We will explore such rounding strategies in detail in the next section.

The convex relaxation yields the optimization problem (21) which can be solved using a cutting-plane method (see Section 4.2.1), an alternating minimization method (see Section 4.2.2) or reformulated as an SDO and solved as such. Since Algorithm 1 is also an outer-approximation scheme, solving the convex relaxation via a cutting-plane method has the additional benefit of producing valid linear lower-approximations of $f(\mathbf{Y})$ to initialize Algorithm 1 with.

4.2.1. Cutting-Plane Methods for Improving the Root Node Bound As mentioned previously, Problem (21) can be solved by a cutting-plane method such as Kelley’s algorithm (see Kelley 1960), which is a continuous analog of Algorithm 1 that solves Problem (10) over $\text{Conv}(\mathcal{Y}_n^k)$, rather than $\mathcal{Y}_n^k$. The main benefit of such a cutting-plane method is that the cuts generated are valid for both $\text{Conv}(\mathcal{Y}_n^k)$ and $\mathcal{Y}_n^k$, and therefore can be used to initialize Algorithm 1 and ensure that its *initial* lower bound is equal to the semidefinite relaxation. As demonstrated by Fischetti et al. (2017) in the context of MICO and facility location problems, this approach often accelerates the convergence of decomposition schemes by orders of magnitude. We present pseudocode in Appendix D.1

Figure 1 Convergence behavior of Kelley’s method and the in-out method for solving the semidefinite relaxation of a synthetic matrix completion instance where $n = 100$ (left), and lower bounds generated by a single-tree implementation of Algorithm 1 for a synthetic matrix completion instance where $n = 10$ (right).

Figure 1’s left panel illustrates the convergence of Kelley’s method and the **in-out** method for solving the semidefinite relaxation of a noiseless matrix completion problem⁵. Note that in our plot of the **in-out** method on the continuous relaxation we omit the time required to first solve the SDO relaxation; this is negligible (38.4s) compared to the time required for either approach to solve the relaxation using cutting planes. Observe that the **in-out** method’s lower bound is both initially better and converges substantially faster to the optimal solution than Kelley’s method. This justifies our use of the **in-out** method over Kelley’s method for a stabilizing cut loop in numerical experiments.

Once the relaxation is solved, the generated cuts are used to initialize Algorithm 1. Figure 1’s right panel displays the convergence profile of the lower bound of Algorithm 1 initialized with cuts from Kelley’s or the **in-out** method (with a limit of 100 cuts). We use a single-tree implementation of Algorithm 1⁶ and again a noiseless matrix completion setting⁷. We also consider the impact of using the SOC inequalities $Y_{i,j}^2 \leq Y_{i,i}Y_{j,j}$ in the master problem formulation. Using the **in-out** method and imposing the SOC inequalities are both vitally important for obtaining high-quality lower bounds from Algorithm 1. Accordingly, we make use of both ingredients in our numerical experiments.

4.2.2. Solving the Semidefinite Relaxation at Scale via Alternating Minimization In preliminary numerical experiments, we found that modern IPM codes such as **Mosek** 9.0 cannot optimize over the Frobenius/nuclear norm penalties when $n > 200$ on a standard laptop. As real-world low-rank problems are often large-scale, we now explore more scalable alternatives for optimizing over these penalties. As scalable alternatives for the nuclear norm penalty have been studied, we focus on the

⁵ The data generation process is detailed in Section 5.2. Here, $n = 100$, $p = 0.25$, $r = 1$, and $\gamma = \frac{20}{p}$.

⁶ We warm-start the upper bound with greedy rounding and the Burer-Monterio local improvement heuristic described in Section 4.3. To mitigate against numerical instability, we opted to be conservative with our parameters, and therefore turned Gurobi’s heuristics off, set FuncPieceError and FuncPieceLength to their minimum possible values (10^{-5} and 10^{-6}), set the MIP gap to 1% and the time limit for each solve to one hour.

⁷ Here, $n = 10$, $p = 0.25$, $r = 1$, and $\gamma = \frac{5}{p}$.

Frobenius penalty, and refer to (Recht et al. 2010) for nuclear norm minimization. We begin our analysis with the following result (proof deferred to Appendix A.9):

LEMMA 9. *For any fixed $\mathbf{X}_t$ in Problem (23), an optimal choice of $\boldsymbol{\theta}$ is given by $\boldsymbol{\theta}^* = \mathbf{X}_t^\top (\mathbf{Y}^*)^\dagger \mathbf{X}_t$, where $\mathbf{Y}^* = \sum_{i=1}^n \rho_i^* \mathbf{u}_i \mathbf{u}_i^\top$, $\mathbf{X}_t = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top$ is an SVD of $\mathbf{X}_t$, and $\boldsymbol{\rho}^*$ is an optimal solution to the following second order cone problem:*

$$\min_{\boldsymbol{\rho} \in [0,1]^n: \mathbf{e}^\top \boldsymbol{\rho} \leq k} \lambda \cdot \mathbf{e}^\top \boldsymbol{\rho} + \sum_{i=1}^n \frac{\sigma(\mathbf{X}_t)^2}{2\gamma \rho_i}. \quad (39)$$

As optimizing over $\mathbf{X}$ for a fixed $\mathbf{Y}_t$ is straightforward, Lemma 9 suggests a viable approach for optimize over the Frobenius norm penalty is alternating minimization (AM; see Beck and Teboulle 2009, for a modern implementation). By specializing Beck and Teboulle (2009)’s implementation of AM to the Frobenius norm penalty, we obtain an efficient numerical strategy for obtaining an optimal solution to (23), which we present in Algorithm 2; we note that since $\langle \mathbf{X} \mathbf{X}, \mathbf{Y}^\dagger \rangle$ is jointly convex in $\mathbf{X}, \mathbf{Y}$ (this follows directly from Lemma 4), alternating minimization converges to an optimal solution to the semidefinite relaxation under standard convergence conditions for block coordinate descent techniques for convex programs (see, e.g., Bertsekas 2016, Section 3.7) such as the introduction of a proximal term.

We now discuss some enhancements to Algorithm 2 which improve its rate of convergence in practice.

- Imposing a proximal regularization term in the objective, namely $+\frac{\tau}{2} \|\mathbf{X} - \mathbf{X}_t\|_F^2$, improves the rate of convergence of the method by stabilizing the iterates; we make use of this in our experiments.
- The method stalls when the eigenvalues of $\mathbf{Y}_t$ are near zero (a) due to numerical instability and (b) because $\mathbf{Y}_t$ is near the boundary of $\text{Conv}(\mathcal{Y}_n^k)$. Therefore, to accelerate convergence, we require that $\lambda_{\min}(\mathbf{Y}) \geq \frac{K}{t}$ at the t th iterate, where $K \approx 10^{-2}$. In practice, this introduces very little error.
- Selecting an optimal $\mathbf{W}^{t+1}$ is generally much cheaper than selecting an optimal $\mathbf{V}^{t+1}$, since the former problem involves optimizing over n eigenvalues, rather than n^2 variables. Therefore, efficient implementations of Algorithm 2 necessarily require efficient methods for obtaining $\mathbf{V}^{t+1}$. In the case of matrix completion, $g(\mathbf{X})$ is a quadratic form, which implies that obtaining $\mathbf{V}^{t+1}$ is equivalent to solving a linear system, which we do iteratively in our numerical experiments.
- We solve for $\mathbf{V}^{t+1}$ by solving the first-order optimality condition using a successive over-relaxation linear technique, or in rare instances where the linear system solver fails to converge we use `Ipopt` to solve the QP’s first-order optimality condition.

To confirm that Algorithm 2 has indeed converged (at least approximately) to an optimal solution, we require a dual certificate. As optimizing over the set of dual variables $\boldsymbol{\alpha}$ for a fixed $\mathbf{Y}_t$ does not supply such a bound, we now invoke strong duality to derive a globally valid lower bound. Formally, we have the following result (proof deferred to Appendix A.10):

LEMMA 10. *Suppose that Assumption 2 holds. Then, strong duality holds between:*

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \max_{\boldsymbol{\alpha} \in \mathbb{R}^{n \times m}} h(\boldsymbol{\alpha}) - \frac{\gamma}{2} \sum_{i=1}^n \sum_{j=1}^n Y_{i,j} \langle \boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j \rangle, \quad (40)$$

Algorithm 2 An Accelerated Alternating Minimization Algorithm (c.f. Beck and Teboulle 2009)

Require: Initial solution $\mathbf{X}_1, \tau_1 \leftarrow 1$
 $t \leftarrow 1, T_{\max}$
repeat

 Compute $\mathbf{W}^{t+1}$ solution of $\operatorname{argmin}_{\mathbf{Y} \in \operatorname{Conv}(\mathcal{Y}_n^k)} g(\mathbf{X}_t) + \frac{1}{2\gamma} \langle \mathbf{X}_t \mathbf{X}_t^\top, \mathbf{Y}^\dagger \rangle$

 Set $\mathbf{Y}^{t+1} = \mathbf{W}^t + \frac{\tau_{t-1}}{\tau_{t+1}} (\mathbf{W}_t - \mathbf{W}_{t-1})$

 Compute $\mathbf{V}^{t+1}$ solution of $\operatorname{argmin}_{\mathbf{X} \in \mathbb{R}^{n \times m}} g(\mathbf{X}) + \frac{1}{2\gamma} \langle \mathbf{X} \mathbf{X}^\top, \mathbf{Y}_t^\dagger \rangle$

 Set $\mathbf{X}^{t+1} = \mathbf{V}^t + \frac{\tau_{t-1}}{\tau_{t+1}} (\mathbf{V}_t - \mathbf{V}_{t-1})$

 Set $\tau_{t+1} = \frac{1 + \sqrt{1 + 4\tau_t^2}}{2}$

 If $t \bmod 20 = 0$ compute dual bound at $\mathbf{Y}^{t+1}$ via Equation (41).

 $t \leftarrow t + 1$
until $t > T_{\max}$ or duality gap $\leq \epsilon$
return $\mathbf{X}_t, \mathbf{Y}_t$

$$\max_{\substack{\boldsymbol{\alpha} \in \mathbb{R}^{n \times m}, \\ \mathbf{U} \succeq \mathbf{0}, t \geq 0}} h(\boldsymbol{\alpha}) - \operatorname{tr}(\mathbf{U}) - kt \quad \text{s.t.} \quad \mathbf{U} + \mathbb{I}t \succeq \frac{\gamma}{2} \boldsymbol{\alpha} \boldsymbol{\alpha}^\top. \quad (41)$$

Lemma 10 demonstrates that Problem (1)’s semidefinite relaxation is equivalent to maximizing the dual conjugate $h(\boldsymbol{\alpha})$, minus the k largest eigenvalues of $\frac{\gamma}{2} \boldsymbol{\alpha} \boldsymbol{\alpha}^\top$. Moreover, as proven in the special case of sparse regression by Bertsimas et al. (2020), one can show that if the k th and $k+1$ th largest eigenvalues of $\boldsymbol{\alpha} \boldsymbol{\alpha}^\top$ in a solution to (40) are distinct then Problem (40)’s lower bound is tight.

4.3. Upper Bounds via Greedy Rounding

We now propose a greedy rounding method for rounding $\mathbf{Y}^*$, an optimal $\mathbf{Y}$ in a semidefinite relaxation of Problem (9), to obtain certifiably near-optimal solutions to Problem (9) quickly. Rounding schemes for approximately solving low-rank optimization problems by rounding their SDO relaxations have received a great deal of attention since they were first proposed by Goemans and Williamson (1995). Our analysis is, however, more general than typically conducted when solving low-rank problems, as it involves rounding a projection matrix $\mathbf{Y}$, rather than rounding $\mathbf{X}$, and therefore is able to generalize to the rank- k case for $k > 1$, which has historically been challenging.

Observe that for any feasible $\mathbf{Y} \in \operatorname{Conv}(\mathcal{Y}_n)$, $0 \leq \lambda_i(\mathbf{Y}) \leq 1$ for each eigenvalue of $\mathbf{Y}$, and $\mathbf{Y}$ is a projection matrix if and only if its eigenvalues are binary. Combining this observation with the Lipschitz continuity of $f(\mathbf{Y})$ in $\mathbf{Y}$ suggests that high-quality feasible projection matrices can be found in the neighborhood of a solution to the semidefinite relaxation, and a good method for obtaining them is to greedily round the eigenvalues of $\mathbf{Y}$. Namely, let $\mathbf{Y}^*$ denote a solution to the semidefinite relaxation (21), $\mathbf{Y}^* = \mathbf{U} \boldsymbol{\Lambda}^* \mathbf{U}^\top$ be a singular value decomposition of $\mathbf{Y}^*$ such that $\boldsymbol{\Lambda}^*$ is a diagonal matrix with on-diagonal entries $\Lambda_{i,i}^*$, and $\boldsymbol{\Lambda}_{\text{greedy}}$ be a diagonal matrix obtained from rounding up (to 1) k of the

highest diagonal coefficients of Λ^* , and rounding down (to 0) the $n - k$ others, with diagonal entries $\Lambda_{i,i} := (\Lambda_{\text{greedy}})_{i,i}$. We then let $\mathbf{Y}_{\text{greedy}} = \mathbf{U}\Lambda_{\text{greedy}}\mathbf{U}^\top$. We now provide guarantees on the quality of the greedily rounded solution (proof deferred to Appendix A.11):

THEOREM 5. *Let $\mathbf{Y}^*$ denote a solution to the semidefinite relaxation (21), $\mathbf{Y}^* = \mathbf{U}\Lambda\mathbf{U}^\top$ be a singular value decomposition of $\mathbf{Y}^*$, $\mathcal{R}$ denote the indices of strictly fractional diagonal entries in Λ , and $\alpha^*(\mathbf{Y})$ denote an optimal choice of α for a given $\mathbf{Y}$, i.e.,*

$$\alpha^*(\mathbf{Y}) \in \arg \max_{\alpha} \left\{ \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} h(\alpha) - \Omega^*(\alpha, \mathbf{Y}, \mathbf{V}_{11}, \mathbf{V}_{22}) \right\}.$$

Suppose that for any $\mathbf{Y} \in \mathcal{Y}_n^k$, we have $\sigma_{\max}(\alpha^(\mathbf{Y})) \leq L$. Then, any valid rounding of $\mathbf{Y}^*$ which preserves the relaxation's eigenbasis, i.e., $\mathbf{Y}_{\text{rounded}} = \mathbf{U}\Lambda_{\text{rounded}}\mathbf{U}^\top$ where $\mathbf{Y}^* = \mathbf{U}\Lambda\mathbf{U}^\top$ and Λ_{rounded} is a diagonal matrix with binary diagonal entries $\Lambda_{i,i}^{\text{rounded}}$ such that $\text{tr}(\Lambda_{\text{rounded}}) \leq k$, satisfies*

$$f(\mathbf{Y}_{\text{rounded}}) - f(\mathbf{Y}^*) \leq \frac{\gamma}{2} L^2 |\mathcal{R}| \max_{\beta \geq \mathbf{0}: \|\beta\|_1 \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i, \quad (42)$$

under the Frobenius penalty and

$$f(\mathbf{Y}_{\text{rounded}}) - f(\mathbf{Y}^*) \leq ML |\mathcal{R}| \max_{\beta \geq \mathbf{0}: \|\beta\|_1 \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i, \quad (43)$$

for the spectral penalty. Moreover, let $\mathbf{Y}_{\text{greedy}} = \mathbf{U}\Lambda_{\text{greedy}}\mathbf{U}^\top$ be an instance of $\mathbf{Y}_{\text{rounded}}$ obtained by setting $\Lambda_{i,i} = 1$ for k of the highest diagonal coefficients in Λ^ . Then, the above bounds imply that $0 \leq f(\mathbf{Y}_{\text{greedy}}) - f(\mathbf{Y}^*) \leq \epsilon$, where $\epsilon = ML \min(|\mathcal{R}|, n - k)$ for the spectral penalty and $\epsilon = \frac{\gamma}{2} \min(|\mathcal{R}|, n - k) L^2$ for the Frobenius penalty.*

This result calls for multiple remarks:

- When the relaxation gap $f(\mathbf{Y}_{\text{greedy}}) - f(\mathbf{Y}^*) = 0$, and the optimal solution to the relaxation, $\mathbf{Y}^*$, is unique, $|\mathcal{R}| = 0$. This justifies retaining $\mathcal{R}$ in the bound, rather than replacing it with n .
- The techniques introduced in Section 3.5 for computing an M so that an optimal solution $\mathbf{X}^*$ obeys $\|\mathbf{X}^*\|_\sigma \leq M$, also apply to computing an explicit L such that $\|\alpha^*\|_\sigma \leq L$ in the above bound.
- The rounding technique is robust, because it minimizes the worst-case Lipschitz upper bound, under the assumption $\sigma_{\max}(\alpha^*) \leq L$ (i.e., we have no information about which coordinate⁸ has the largest Lipschitz upper bound). For instance, under Frobenius regularization the bound is

$$f(\mathbf{Y}_{\text{rounded}}) - f(\mathbf{Y}^*) \leq \frac{\gamma}{2} L^2 |\mathcal{R}| \max_{\beta \geq \mathbf{0}: \|\beta\|_1 \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i, \quad (44)$$

which is minimized over $\Lambda^{\text{rounded}} : \text{tr}(\Lambda^{\text{rounded}}) \leq k$ by solving:

$$\min_{\lambda \in \mathcal{S}_n^k} \max_{\beta \geq \mathbf{0}: \|\beta\|_1 \leq 1} \frac{\gamma}{2} L^2 |\mathcal{R}| \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \lambda_i) \beta_i, \quad (45)$$

i.e., rounding greedily. This interpretation suggests that greedy rounding never performs too badly.

⁸ If we had this information then, as the proof of Theorem 5 suggests, we would greedily round to one k of the indices with the largest values of $L_i \Lambda_{i,i}^*$.

To improve the greedily rounded solution, we implement a local search strategy which obtains even higher quality warm-starts. Namely, a variant of the popular Burer-Monterio (BM) heuristic (Burer and Monteiro 2003), which seeks low-rank solutions $\mathbf{X}$ by applying a non-linear factorization $\mathbf{X} = \mathbf{U}\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{n \times l}$, $\mathbf{V} \in \mathbb{R}^{m \times k}$ and iteratively optimizing over $\mathbf{U}$ for a fixed $\mathbf{V}$ (resp. $\mathbf{V}$ for a fixed $\mathbf{U}$) until convergence to a local optima occurs. This strategy improves our greedily rounded solution because we initially set $\mathbf{U}$ to be the square root of $\mathbf{Y}_{greedy}$ and optimize over $\mathbf{V}$; recall that if $\mathbf{Y}$ is a projection matrix we have $\mathbf{Y} = \mathbf{U}\mathbf{U}^\top$ and $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ for some singular value decomposition $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}^\top$.

5. Numerical Experiments

In this section, we evaluate the algorithmic strategies derived in the previous section, implemented in Julia 1.3 using JuMP.jl 0.20.1, Gurobi 9.0.1 to solve the non-convex QCQO master problems⁹, and Mosek 9.1 to solve the conic subproblems/continuous relaxations. Except where indicated otherwise, all experiments were performed on a Intel Xeon E5—2690 v4 2.6GHz CPU core using 32 GB RAM. To bridge the gap between theory and practice, we have made our code freely available on Github at github.com/ryancorywright/MixedProjectionSoftware.

We evaluate the different ingredients of our numerical strategy on a matrix completion example: First, we solve the semidefinite relaxation by implementing Algorithm 2 and demonstrate its increased scalability over Mosek’s IPM in Section 5.1. From the solution of the relaxation, our rounding and local search heuristics then provide near-optimal solutions that outperform state-of-the-art heuristic methods, as discussed in Section 5.2. We implement Algorithm 1, benchmark its performance and, for the first time, solve low-rank matrix completion to certifiable optimality in Section 5.3. In Section 5.4, we explore the role which regularization plays in our numerical strategy, by showing that increasing the amount of regularization in Problem (1) decreases the relative gap, the problem’s complexity, and the amount of time required to solve the problem to optimality. Finally, in Section 5.5 we solve sensor location problems to certifiable optimality.

5.1. Exploring the Scalability of the Convex Relaxations

In this section, we explore the relative scalability of Mosek’s interior point method and Algorithm 2.

We consider convex relaxations of matrix completion problems. Similarly to Candès and Plan (2010), we generate two low-rank matrices $\mathbf{M}_L, \mathbf{M}_R \in \mathbb{R}^{n \times r}$ with i.i.d. $\mathcal{N}(0, 1)$ entries, and attempt to recover the matrix $\mathbf{M} = \mathbf{M}_L \mathbf{M}_R^\top$ given a proportion p of its observations. Here, we fix $p = 0.25$ and $k = r = 5$, vary n , and set $\gamma = \frac{20}{p}$ where we scale γ proportionally to $1/p$ so that the relative importance of $\|\mathbf{X}\|_F^2$ and $\sum_{(i,j) \in \Omega} (X_{i,j} - A_{i,j})^2$ remains constant with p .

⁹ We remark that Gurobi solves the non-convex QCQO master problems by translating them to piecewise linear optimization problems. Since rank constraints are not MICO representable, this introduces some error. To mitigate against this error, we set the Gurobi parameters FuncPieceError and FuncPieceLength to their minimum possible values (10^{-6} and 10^{-5} respectively). Additionally, we set NonConvex to 2, and otherwise use default Gurobi/Mosek parameters.

We solve the continuous relaxation

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}, \mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k), \boldsymbol{\theta} \in S^n} \frac{1}{2\gamma} \text{tr}(\boldsymbol{\theta}) + \sum_{(i,j) \in \Omega} (X_{i,j} - A_{i,j})^2 \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{\theta} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{Y} \end{pmatrix} \succeq \mathbf{0}. \quad (46)$$

Table 3 reports the time required by Algorithm 2 to obtain a solution with a relative duality gap of 0.1%. To evaluate numerical stability, we also report the relative MSE of the greedily rounded solution; experiments where $n \leq 250$ were run on a standard MacBook pro with 16GB RAM, while larger experiments were run on the previously described cluster with 100GB RAM.

Table 3 Scalability of convex relaxations, averaged over 5 matrices. Problem is regularized with Frobenius norm and $\gamma = \frac{20}{p}$. “-” indicates an instance could not be solved with the supplied memory budget.

n	Mosek		Algorithm 2		n	Algorithm 2	
	Relative MSE	Time (s)	Relative MSE	Time (s)		Relative MSE	Time (s)
50	0.429	2.28	0.438	17.28	350	0.058	6,970
100	0.138	47.20	0.139	79.01	400	0.056	8,096
150	0.082	336.1	0.081	228.7	450	0.055	26,350
200	0.0675	1,906	0.067	841.7	500	0.054	28,920
250	-	-	0.062	1,419	550	0.0536	39,060
300	-	-	0.059	2,897	600	0.0525	38,470

Our results demonstrate the efficiency of Algorithm 2: the relative MSE is comparable to **Mosek**’s, but computational time does not explode with n . Since it does not require solving any SDOs and avoids the computational burden of performing the Newton step in an IPM, Algorithm 2 scales beyond $n = 600$ (1,440,000 decision variables), compared to $n = 200$ for IPMs (80,000 decision variables).

5.2. Numerical Evaluation of Greedy Rounding on Matrix Completion Problems

In this section, we compare the greedy rounding method with state-of-the-art heuristic methods, and demonstrate that, by combining greedy rounding with the local search heuristic of (Burer and Monteiro 2003), our approach outperforms state-of-the-art heuristic methods and therefore should be considered as a viable and efficient warm-start for Algorithm 1.

We consider the previous matrix completion problems and assess the ability to recover the low-rank matrix $\mathbf{M}$ (up to a relative MSE of 1%), for varying fraction of observed entries p and rank r , with $n = 100$ fixed. Note that, other than the inclusion of a Frobenius regularization term, this is the same experimental setup considered by Candès and Recht (2009), Recht et al. (2010) among others.

We compare the performance of four methods: the greedy rounding method, both with and without the local improvement heuristic from Burer and Monteiro (2003), against the local improvement heuristic alone (with a thresholded-SVD initialization point) and the nuclear norm approach. Specifically, the greedy rounding method takes the solution of the previous convex relaxation with $\gamma = \frac{500}{p}$ and rounds its singular values to generate a feasible solution $\mathbf{Y}_{\text{greedy}}$. For the local improvement heuristic, we solve:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}, \mathbf{U}, \mathbf{V} \in \mathbb{R}^{n \times k}} \frac{1}{2\gamma} \|\mathbf{X}\|_2^2 + \sum_{(i,j) \in \Omega} (X_{i,j} - A_{i,j})^2 \quad \text{s.t.} \quad \mathbf{X} = \mathbf{U}\mathbf{V}^\top,$$

Figure 2 Prop. matrices recovered with $\leq 1\%$ relative MSE (higher is better), for different values of p (x-axis) and $r(2n-r)/pn^2 \propto 1/n$ (y-axis), averaged over 25 rank- r matrices.

for $\gamma = \frac{500}{p}$ and $k = r$, and iteratively optimize over $\mathbf{U}$ and $\mathbf{V}$ using *Mosek*. We provide an initial value for $\mathbf{U}$ by either taking the first k left-singular vectors of a matrix $\mathbf{A}$ where unobserved entries are replaced by 0, or taking the square root of $\mathbf{Y}_{\text{greedy}}$. For the nuclear norm regularization strategy, since our observations are noiseless, we solve: $\min_{\mathbf{X} \in \mathbb{R}^{n \times n}} \|\mathbf{X}\|_*$ s.t. $X_{i,j} = A_{i,j} \quad \forall (i,j) \in \Omega$.

Figure 2 depicts the proportion of times the matrix was recovered exactly (averaged over 25 samples per tuple of (n, p, r)), while Figure 3 depicts the relative average MSE over all instances. As in Candès and Recht (2009), Recht et al. (2010), we vary p between 0 and 1 and consider all possible ranks r such that $r(2n-r) \leq pn^2$. From this set of experiments, we make several observations: First, greedy rounding and the local improvement heuristic outperform nuclear norm minimization both in terms of average relative MSE and amount of data required to recover the matrix. Second, the local improvement heuristic improves upon greedy rounding. In terms of its ability to recover the underlying matrix exactly, it performs equally well with either initialization strategy. However, initialization with the greedy rounding supplies dramatically lower average MSEs in instances where no approach recovers the true matrix exactly. This suggests that initialization strategies for the Burer-Monterio heuristic should be revisited and greedy rounding considered as a viable and more accurate alternative than selecting a random feasible point.

5.3. Benchmarking Algorithm 1 on Synthetic Matrix Completion Problems

We now benchmark Algorithm 1 on matrix completion problems where $n \in \{10, 20, 30\}$.

Figure 3 Average relative MSE (lower is better), averaged over 25 rank- r matrices. We cap the relative MSE at 1.0.

We first compare the two different implementations of Algorithm 1, single- and multi-tree, with solving the problem directly as a QCQO in **Gurobi** 9.0 (Section 4.1.2). In Algorithm 1, the lower bounds are warm-started with 200 cuts from the in-out method, and greedy rounding with local search improvement is used for the upper bounds; if a single-tree instance fails to find a feasible solution (due to numerical instability in **Gurobi**) we return the gap between the warm-start and the semidefinite relaxation. At the t th iteration, we impose a time limit of $10t$ seconds for generating the new cut so as to increase numerical precision as the solver progresses. We also impose a limit of 20 cuts for the multi-tree approach, a time limit of 30,000s for the single-tree approach¹⁰, a time limit of 3,600s for **Gurobi**, and an optimality gap of 1%¹¹. Average runtime, number of nodes, and optimality gap are reported in Table 4. Note that the same random instances were solved by all three approaches (by fixing the random seeds), to facilitate a less noisy comparison.

We observe that Algorithm 1 drastically improves upon **Gurobi** both in terms of computational time (reduced by up to an order of magnitude) and accuracy (absolute gap reduced by around an order of

¹⁰ We require a larger time limit than 3,600s, since **Gurobi** often fails to find any feasible solutions within this time limit due to the numerical difficulties inherent in integrating lazy constraint callbacks and a non-convex master problem.

¹¹ We report the absolute gap between the better of **Gurobi**'s lower bound and the semidefinite lower bound, compare to the objective value which we evaluate directly; this is sometimes 1 – 2% even when **Gurobi** reports that it has found an optimal solution, due to numerical instability in **Gurobi**. Note that we report the absolute, rather than relative, gap since the relative gap depends on the quality of **Gurobi**'s approximation of $\mathcal{Y}_n^k$, which is controlled by the parameter `FuncPieceError` and cannot be set lower than 10^{-6} ; also note that the objective values are on the order of 0.5-5.0 for the problems reported in Table 4.

magnitude). Multi-tree dominates single-tree and Gurobi in terms of runtime and the quality of the solution found, although single-tree occasionally has a smaller gap at termination. Moreover, multi-tree consistently finds high-quality feasible solutions earlier than single tree and accepts our warm-start more consistently, which suggests it may scale better to high-dimensional settings.

Table 4 Scalability of Algorithm 1 vs. Gurobi for solving rank-1 matrix completion problems to certifiable optimality, averaged over 20 random matrices per row. In multi-tree, Nodes denotes the number of nodes expanded in the final branch-and-cut tree, while Nodes (t) denotes the number of nodes expanded over all trees for the multi-tree implementation.

n	p	γ	Gurobi (direct)			Algorithm 1 (single-tree)				Algorithm 1 (multi-tree)				
			Time(s)	Nodes	Gap	Time(s)	Nodes	Gap	Cuts	Time(s)	Nodes	Nodes (t)	Gap	Cuts
10	0.1	$20/p$	> 3,600	313,700	0.0301	10,310	40,060	0.0004	23,460	252.3	630.2	10,099	0.0019	2.95
10	0.2	$20/p$	> 3,600	299,200	0.0854	19,440	28,430	0.0229	19,370	1,672	2,277	28,895	0.0104	11.0
10	0.3	$20/p$	> 3,600	274,500	0.1167	20,368	25,480	0.0433	20,290	2,319	2,684	30,906	0.0317	15.4
10	0.1	$100/p$	> 3,600	281,800	0.0068	18,580	62,680	0.0015	42,200	239.5	405	7,499	0.0003	3.20
10	0.2	$100/p$	> 3,600	271,300	0.0178	27,990	39,330	0.0492	31,060	1,269	1,931	21,420	0.0042	8.40
10	0.3	$100/p$	3,239	237,000	0.0178	25,750	29,350	0.0434	23,390	2,472	2,134	32,196	0.0098	19.6
20	0.1	$20/p$	> 3,600	80,760	0.8915	> 30,000	13,110	0.741	13,070	2,917	413.2	4,835	0.0166	18.5
20	0.2	$20/p$	> 3,600	65,310	4.094	> 30,000	7,023	0.1816	7,008	3,512	143.0	1,746	0.247	20.0
20	0.3	$20/p$	> 3,600	64,850	4.745	28,700	6,914	0.1066	6,828	3,287	204.9	2,081	0.253	19.6
20	0.1	$100/p$	> 3,600	60,830	0.428	> 30,000	13,790	0.7714	13,799	3,013	508.0	6,152	0.0072	17.6
20	0.2	$100/p$	> 3,600	43,850	1.421	> 30,000	6,395	0.0543	6,395	3,106	80.9	1,027	0.0903	17.6
20	0.3	$100/p$	> 3,600	55,150	2.810	29,530	6,538	0.0271	6,510	2,910	62.7	744.2	0.1368	17.0

Next, we evaluate the performance of the multi-tree implementation of Algorithm 1 on a more extensive test-set, including instances where $\text{Rank}(\mathbf{M}) > 1$, in Table 5. Note that when $r = 1$ we use the same experimental setup (although we impose a time limit of $30t$ seconds, or 7200 seconds if there has been no improvement for two consecutive iterations, a cut limit of 50 cuts when $n > 20$), and when $r > 1$ we increase the time limit per iteration to $300t$ seconds (or 7200 seconds if there has been no improvement for two consecutive iterations), and allow up to 100 PSD cuts per iteration to be added at the root node via a user cut callback, in order to strengthen the approximation of the PSD constraint $\mathbf{Y} \succeq \mathbf{0}$. We observe that the problem’s complexity increases with the rank, although not too excessively. Moreover, when $r > 1$ the bound gap is actually smaller when $\gamma = \frac{100}{p}$ than when $\gamma = \frac{20}{p}$. We believe this is because Gurobi cannot represent the SDO constraint $\mathbf{Y} \succeq \mathbf{0}$ and its SOC approximation is inexact (even with PSD cuts), and in some cases refining this approximation is actually harder than refining our approximation of $g(\mathbf{X})$.

Note that the main bottleneck inhibiting solving matrix completion problems where $n \geq 50$ is Gurobi itself, as the non-convex solver takes increasing amounts of time to process warm-starts (sometimes in the 100s or 1000s of seconds) when n increases. We believe this may be because of the way Gurobi translates orthogonal projection matrices to a piecewise linear formulation. Encouragingly, this suggests that our approach may successfully scale to 100×100 matrices as Gurobi improves their solver.

Table 5 Scalability of Algorithm 1 (multi-tree) for solving low-rank matrix completion problems to certifiable optimality, averaged over 20 random matrices per row.

n	p	γ	Rank-1				Rank-2				Rank-3			
			Time(s)	Nodes	Gap	Cuts	Time(s)	Nodes	Gap	Cuts	Time(s)	Nodes	Gap	Cuts
10	0.1	$20/p$	182.1	9,755	0.0005	2.56	24,220	35,670	0.0034	5.78	37,780	39,870	0.0071	9.28
10	0.2	$20/p$	3,508	21,060	0.0026	10.8	209,900	108,000	0.0252	35.3	135,260	35,870	0.031	26.2
10	0.3	$20/p$	5,488	30,970	0.0039	13.1	302,200	70,500	0.0866	50.0	302,100	31,870	0.0197	50.0
10	0.1	$100/p$	656.5	28,870	0.0001	2.14	676.1	25,493	0.0009	1.83	842.7	20,700	0.0024	1.79
10	0.2	$100/p$	1,107	10,010	0.0009	4.29	2,065	42,490	0.0019	5.61	57,530	36,910	0.0124	10.7
10	0.3	$100/p$	3,364	48,730	0.0022	6.30	272,300	33,150	0.0195	44.7	249,700	35,530	0.0499	42.2
20	0.1	$20/p$	2,017	4,756	0.0061	8.20	253,900	8,030	0.0279	42.7	255,400	3,015	0.0309	43.2
20	0.2	$20/p$	6,369	6,636	0.0136	15.0	298,700	3,342	0.549	50.0	295,500	236.5	0.879	50.0
20	0.3	$20/p$	6,687	4,187	0.0082	18.4	296,500	3,175	1.123	50.0	291,100	41.35	2.147	50.0
20	0.1	$100/p$	1,266	8,792	0.0087	8.35	211,700	6,860	0.0073	34.24	171,900	2,350	0.0131	29.8
20	0.2	$100/p$	1,220	2,710	0.0104	7.80	302,800	2,426	0.123	50.0	298,800	221.4	0.123	50.0
20	0.3	$100/p$	1,272	1,837	0.0064	3.14	299,000	2,518	0.264	50.0	293,500	43.0	0.659	50.0
30	0.1	$20/p$	300,300	2,735	0.0905	50.0	304,300	164.0	0.790	50.0	303,100	1.10	0.365	50.0
30	0.2	$20/p$	298,700	1,511	0.136	50.0	301,700	9.62	3.105	50.0	302,600	1.00	5.581	50.0
30	0.3	$20/p$	183,800	1,743	0.0476	36.9	303,000	1.63	5.232	50.0	305,000	0.70	14.60	50.0
30	0.1	$100/p$	305,600	2,262	0.0273	50.0	302,800	97.40	0.0973	50.0	305,000	1.90	0.0967	50.0
30	0.2	$100/p$	246,300	3,285	0.0315	43.6	304,300	6.17	0.697	50.0	302,600	1.00	1.419	50.0
30	0.3	$100/p$	25,970	11,020	0.0089	17.1	304,000	1.00	0.923	50.0	304,700	1.00	3.221	50.0

(a) $n = 25$ (b) $n = 50$ **Figure 4** Average relative MSE for nuclear norm (NN), greedy rounding (GD), Burer-Monterio (BM), and outer-approximation (OA) when imputing a rank-1 $n \times n$ matrix. All results are averaged over 25 matrices.

Finally, we compare the solution from the exact formulation (9) solved using Algorithm 1 (multi-tree) with the initial warm-start we proposed and two state-of-the-art heuristics, namely nuclear norm minimization and the Burer-Monterio approach, as in Section 5.2. Here, we take $n \in \{25, 50\}$, $r = 1$, p ranging from 0 to 0.4, and $\gamma = \frac{100}{p}$. Figure 4 depicts the average relative MSE over the entire matrix, averaged over 25 random instances per value of p . When $p \geq 0.2$, the exact method supplies an out-of-sample relative MSE around 0.6% lower than Burer-Monterio¹².

¹² Because we ran all methods on the same random instances, this difference is statistically significant, with a p-value of 2×10^{-51} (resp. 2×10^{-129}) that the relative MSE is lower for the exact method when $n = 25$ (resp. $n = 50$).

Figure 5 Average relative in-sample bound gap (%), averaged over 25 rank- r matrices.

5.4. Exploring the Impact of Regularization on Problem Complexity

We now examine the impact of the regularization term $\frac{1}{2\gamma}\|\mathbf{X}\|_F^2$ on the problem complexity, as captured by the relative in-sample duality gap between the semidefinite relaxation and the objective value of the greedy solution with a BM local improvement heuristic. We generate the problem data in the same manner as the previous experiment, and display results for four values of γ in Figure 5. Observe that as γ increases, both the duality gap and the problem's complexity increase. This observation confirms similar results on the impact of regularization in mixed-integer conic optimization problems (c.f. Bertsimas and Cory-Wright 2018, Bertsimas et al. 2019). Additionally, when $\gamma = \frac{500}{p}$ in Figure 5(d), the region where the in-sample duality gap is zero corresponds to exactly recovering the underlying matrix with high probability, while a strictly positive duality gap corresponds to instances with partial recovery only (see Figure 2). This suggests a deep connection between relaxation tightness and statistical recovery.

While the relative in-sample semidefinite relaxation gap is a theoretical measure of problem difficulty, it does not indicate how fast Algorithm 1 converged in practice. In this direction, we solve the 20 synthetic matrix completion problems considered in Table 4 where $n \in \{10, 20\}$, $r = 1$, $p \in \{0.2, 0.3\}$ for 20 different values of $\gamma \in [10^0, 10^4]$ (distributed uniformly on a log-scale), and compare the relative in-sample semidefinite gap (greedily rounded solution vs. semidefinite bound) with Algorithm 1's runtimes in Figure 6, for the single-tree (left panel) and multi-tree (right panel) implementation. Results are averaged over 20 random synthetic instances per value of γ . We observe that the relaxation gap does correlate with runtime for single-tree. Yet, the relationship between the relaxation gap and runtime

is less straightforward for multi-tree, as it depends on how Gurobi balances cut generation and node expansion, and the conditioning of the problem.

Figure 6 Average runtime against relative semidefinite relaxation gap for Algorithm 1 single-tree (left) and multi-tree (right) over 20 synthetic matrix completion instances per data point, where $p \in \{0.2, 0.3\}$, $r = 1$, $n \in \{10, 20\}$.

The regularizer γ also impact the bias term $\frac{1}{2\gamma} \|\mathbf{X}\|_F^2$ added to the objective function, hence the suboptimality of the solution. To further illustrate the impact of the regularizer γ on solve times and the trade-off between tractability and sub-optimality, Figure 7 reports the average runtime and MSE for the previously solved instances, as a function of γ . Figure 7 illustrates how γ balances tractability (runtime, top row) and optimality of the solution (MSE, bottom row). Also, single-tree (left panel) is one order of magnitude slower than multi-tree (right panel), and is also more numerically unstable when γ increases, largely because of the difficulty of combining a non-convex master problem and lazy constraint callbacks (which imposes many cuts, without processing the implications of these cuts as quickly). Echoing our findings in the previous section, this suggests that, while in MICO single-tree typically outperforms multi-tree, at the current state of technology multi-tree should be considered as a viable and potentially more efficient alternative for matrix completion problems which have non-convex master problems. However, as the algorithmic implementations of non-convex QCQP solvers mature, this finding should be revisited.

5.5. Benchmarking Algorithm 1 on Synthetic Coordinate Recovery Problems

We now benchmark the performance of Algorithm 1 on anchor-free synthetic coordinate recovery problems, as previously studied by Biswas and Ye (2004), Luo et al. (2010) among others.

Specifically, we sample n coordinates $\mathbf{x}_i$ uniformly over $[-0.5, 0.5]^k$ for $k \in \{2, 3\}$, and attempt to recover a noisy Gram matrix $\mathbf{G} \in S_+^n$ of the $\mathbf{x}_i$'s, given a subset of observations of the underlying matrix. Similarly to Biswas and Ye (2004), we supply the distance between the points $D_{i,j} = \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 + z$, where $z \sim \mathcal{N}(0, 0.01)$, if and only if the radio range between the two points is such that $D_{i,j} \leq d_{radio}^2$.

Figure 7 Average runtime (top) and MSE (bottom) vs. γ for Algorithm 1 single-tree (left) and multi-tree (right) implementations over 20 synthetic matrix completion instances where $p \in \{0.2, 0.3\}$, $r = 1$ and $n \in \{10, 20\}$. The same random seeds were used to generate random matrices completed by single-tree and multi-tree.

Note that we solve these problems in precisely the same fashion as the largest matrix completion problems solved in the previous section (multi-tree, with a limit of 50 cut passes etc.)

Formally, in order to account for noise in the observed entries, we solve the following problem:

$$\min_{\mathbf{Y} \in \mathcal{Y}_h^k} \min_{\mathbf{G} \in \mathcal{S}_+^n} \frac{1}{2\gamma} \|\mathbf{G}\|_F^2 + \text{tr}(\mathbf{G}) + \lambda \cdot \|\boldsymbol{\xi}\|_1 \quad \text{s.t.} \quad G_{i,i} + G_{j,j} - 2G_{i,j} + \xi_{i,j} = D_{i,j} \quad \forall (i,j) \in \Omega, \quad \mathbf{G} = \mathbf{Y}\mathbf{G},$$

where $\lambda > 0$ is a penalty term which encourages robustness, and the Frobenius norm objective likewise encourages robustness against noise in $\mathbf{G}$. The performance of Algorithm 1 (multi-tree) on various synthetic instances is reported in Table 6, for $\gamma, n, d_{\text{radio}}, k$ varying.

We observe that the problem's complexity increases with the rank and with the dimensionality of the Gram matrix, although not too excessively. Indeed, Algorithm 1 can solve coordinate recovery problems with tens of data points to certifiable optimality in hours.

5.6. Summary of Findings from Numerical Experiments

Our main findings from the numerical experiments in this section are as follows:

Table 6 Scalability of Algorithm 1 (multi-tree) for solving sensor location problems to certifiable optimality, averaged over 20 random instances per row. A “-” denotes an instance that cannot be solved within the time budget, because Gurobi fails to accept our warm-start and cannot find a feasible solution. We let $\lambda = n^2$ for all instances.

n	d_{radio}	γ	Rank-2				Rank-3			
			Time(s)	Nodes	Gap	Cuts	Time(s)	Nodes	Gap	Cuts
10	0.1	$1/p$	135.3	6,926	0.0001	1.00	45.14	0.02	0.0000	1.00
10	0.2	$1/p$	3,189	5,249	0.0024	11.5	216.8	7,819	0.0022	1.00
10	0.1	$100/p$	76.2	1,155	0.0000	1.00	140.6	950	0.0000	1.00
10	0.2	$100/p$	480.6	0.05	0.0001	21.7	92.6	139	0.0000	1.14
20	0.1	$1/p$	3,475	4,548	0.0007	13.0	3,090	9,740	0.0001	1.00
20	0.2	$1/p$	73,000	0.50	0.0149	50.0	7,173	5,313	0.0038	1.20
20	0.1	$100/p$	1,878	0.00	0.0000	3.91	64.9	0.00	0.0000	1.07
20	0.2	$100/p$	67,530	0.20	0.0044	50.0	55.7	0.00	0.0002	1.00

- As demonstrated in Section 5.1, Algorithm 2 successfully solves convex relaxations of low-rank problems where $n = 100$ s, in a faster and more scalable fashion than state-of-the-art interior point codes such as *Mosek*.
- As demonstrated in Section 5.2, the following strategy is almost as effective as solving a low-rank problem exactly: solving the convex relaxation using Algorithm 2, greedily rounding the solution to the convex relaxation, and using this greedily rounded solution as a warm-start for a local method such as the method of Burer and Monteiro (2003, 2005). The success of this strategy can be explained because solving a relaxation and rounding is a global strategy which matches the low-order moments of the set of optimal solutions to obtain a solution near an optimal solution, while local methods polish a solution by seeking the best solution within some neighborhood of an initial point.
- As demonstrated in Section 5.4, increasing the amount of regularization in a low-rank problem by decreasing γ decreases the duality gap between a low-rank problem with Frobenius or spectral norm problem, and its convex relaxation. Therefore, increasing the amount of regularization makes the problem easier in a practical sense (although not necessarily in a complexity-theoretic sense).
- As demonstrated in Sections 5.3, 5.5, Algorithm 1 scales to solve problems where n is in the tens, i.e., hundreds or thousands of decision variables, in hours. Moreover, the main bottleneck inhibiting solving problems where n is in the hundreds or thousands is that we solve our master problems using Gurobi, a QCQO solver which translates the orthogonal projection matrix constraint into many piecewise linear constraints. This suggests that a custom branch-and-bound solver which explicitly models orthogonal projection matrices constitutes a promising area for future work.

6. Conclusion

In this paper, we introduced Mixed-Projection Conic Optimization, a new framework for modeling rank constrained optimization problems that, for the first time, solves low-rank problems to certifiable optimality at moderate problem sizes. We also provided a characterization of the complexity of rank

constraints, and proposed new convex relaxations and rounding methods that lead to viable and more accurate solutions than those obtained via existing techniques such as the log-det or nuclear norm heuristics. Inspired by the collective successes achieved in mixed-integer optimization, we hope that MPCO constitutes an exciting new research direction for the optimization community. For instance, we believe that custom branch-and-bound solvers that explicitly model orthogonal projection matrices could further enhance the scalability of the MPCO framework.

Acknowledgments: We thank the two anonymous referees and the associate editor for many valuable comments which improved the paper substantially.

References

- T. Achterberg and E. Towle. Gurobi webinar: Non-convex quadratic optimization. <https://www.gurobi.com/resource/non-convex-quadratic-optimization/>, 2020. Accessed: 2021-02-09.
- A. A. Ahmadi and A. Majumdar. DSOS and SDSOS optimization: More tractable alternatives to sum of squares and semidefinite optimization. *SIAM J. Appl. Alg. Geom.*, 3(2):193–230, 2019.
- M. S. Aktürk, A. Atamtürk, and S. Gürel. A strong conic quadratic reformulation for machine-job assignment with controllable processing times. *Oper. Res. Letters*, 37(3):187–191, 2009.
- A. Atamtürk and A. Gomez. Rank-one convexification for sparse regression. *arXiv:1901.10334*, 2019.
- C. Audet, P. Hansen, B. Jaumard, and G. Savard. A branch and cut algorithm for nonconvex quadratically constrained quadratic programming. *Math. Prog.*, 87(1):131–152, 2000.
- L. Bai, J. E. Mitchell, and J.-S. Pang. On conic QPCCs, conic QCQPs and completely positive programs. *Math. Prog.*, 159(1):109–136, 2016.
- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imag. Sci.*, 2(1):183–202, 2009.
- R. M. Bell and Y. Koren. Lessons from the Netflix prize challenge. Technical report, AT&T Bell Laboratories, 2007.
- P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan. Mixed-integer nonlinear optimization. *Acta Numerica*, 22:1, 2013.
- W. Ben-Ameur and J. Neto. Acceleration of cutting-plane and column generation algorithms: Applications to network design. *Networks*, 49(1):3–17, 2007.
- A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization*, volume 2. SIAM, 2001.
- D. P. Bertsekas. *Nonlinear programming*. Athena Scientific Belmont, 3rd edition, 2016.
- D. Bertsimas and R. Cory-Wright. A scalable algorithm for sparse portfolio selection. *arXiv preprint arXiv:1811.00138*, 2018.
- D. Bertsimas and R. Cory-Wright. On polyhedral and second-order cone decompositions of semidefinite optimization problems. *Oper. Res. Lett.*, 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, A. King, and R. Mazumder. Best subset selection via a modern optimization lens. *Ann. Stat.*, pages 813–852, 2016.
- D. Bertsimas, M. S. Copenhaver, and R. Mazumder. Certifiably optimal low rank factor analysis. *J. Mach. Learn. Res.*, 18(1):907–959, 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.

- D. Bertsimas, J. Pauphilet, and B. van Parys. Sparse regression: Scalable algorithms and empirical performance. *Statistical Science*, 35(4):555–578, 2020.
- S. Bi, S. Pan, and D. Sun. A multi-stage convex relaxation approach to noisy structured low-rank matrix recovery. *Math. Prog. Comp.*, 12(4):569–602, 2020.
- P. Biswas and Y. Ye. Semidefinite programming for ad hoc wireless sensor network localization. In *Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 46–54. ACM, 2004.
- G. Blekherman, P. A. Parrilo, and R. R. Thomas. *Semidefinite optimization and convex algebraic geometry*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2012.
- P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter. An algorithmic framework for convex mixed integer nonlinear programs. *Disc. Optim.*, 5(2):186–204, 2008.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. SIAM Philadelphia, PA, 1994.
- S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Prog.*, 95(2):329–357, 2003.
- S. Burer and R. D. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Math. Prog.*, 103(3):427–444, 2005.
- E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comp. Math.*, 9(6):717, 2009.
- J. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. Symp. T. Comp.*, pages 460–467. ACM, 1988.
- A. L. Chistov and D. Y. Grigoriev. Complexity of quantifier elimination in the theory of algebraically closed fields. In *International Symposium on Mathematical Foundations of Computer Science, volume 176 of Lecture Notes in Computer Science*, pages 17–31. Springer Verlag, 1984.
- C. Coey, M. Lubin, and J. P. Vielma. Outer approximation with conic certificates for mixed-integer convex problems. *Math. Prog. Comp.*, 12:249–293, 2020.
- D. Cox, J. Little, and D. O’Shea. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*. Springer Science & Business Media, 2013.
- A. d’Aspremont. A semidefinite representation for some minimum cardinality problems. In *42nd IEEE International Conference on Decision and Control*, volume 5, pages 4985–4990. IEEE, 2003.
- A. d’Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. Lanckriet. A direct formulation for sparse pca using semidefinite programming. *SIAM Rev.*, 49(3):434–448, 2007.
- C. Ding, D. Sun, and J. Y. Jane. First order optimality conditions for mathematical programs with semidefinite cone complementarity constraints. *Math. Prog.*, 147(1):539–579, 2014.
- H. Dong, K. Chen, and J. Linderoth. Regularization vs. relaxation: A conic optimization perspective of statistical variable selection. *arXiv:1510.06083*, 2015.
- J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American mathematical Society*, 82(2):421–439, 1956.
- M. A. Duran and I. E. Grossmann. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. *Math. Prog.*, 36(3):307–339, 1986.
- J. Eckstein and D. P. Bertsekas. On the Douglas—Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math. Prog.*, 55(1):293–318, 1992.
- J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.

- M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002.
- M. Fischetti, I. Ljubić, and M. Sinnl. Redesigning Benders decomposition for large-scale facility location. *Mang. Sci.*, 63(7):2146–2162, 2017.
- A. Frangioni and C. Gentile. SDP diagonalizations and perspective cuts for a class of nonseparable MIQP. *Oper. Res. Lett.*, 35(2):181–185, 2007.
- F. Glover. Improved linear integer programming formulations of nonlinear integer problems. *Mang. Sci.*, 22(4):455–460, 1975.
- 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.
- I. E. Grossmann. Review of nonlinear mixed-integer and disjunctive programming techniques. *Opt. & Eng.*, 3(3):227–252, 2002.
- M. Grötschel and Y. Wakabayashi. A cutting plane algorithm for a clustering problem. *Math. Prog.*, 45(1-3):59–96, 1989.
- O. Günlük and J. Linderoth. Perspective reformulation and applications. In *Mixed Integer Nonlinear Programming*, pages 61–89. Springer, 2012.
- H. Hazimeh, R. Mazumder, and A. Saab. Sparse regression at scale: Branch-and-bound rooted in first-order optimization. *arXiv preprint arXiv:2004.06152*, 2020.
- R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, New York, 1985.
- P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proc. forty-fifth ACM symp. T. Comput.*, pages 665–674, 2013.
- R. C. Jeroslow. There cannot be any algorithm for integer programming with quadratic constraints. *Oper. Res.*, 21(1):221–224, 1973.
- J. E. Kelley, Jr. The cutting-plane method for solving convex programs. *J. Soc. Ind. Appl. Math.*, 8(4):703–712, 1960.
- T. Kleinert, M. Labbé, F. a. Plein, and M. Schmidt. There’s no free lunch: on the hardness of choosing a correct big-M in bilevel optimization. *Oper. Res.*, 68(6):1716–1721, 2020.
- B. Kocuk, S. S. Dey, and X. A. Sun. Strong SOCP relaxations for the optimal power flow problem. *Oper. Res.*, 64(6):1177–1196, 2016.
- J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Opt.*, 11(3):796–817, 2001.
- J. Lavaei and S. H. Low. Zero duality gap in optimal power flow problem. *IEEE Trans. Power Sys.*, 27(1):92–107, 2011.
- J. Lee and B. Zou. Optimal rank-sparsity decomposition. *J. Glob. Opt.*, 60(2):307–315, 2014.
- J. Linderoth. A simplicial branch-and-bound algorithm for solving quadratically constrained quadratic programs. *Math. Prog.*, 103(2):251–282, 2005.
- M. Lubin, I. Zadik, and J. P. Vielma. Mixed-integer convex representability. *Math. Oper. Res.*, to appear, 2021.
- Z.-Q. Luo, J.-S. Pang, and D. Ralph. *Mathematical programs with equilibrium constraints*. Cambridge University Press, 1996.
- Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang. Semidefinite relaxation of quadratic optimization problems. *IEEE Sig. Proc. Mag.*, 27(3):20–34, 2010.
- G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part i—convex underestimating problems. *Math. Prog.*, 10(1):147–175, 1976.
- A. Meenakshi and C. Rajian. On a product of positive semidefinite matrices. *Lin. Alg. Appl.*, 295(1-3):3–6, 1999.
- A. Mutapcic and S. Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Opt. Meth. Soft.*, 24(3):381–406, 2009.
- S. Naldi. Solving rank-constrained semidefinite programs in exact arithmetic. *J. Symb. Comput.*, 85:206–223, 2018.
- A. Nemirovski, C. Roos, and T. Terlaky. On maximization of quadratic form over intersection of ellipsoids with common center. *Math. Prog.*, 86(3):463–473, 1999.

- Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Y. Nesterov. Smoothing technique and its applications in semidefinite optimization. *Mathematical Programming*, 110(2):245–259, 2007.
- L. T. Nguyen, J. Kim, and B. Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.
- M. L. Overton and R. S. Womersley. On the sum of the largest eigenvalues of a symmetric matrix. *SIAM J. Matr. Anal. Appl.*, 13(1):41–45, 1992.
- J. Peng and Y. Wei. Approximating k-means-type clustering via semidefinite programming. *SIAM J. Opt.*, 18(1):186–205, 2007.
- J. Peng and Y. Xia. A new theoretical framework for k-means-type clustering. In *Found. Adv. data mining*, pages 79–96. Springer, 2005.
- M. Pilanci, M. J. Wainwright, and L. El Ghaoui. Sparse learning via boolean relaxations. *Math. Prog.*, 151(1):63–87, 2015.
- S. Puntanen and G. P. H. Styan. Historical introduction: Issai Schur and the early development of the Schur complement. In F. Zhang, editor, *The Schur Complement and Its Applications*, pages 1–16. Springer US, Boston, MA, 2005.
- 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.
- B. Recht and C. Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. *Math. Prog. Computation*, 5(2):201–226, 2013.
- B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, 2010.
- J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. part i-iii. *J. Symb. Comp.*, 13(3):255–352, 1992.
- M. Schaefer. Realizability of graphs and linkages. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 461–482. Springer New York, New York, NY, 2013.
- H. D. Sherali and B. M. Fraticelli. Enhancing RLT relaxations via a new class of semidefinite cuts. *J. Glob. Opt.*, 22(1-4):233–261, 2002.
- M. Udell, C. Horn, R. Zadeh, and S. Boyd. Generalized low rank models. *Foundations and Trends® in Machine Learning*, 9(1):1–118, 2016.
- L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Math. Prog. Comp.*, 4(4):333–361, 2012.
- H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of semidefinite programming*. Springer US, 2012.
- C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Ann. Stat.*, 38(2):894–942, 2010.
- H. Zhang, Z. Lin, and C. Zhang. A counterexample for the validity of using nuclear norm as a convex surrogate of rank. In *Euro. Conf. Mach. Learn. Know. Disc. Data.*, pages 226–241. Springer, 2013.
- X. Zheng, X. Sun, and D. Li. Improving the performance of MIQP solvers for quadratic programs with cardinality and minimum threshold constraints: A semidefinite program approach. *INFORMS J. Comput.*, 26(4):690–703, 2014.

A. Omitted Proofs

In this section, we supply omitted proofs of the results stated in the manuscript, in the order in which the results were stated.

A.1. Proof of Theorem 1

Proof of Theorem 1 By (Blekherman et al. 2012, Theorem 2.49), a set of distances $d_{i,j}$ can be embedded in a Euclidean space of dimension k if and only if there exists some Gram matrix $\mathbf{G}$ such that $\mathbf{D} = \text{Diag}(\mathbf{G})\mathbf{e}^\top + \mathbf{e}\text{Diag}(\mathbf{G})^\top - 2\mathbf{G}$ where $\mathbf{G} \succeq \mathbf{0}$ and $\text{Rank}(\mathbf{G}) \leq k$. Therefore, Proposition 1's $\exists\mathbb{R}$ -hard problem is a special case of:

$$\min_{\mathbf{G} \in S_+^n} \text{Rank}(\mathbf{G}) \quad \text{s.t.} \quad \text{Diag}(\mathbf{G})\mathbf{e}^\top + \mathbf{e}\text{Diag}(\mathbf{G})^\top - 2\mathbf{G} = \mathbf{D},$$

where $k = 2$, $D_{i,j} = 1 \ \forall (i,j) \in E$, and we do not impose the (i,j) th equality otherwise. $\square$

A.2. Proof of Theorem 2

Proof of Theorem 2 To establish this result, it suffices to perform a reduction from the following feasibility system to a polynomially sized system of polynomial equalities and inequalities:

$$\exists \mathbf{X}: \quad \text{Rank}(\mathbf{X}) \leq k, \quad \langle \mathbf{A}_i, \mathbf{X} \rangle = b_i \quad \forall i \in [m], \quad \mathbf{X} \succeq \mathbf{0},$$

because testing the feasibility of the above system a polynomial number of times for $k \in [n]$ certainly solves Problem (9).

By Proposition 2, this system is feasible if and only if the following system is also feasible:

$$\exists \mathbf{X}, \mathbf{Y}: \quad \mathbf{Y}^2 = \mathbf{Y}, \quad \mathbf{X} = \mathbf{X}\mathbf{Y}, \quad \text{tr}(\mathbf{Y}) \leq k, \quad \langle \mathbf{A}_i, \mathbf{X} \rangle = b_i \quad \forall i \in [m], \quad \mathbf{X} \succeq \mathbf{0},$$

The result then follows by observing that semidefinite constraints are indeed semialgebraic constraints, by the Tarski-Seidenberg theorem (see Blekherman et al. 2012, Chapter A.4.4), and therefore this system is equivalent to a polynomially sized system of polynomial equalities and inequalities. $\square$

A.3. Complexity of Low-Rank Integer Optimization and Hilbert's 10th Problem

We now demonstrate that imposing the constraint $\mathbf{X} \in \mathbb{Z}^{n \times n}$ changes Problem (1)'s complexity status, by making it undecidable. Formally, we have:

THEOREM 6. *The following problem is undecidable, even when $k = 1$ and its objective is binary:*

$$\min_{\mathbf{X} \in \mathbb{Z}^{n \times n}} \langle \mathbf{C}, \mathbf{X} \rangle \quad \text{s.t.} \quad \langle \mathbf{A}_i, \mathbf{X} \rangle = b_i \quad \forall i \in [m], \quad \text{Rank}(\mathbf{X}) \leq k, \quad \mathbf{X} \succeq \mathbf{0}.$$

REMARK 8. Theorem 6's reduction does not hold in the presence of a regularizer. Indeed, imposing either a spectral or Frobenius norm regularizer ensures the boundedness of the problem's level sets, which allows the problem to be solved in finite time via branch-and-bound. Nonetheless, Theorem 6 shows that imposing an integrality constraint makes Problem (1) *much* harder.

Proof of Theorem 6 We perform a reduction from integer optimization with quadratic constraints, which is undecidable when the objective is binary (Jeroslow 1973), by reduction from Hilbert's 10th problem. Recall that integer optimization with quadratic constraints is definitionally:

$$\min_{\mathbf{x} \in \mathbb{Z}^n} \langle \mathbf{c}, \mathbf{x} \rangle \quad \text{s.t.} \quad \langle \mathbf{x}, \mathbf{Q}_i \mathbf{x} \rangle + \langle \mathbf{a}_i, \mathbf{x} \rangle \leq b_i \quad \forall i \in [m]. \quad (47)$$

We now show that this problem is equivalent to:

$$\min_{\mathbf{x} \in \mathbb{Z}^n, \mathbf{X} \in \mathbb{Z}^{n \times n}} \langle \mathbf{c}, \mathbf{x} \rangle \quad \text{s.t.} \quad \langle \mathbf{Q}_i, \mathbf{X} \rangle + \langle \mathbf{a}_i, \mathbf{x} \rangle \leq b_i \quad \forall i \in [m], \quad \text{Rank} \begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{X} \end{pmatrix} \leq 1, \quad \begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{X} \end{pmatrix} \succeq \mathbf{0}.$$

To establish the result, it suffices to show that $\mathbf{X} = \mathbf{x}\mathbf{x}^\top$ in any feasible solution to the second problem. We now show this, by appealing to the Guttman rank identity (see Lemma C.2). We remind the reader that an equivalent form of the Guttman rank formula is the identity:

$$\text{Rank} \begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{X} \end{pmatrix} = \text{Rank}(\mathbf{X}) + \text{Rank}(\mathbf{X} - \mathbf{x}\mathbf{x}^\top),$$

Since $\mathbf{X} \succeq \mathbf{0}$ and the left hand side of the above expression is either 0 or 1, this formula implies that either $\mathbf{X} = \mathbf{x}\mathbf{x}^\top$ or $\mathbf{X} = \mathbf{0}$. However, the latter case can only hold if $\mathbf{x} = \mathbf{0}$, since $\mathbf{X} \succeq \mathbf{x}\mathbf{x}^\top$ by Schur complements. Therefore, $\mathbf{X} = \mathbf{x}\mathbf{x}^\top$ and the result holds. $\square$

A.4. Proof of Lemma 4

Proof of Lemma 4 Let us fix $\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)$. Then, we have that:

$$\begin{aligned} \max_{\alpha} h(\alpha) - \frac{\gamma}{2} \langle \alpha, \mathbf{Y} \alpha \rangle &= \max_{\alpha, \beta} h(\alpha) - \frac{\gamma}{2} \langle \beta, \mathbf{Y} \beta \rangle \quad \text{s.t.} \quad \beta = \alpha, \\ &= \max_{\alpha, \beta} \min_{\mathbf{X}} h(\alpha) - \frac{\gamma}{2} \langle \beta, \mathbf{Y} \beta \rangle - \langle \mathbf{X}, \beta - \alpha \rangle, \\ &= \min_{\mathbf{X}} \max_{\alpha} \underbrace{[h(\alpha) + \langle \mathbf{X}, \alpha \rangle]}_{(-h)^*(\mathbf{X}) = g(\mathbf{X})} + \max_{\beta} \left[\frac{-\gamma}{2} \langle \mathbf{Y} \beta, \beta \rangle - \langle \mathbf{X}, \beta \rangle \right]. \end{aligned}$$

Finally, the optimality condition with respect to β is $\mathbf{Y} \beta = \frac{-1}{\gamma} \mathbf{X}$, which implies the later term is

$$\begin{aligned} \max_{\mathbf{W}} \left[\frac{1}{2\gamma} \langle \mathbf{X}, \mathbf{Y}^\dagger \mathbf{X} \rangle - \frac{1}{2} \langle \mathbf{X}, (\mathbb{I} - \mathbf{Y}^\dagger \mathbf{Y}) \mathbf{W} \rangle \right] &= \max_{\mathbf{W}} \left[\frac{1}{2\gamma} \langle \mathbf{X}, \mathbf{Y}^\dagger \mathbf{X} \rangle - \frac{1}{2} \langle \mathbf{W}, (\mathbb{I} - \mathbf{Y}^\dagger \mathbf{Y}) \mathbf{X} \rangle \right] \\ &= \begin{cases} \frac{1}{2\gamma} \langle \mathbf{X}, \mathbf{Y}^\dagger \mathbf{X} \rangle & \text{if } \mathbf{Y} \in \text{Span}(\mathbf{X}), \\ +\infty & \text{otherwise.} \end{cases} \end{aligned}$$

We therefore conclude that the later term is equal to $\frac{1}{2\gamma} \langle \mathbf{X}, \mathbf{Y}^\dagger \mathbf{X} \rangle$ whenever the constraint $\mathbf{Y}^\dagger \mathbf{Y} \mathbf{X} = \mathbf{X}$ holds. By the generalized Schur complement lemma C.1, this expression is equivalent to introducing a new matrix $\boldsymbol{\theta}$, imposing the term $\frac{1}{2\gamma} \text{tr}(\boldsymbol{\theta})$ and requiring that $\begin{pmatrix} \boldsymbol{\theta} & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{Y} \end{pmatrix} \succeq \mathbf{0}$. $\square$

A.5. Proof of Lemma 5

Proof of Lemma 5 Let us fix $\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)$. Then, we have that:

$$\begin{aligned} \max_{\alpha \in S^n, \mathbf{W}_+, \mathbf{W}_- \succeq \mathbf{0}} h(\alpha) - M \langle \mathbf{Y}, \mathbf{W}_+ - \mathbf{W}_- \rangle \quad \text{s.t.} \quad \alpha = \mathbf{W}_+ - \mathbf{W}_- \\ &= \max_{\alpha \in S^n, \mathbf{W}_+, \mathbf{W}_- \succeq \mathbf{0}} \min_{\mathbf{X} \in S^n} h(\alpha) - M \langle \mathbf{Y}, \mathbf{W}_+ - \mathbf{W}_- \rangle + \langle \mathbf{X}, \alpha - \mathbf{W}_+ + \mathbf{W}_- \rangle \\ &= \min_{\mathbf{X} \in S^n} \max_{\alpha \in S^n, \mathbf{W}_+, \mathbf{W}_- \succeq \mathbf{0}} h(\alpha) - M \langle \mathbf{Y}, \mathbf{W}_+ - \mathbf{W}_- \rangle + \langle \mathbf{X}, \alpha - \mathbf{W}_+ + \mathbf{W}_- \rangle \\ &= \min_{\mathbf{X} \in S^n} \max_{\alpha \in S^n} \underbrace{[h(\alpha) + \langle \mathbf{X}, \alpha \rangle]}_{(-h)^*(\mathbf{X}) = g(\mathbf{X})} + \max_{\mathbf{W}_+, \mathbf{W}_- \succeq \mathbf{0}} [-M \langle \mathbf{Y}, \mathbf{W}_+ - \mathbf{W}_- \rangle + \langle \mathbf{X}, -\mathbf{W}_+ + \mathbf{W}_- \rangle]. \end{aligned}$$

Finally, the optimality conditions with respect to $\mathbf{W}_+, \mathbf{W}_-$ imply that $-M\mathbf{Y} \preceq \mathbf{X} \preceq M\mathbf{Y}$. $\square$

A.6. Proof of Lemmas 7 and 8

Proof of Lemma 7 In Problem (29), it is not too hard to see that for any $\mathbf{X}$ an optimal choice of $\mathbf{Y}$ is $\mathbf{Y} = \frac{1}{M}\mathbf{X}_+ + \frac{1}{M}\mathbf{X}_-$, where $\mathbf{X}_+, \mathbf{X}_-$ are orthogonal positive semidefinite matrices such that $\mathbf{X} = \mathbf{X}_+ - \mathbf{X}_-$. Therefore, the result follows by observing that $\text{tr}(\mathbf{X}_+ + \mathbf{X}_-) = \|\mathbf{X}\|_*$. $\square$

Proof of Lemma 8 In Problem (31), for any feasible $\mathbf{X}$ we have $\|\mathbf{X}\|_\sigma \leq M$. Under this constraint, it follows that for any $\mathbf{X}$ an optimal choice of $\mathbf{Y}, \mathbf{Y}'$ is $\mathbf{Y} = \mathbf{U}\Sigma\mathbf{U}^\top$, $\mathbf{Y}' = \mathbf{V}\Sigma\mathbf{V}^\top$, where $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$ is an SVD of $\mathbf{X}$. The result follows as $\text{tr}(\mathbf{Y}) = \text{tr}(\Sigma) = \|\mathbf{X}\|_*$. $\square$

A.7. Proof of Lemma 6

Proof of Lemma 6 Observe that, by the Generalized Schur Complement Lemma (see, e.g., lemma C.1), an optimal choice of $\boldsymbol{\theta}$ in Problem (26) is $\boldsymbol{\theta} = \mathbf{X}\mathbf{Y}^\dagger\mathbf{X}^\top$. Therefore, we can eliminate $\boldsymbol{\theta}$ from Problem (26), to obtain the equivalent objective:

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)} \min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \lambda \cdot \text{tr}(\mathbf{Y}) + g(\mathbf{X}) + \frac{1}{2\gamma} \langle \mathbf{X}\mathbf{X}^\top, \mathbf{Y}^\dagger \rangle.$$

Moreover, by the rank-nullity theorem (see, e.g., Horn and Johnson 1985, Chapter 0.2.3), we can split the columns of $\mathbf{Y}$ into columns in the span of the columns of $\mathbf{X}$ and columns orthogonal to the columns of $\mathbf{X}$. Since the columns orthogonal to the columns of $\mathbf{X}$ do not affect the objective value, it follows that we can write $\mathbf{Y}^\dagger = \sum_{i=1}^n \frac{1}{\theta_i} \mathbf{u}_i \mathbf{u}_i^\top$ without loss of optimality, where $\mathbf{X}\mathbf{X}^\top = \mathbf{U}\Sigma\mathbf{U}^\top$ is an SVD of $\mathbf{X}\mathbf{X}^\top$, and $0 \leq \theta_i \leq 1$ for each θ_i , because $\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n)$. Problem (26) then becomes:

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}, \boldsymbol{\theta} \in \mathbb{R}^n: \mathbf{0} \leq \boldsymbol{\theta} \leq \mathbf{e}} g(\mathbf{X}) + \sum_{i=1}^n \left(\lambda \theta_i + \frac{\sigma_i(\mathbf{X})^2}{2\gamma \theta_i} \right).$$

The result then follows because, for any $\lambda > 0$, (c.f. Pilanci et al. 2015, Equation (30))

$$\min_{0 \leq \theta \leq 1} \left[\lambda \theta + \frac{t^2}{\theta} \right] = \begin{cases} 2\sqrt{\lambda}|t|, & \text{if } |t| \leq \sqrt{\lambda}, \\ t^2 + \lambda, & \text{otherwise.} \end{cases} \quad \square$$

A.8. Proof of Theorem 4

Proof of Theorem 4 We only detail the proof of ϵ -optimality; the proof of ϵ -feasibility is almost identical (see Mutapcic and Boyd 2009).

Suppose that at some iteration $k > 1$, Algorithm 1 has not converged. Then,

$$\theta_k - f(\mathbf{Y}_k) < -\epsilon, \quad \text{and} \quad \theta_k \geq f(\mathbf{Y}_i) + \langle \mathbf{H}_i, \mathbf{Y}_k - \mathbf{Y}_i \rangle \quad \forall i < k.$$

But $\theta_k \leq f(\mathbf{Y}_i)$, since θ_k and $f(\mathbf{Y}_i)$ are respectively valid lower and upper bounds on the optimal objective. Therefore, $\langle \mathbf{H}_i, \mathbf{Y}_k - \mathbf{Y}_i \rangle \geq 0$. Putting the two inequalities together then implies that

$$f(\mathbf{Y}_k) - f(\mathbf{Y}_i) > \epsilon + \langle \mathbf{H}_i, \mathbf{Y}_k - \mathbf{Y}_i \rangle \geq \epsilon, \text{ or equivalently } \epsilon < f(\mathbf{Y}_k) - f(\mathbf{Y}_i) \leq L\|\mathbf{Y}_i - \mathbf{Y}_k\|_F,$$

where the second inequality holds by Lipschitz continuity. Rearranging this inequality implies that $\|\mathbf{Y}_i - \mathbf{Y}_k\|_F > \frac{\epsilon}{L}$, i.e., Algorithm 1 never visits any point within a ball of radius $\frac{\epsilon}{L}$ (with respect to the

Frobenius norm) twice. Moreover, by iteration k , Algorithm 1 visits k points within non-overlapping balls with combined volume

$$k \frac{\pi^{\frac{n^2}{2}}}{\Gamma(\frac{n^2}{2} + 1)} \left(\frac{\epsilon}{L}\right)^{n^2},$$

and these balls are centered at feasible points, i.e., contained within a ball of radius $K + \frac{\epsilon}{L}$ with volume

$$\frac{\pi^{\frac{n^2}{2}}}{\Gamma(\frac{n^2}{2} + 1)} \left(K + \frac{\epsilon}{L}\right)^{n^2}.$$

That is, if Algorithm 1 has not converged at iteration k , we have: $k < \left(\frac{LK}{\epsilon} + 1\right)^{n^2}$, which implies that we converge to an ϵ -optimal solution within $k \leq \left(\frac{LK}{\epsilon} + 1\right)^{n^2}$ iterations. $\square$

A.9. Proof of Lemma 9

Proof of Lemma 9 The equality $\boldsymbol{\theta}^* = \mathbf{X}_t^\top (\mathbf{Y}^*)^\dagger \mathbf{X}_t$ is immediate from the Schur complement lemma C.1. Therefore, we focus on deriving an optimal $\mathbf{Y}$ for a fixed $\mathbf{X}_t$.

From the second equality in the Schur complement lemma C.1, we must have $\mathbf{X} = \mathbf{Y}\mathbf{Y}^\dagger \mathbf{X}$ for feasibility. Therefore, $\text{span}(\mathbf{X}) \subseteq \text{span}(\mathbf{Y})$. Moreover, columns of $\mathbf{Y}$ which are in $\text{null}(\mathbf{X})$ do not contribute to the optimal objective and can therefore be omitted without loss of optimality. Therefore, $\mathbf{Y} = \sum_{i=1}^n \rho_i \mathbf{u}_i \mathbf{u}_i^\top$ for some $\boldsymbol{\rho}$, where $\mathbf{X}_t = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ is an SVD of $\mathbf{X}_t$. The result follows from observing that $\mathbf{0} \leq \boldsymbol{\rho} \leq \mathbf{e}$ and $\mathbf{e}^\top \boldsymbol{\rho} \leq k$, since $\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)$. $\square$

A.10. Proof of Lemma 10

Proof of Lemma 10 As Assumption 2 holds, we can exchange the minimization and maximization operators in Problem (40). Therefore, Problem (40) has the same optimal objective as:

$$\max_{\boldsymbol{\alpha}} h(\boldsymbol{\alpha}) - \max_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k)} \frac{\gamma}{2} \sum_{i=1}^n \sum_{j=1}^n Y_{i,j} \langle \boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j \rangle. \quad (48)$$

Therefore, to establish the result, it suffices to show that we obtain Problem (41) after taking the dual of Problem (48)'s inner problem. This is indeed the case, because $\text{Conv}(\mathcal{Y}_n^k)$ is a convex compact set with non-empty relative interior, and therefore strong duality holds between the following two problems:

$$\begin{aligned} \max_{\mathbf{Y} \succeq \mathbf{0}} \quad & \frac{\gamma}{2} \sum_{i=1}^n \sum_{j=1}^n Y_{i,j} \langle \boldsymbol{\alpha}_i, \boldsymbol{\alpha}_j \rangle & \text{s.t.} \quad & \mathbf{Y} \preceq \mathbb{I}, [\mathbf{U}] \langle \mathbb{I}, \mathbf{Y} \rangle \leq k, [t], \\ \min_{\mathbf{U} \succeq \mathbf{0}, t \geq 0} \quad & \text{tr}(\mathbf{U}) + kt & \text{s.t.} \quad & \mathbf{U} + \mathbb{I}t \succeq \frac{\gamma}{2} \boldsymbol{\alpha} \boldsymbol{\alpha}^\top. \quad \square \end{aligned}$$

A.11. Proof of Theorem 5

Proof of Theorem 5 To establish the result, we establish the first half of the following inequalities:

$$f(\mathbf{Y}_{\text{rounded}}) - f(\mathbf{Y}^*) \leq \frac{\gamma}{2} L^2 \max_{\boldsymbol{\beta} \geq \mathbf{0}: \|\boldsymbol{\beta}\|_\infty \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i \leq \frac{\gamma}{2} L^2 |\mathcal{R}| \max_{\boldsymbol{\beta} \geq \mathbf{0}: \|\boldsymbol{\beta}\|_1 \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i,$$

under the Frobenius penalty and

$$f(\mathbf{Y}_{\text{rounded}}) - f(\mathbf{Y}^*) \leq ML \max_{\boldsymbol{\beta} \geq \mathbf{0}: \|\boldsymbol{\beta}\|_\infty \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i \leq ML |\mathcal{R}| \max_{\boldsymbol{\beta} \geq \mathbf{0}: \|\boldsymbol{\beta}\|_1 \leq 1} \sum_{i \in \mathcal{R}} (\Lambda_{i,i}^* - \Lambda_{i,i}^{\text{rounded}}) \beta_i,$$

for the spectral penalty—the second half of both inequalities follows readily from the fact that $\|\beta\|_1 \leq |\mathcal{R}||\beta|_\infty \leq |\mathcal{R}|$ which allows us to replace $\|\beta\|_\infty \leq 1$ with $\|\beta\|_1 \leq |\mathcal{R}|$ and move $|\mathcal{R}|$ outside the bound. Indeed, after establishing these inequalities, the result follows by observing that $\mathbf{Y}_{greedy}$ minimizes the right-hand-side of (42)-(43) over the class of projection matrices $\mathbf{Y}_{rounded}$.

Under a Frobenius penalty, by Lipschitz continuity, we have

$$f(\mathbf{Y}_{rounded}) - f(\mathbf{Y}^*) \leq \frac{\gamma}{2} \langle \alpha^*(\mathbf{Y}) \alpha^*(\mathbf{Y})^\top, \mathbf{U}(\mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded}) \mathbf{U}^\top \rangle = \frac{\gamma}{2} \langle \mathbf{U}^\top \alpha^*(\mathbf{Y}) \alpha^*(\mathbf{Y})^\top \mathbf{U}, \mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded} \rangle.$$

Moreover, since $\mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded}$ is a diagonal matrix we need only include the diagonal terms in the inner product. Therefore, since

$$(\mathbf{U}^\top \alpha^*(\mathbf{Y}) \alpha^*(\mathbf{Y})^\top \mathbf{U})_{i,i} = \langle \alpha^*(\mathbf{Y})^\top \alpha^*(\mathbf{Y}), \mathbf{U}_i \mathbf{U}_i^\top \rangle \leq \lambda_{\max}(\alpha^*(\mathbf{Y})^\top \alpha^*(\mathbf{Y})) \leq L^2,$$

where the second-to-last inequality holds because $\|\mathbf{U}_i\|_2 = 1$, the bound on $f(\mathbf{Y}_{rounded}) - f(\mathbf{Y}^*)$ holds.

Alternatively, under spectral norm regularization, by Lipschitz continuity we have

$$\begin{aligned} f(\mathbf{Y}_{rounded}) - f(\mathbf{Y}^*) &\leq M \langle \mathbf{V}_{11}^*(\mathbf{Y}) + \mathbf{V}_{22}^*(\mathbf{Y}), \mathbf{U}(\mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded}) \mathbf{U}^\top \rangle \\ &= M \langle \mathbf{U}^\top (\mathbf{V}_{11}^*(\mathbf{Y}) + \mathbf{V}_{22}^*(\mathbf{Y})) \mathbf{U}, \mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded} \rangle. \end{aligned}$$

Moreover, $\mathbf{\Lambda}^* - \mathbf{\Lambda}_{rounded}$ is a diagonal matrix and therefore

$$(\mathbf{U}^\top (\mathbf{V}_{11}^*(\mathbf{Y}) + \mathbf{V}_{22}^*(\mathbf{Y})) \mathbf{U})_{i,i} = \langle \mathbf{U}_i \mathbf{U}_i^\top, \mathbf{V}_{11}^*(\mathbf{Y}) + \mathbf{V}_{22}^*(\mathbf{Y}) \rangle \leq \lambda_{\max}(\alpha^*(\mathbf{Y})) \leq L,$$

where the last inequality follows since $\mathbf{V}_{11}, \mathbf{V}_{22}$ are orthogonal at optimality, meaning $\mathbf{V}_{11} + \mathbf{V}_{22}$'s leading eigenvalue equals α^* 's leading singular value. Therefore, the bound on $f(\mathbf{Y}_{rounded}) - f(\mathbf{Y}^*)$ holds. $\square$

B. Derivations for the conjugate of the regularizer

In this section, we derive the conjugates of the penalties stated in Table 2, in order to complete our proof of Lemma 2. We first derive our results for rectangular matrices under the formulation $\mathbf{X} = \mathbf{Y} \mathbf{X} \mathbf{Y}'$ for appropriate projection matrices $\mathbf{Y}, \mathbf{Y}'$, before specializing our results by setting $\mathbf{Y}' = \mathbb{I}$. For completeness, we first prove that this reformulation is indeed a valid reformulation of a rank constraint.

PROPOSITION 3. *For any $\mathbf{X} \in \mathbb{R}^{n \times m}$, $\text{Rank}(\mathbf{X}) \leq k \iff \exists \mathbf{Y} \in \mathcal{Y}_n, \mathbf{Y}' \in \mathcal{Y}_m$ s.t. $\text{tr}(\mathbf{Y}), \text{tr}(\mathbf{Y}') \leq k$ and $\mathbf{X} = \mathbf{Y} \mathbf{X} \mathbf{Y}'$, where $\mathcal{Y}_n := \{\mathbf{P} \in S^n : \mathbf{P}^2 = \mathbf{P}\}$ is the set of projection matrices.*

Proof of Proposition 3 We prove the two implications successively.

($\implies$) Let $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ be a singular value decomposition of $\mathbf{X}$. Since $\text{Rank}(\mathbf{X}) \leq k$, we can let $\mathbf{U} \in \mathbb{R}^{n \times k}, \mathbf{\Sigma} \in \mathbb{R}^{k \times k}, \mathbf{V} \in \mathbb{R}^{k \times m}$ without loss of generality. Define $\mathbf{Y} = \mathbf{U}(\mathbf{U}^\top \mathbf{U})^{-1} \mathbf{U}$ and $\mathbf{Y}' = \mathbf{V}(\mathbf{V}^\top \mathbf{V})^{-1} \mathbf{V}$. By construction, $\mathbf{Y} \mathbf{X} \mathbf{Y}' = \mathbf{X}$. In addition, $\text{tr}(\mathbf{Y}) = \text{rank}(\mathbf{Y}) = \text{rank}(\mathbf{U}) \leq k$, and $\text{tr}(\mathbf{Y}') = \text{rank}(\mathbf{Y}') = \text{rank}(\mathbf{V}) \leq k$.

($\impliedby$) Since $\mathbf{X} = \mathbf{Y} \mathbf{X} \mathbf{Y}'$, $\text{rank}(\mathbf{X}) \leq \min(\text{rank}(\mathbf{Y}), \text{rank}(\mathbf{Y}')) \leq \min(\text{tr}(\mathbf{Y}), \text{tr}(\mathbf{Y}')) \leq k$. $\square$

B.1. Rectangular Case

In this section, we derive a dual reformulation for the conjugate of the regularization term in (9). More precisely, for all regularizers of interest $\Omega(\cdot)$, we show that for any matrix α of similar dimension as $\mathbf{X}$ and any projection matrices $\mathbf{Y} \in \mathcal{Y}_n^k$ and $\mathbf{Y}' \in \mathcal{Y}_m^k$,

$$\min_{\mathbf{X}} [\Omega(\mathbf{Y}\mathbf{X}\mathbf{Y}') + \langle \alpha, \mathbf{Y}\mathbf{X}\mathbf{Y}' \rangle] = \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\Omega^*(\alpha, \mathbf{Y}, \mathbf{Y}', \mathbf{V}_{11}, \mathbf{V}_{22}),$$

where $\Omega^*(\cdot)$ is notably linear in $\mathbf{Y}, \mathbf{Y}'$.

Lemma B.1 *Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be a rectangular matrix, $\mathbf{Y} \in \mathcal{Y}_n$, $\mathbf{Y}' \in \mathcal{Y}_m$ be projection matrices and $\gamma > 0$ be a positive scalar. Then*

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m}} \frac{1}{2\gamma} \|\mathbf{Y}\mathbf{X}\mathbf{Y}'\|_F^2 + \langle \mathbf{A}, \mathbf{Y}\mathbf{X}\mathbf{Y}' \rangle = -\frac{\gamma}{2} \text{tr}(\mathbf{Y}\mathbf{A}\mathbf{Y}'\mathbf{A}^\top).$$

Proof of Lemma B.1 Any solution to the minimization problem satisfies the first-order condition $\frac{1}{\gamma} \mathbf{Y}\mathbf{X}\mathbf{Y}' + \mathbf{Y}\mathbf{A}\mathbf{Y}' = 0$. Hence, $\mathbf{X}^* = -\gamma \mathbf{Y}\mathbf{A}\mathbf{Y}'$ is optimal and the objective is $-\frac{\gamma}{2} \text{tr}(\mathbf{Y}\mathbf{A}\mathbf{Y}'\mathbf{A}^\top)$. $\square$

Lemma B.2 *Let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be a rectangular matrix, $\mathbf{Y} \in \mathcal{Y}_n$, $\mathbf{Y}' \in \mathcal{Y}_m$ be projection matrices and $M > 0$ be a positive scalar. Then*

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{R}^{n \times m} : \|\mathbf{Y}\mathbf{X}\mathbf{Y}'\|_\sigma \leq M} \langle \mathbf{A}, \mathbf{Y}\mathbf{X}\mathbf{Y}' \rangle &= \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\frac{M}{2} \langle \mathbf{V}_{11}, \mathbf{Y} \rangle - \frac{M}{2} \langle \mathbf{V}_{22}, \mathbf{Y}' \rangle \\ \text{s.t.} \quad &\begin{pmatrix} \mathbf{V}_{11} & \mathbf{A} \\ \mathbf{A}^\top & \mathbf{V}_{22} \end{pmatrix} \succeq \mathbf{0}, \end{aligned}$$

Proof of Lemma B.2 Since the trace is invariant by cyclic permutation and the matrices $\mathbf{Y}$ and $\mathbf{Y}'$ are symmetric,

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times m} : \|\mathbf{Y}\mathbf{X}\mathbf{Y}'\|_\sigma \leq M} \langle \mathbf{A}, \mathbf{Y}\mathbf{X}\mathbf{Y}' \rangle = \min_{\mathbf{X} \in \mathbb{R}^{n \times m} : \|\mathbf{Y}\mathbf{X}\mathbf{Y}'\|_\sigma \leq M} \langle \mathbf{Y}\mathbf{A}\mathbf{Y}', \mathbf{X} \rangle.$$

The spectral norm penalty is equivalent to

$$\begin{aligned} \|\mathbf{Y}\mathbf{X}\mathbf{Y}'\|_\sigma \leq M &\iff \mathbf{Y}\mathbf{X}\mathbf{Y}'\mathbf{Y}'\mathbf{X}^\top\mathbf{Y} \preceq M^2\mathbf{I}_n \\ &\iff (\mathbf{Y}\mathbf{X}\mathbf{Y}')\mathbf{Y}'(\mathbf{Y}'\mathbf{X}^\top\mathbf{Y}) \preceq M^2\mathbf{I}_n \\ &\iff (\mathbf{Y}\mathbf{X}\mathbf{Y}')\mathbf{Y}'(\mathbf{Y}'\mathbf{X}^\top\mathbf{Y}) \preceq M^2\mathbf{Y}, \end{aligned}$$

where the last inequality follows from the fact that the matrix on the left-hand side is equal to zero over $\text{Im}(\mathbf{Y})^\top$. By Schur complements, the final semidefinite inequality is equivalent to

$$\begin{pmatrix} M\mathbf{Y} & \mathbf{Y}\mathbf{X}\mathbf{Y}' \\ \mathbf{Y}'\mathbf{X}^\top\mathbf{Y} & M\mathbf{Y}' \end{pmatrix} \succeq \mathbf{0} \quad \left[\begin{pmatrix} \mathbf{W}_{11} & \mathbf{W}_{12} \\ \mathbf{W}_{12}^\top & \mathbf{W}_{22} \end{pmatrix} \right].$$

We associate a matrix of dual variables in square brackets. Therefore, the dual problem is

$$\begin{aligned} \max_{\mathbf{W}_{11}, \mathbf{W}_{12}, \mathbf{W}_{22}} & -M\langle \mathbf{W}_{11}, \mathbf{Y} \rangle - M\langle \mathbf{W}_{22}, \mathbf{Y}' \rangle \\ \text{s.t.} \quad & \mathbf{Y}\mathbf{A}\mathbf{Y}' = 2\mathbf{Y}\mathbf{W}_{12}\mathbf{Y}', \quad \begin{pmatrix} \mathbf{W}_{11} & \mathbf{W}_{12} \\ \mathbf{W}_{12}^\top & \mathbf{W}_{22} \end{pmatrix} \succeq \mathbf{0}, \end{aligned}$$

$$\text{or equivalently,} \quad \max_{\mathbf{V}_{11}, \mathbf{V}_{22}} -\frac{M}{2} \langle \mathbf{V}_{11}, \mathbf{Y} \rangle - \frac{M}{2} \langle \mathbf{V}_{22}, \mathbf{Y}' \rangle \quad \text{s.t.} \quad \begin{pmatrix} \mathbf{V}_{11} & \mathbf{A} \\ \mathbf{A}^\top & \mathbf{V}_{22} \end{pmatrix} \succeq \mathbf{0}. \quad \square$$

B.2. Square and Symmetric Case

When the matrix $\mathbf{X}$ is square and symmetric we can take $\mathbf{Y} = \mathbf{Y}'$ and apply the previous results. Alternatively, for the spectral norm penalty, we can further simplify

Lemma B.3 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a square symmetric matrix and $\mathbf{Y} \in \mathcal{Y}_n$ be a projection matrix. Then*

$$\min_{\mathbf{X} \in \mathbb{S}_n: \|\mathbf{Y}\mathbf{X}\mathbf{Y}\|_\sigma \leq M} \langle \mathbf{A}, \mathbf{Y}\mathbf{X}\mathbf{Y} \rangle = \max_{\mathbf{V}_+, \mathbf{V}_- \succeq 0} -M \langle \mathbf{V}_+ + \mathbf{V}_-, \mathbf{Y} \rangle \quad \text{s.t.} \quad \mathbf{A} = \mathbf{V}_+ - \mathbf{V}_-.$$

Proof of Lemma B.3 The constraint $\|\mathbf{Y}\mathbf{X}\mathbf{Y}\|_\sigma \leq M$ can be rewritten as $-M\mathbf{Y} \preceq \mathbf{Y}\mathbf{X}\mathbf{Y} \preceq M\mathbf{Y}$. By strong semidefinite duality (which holds as the minimization problem has non-empty interior; see (Wolkowicz et al. 2012, Chapter 4.1)):

$$\min_{\mathbf{X} \in \mathbb{S}_n: \|\mathbf{Y}\mathbf{X}\mathbf{Y}\|_\sigma \leq M} \langle \mathbf{A}, \mathbf{Y}\mathbf{X}\mathbf{Y} \rangle = \max_{\mathbf{W}_+, \mathbf{W}_- \succeq 0} -M \langle \mathbf{Y}, \mathbf{W}_+ + \mathbf{W}_- \rangle \quad \text{s.t.} \quad \mathbf{Y}\mathbf{A}\mathbf{Y} = \mathbf{Y}\mathbf{W}_+\mathbf{Y} - \mathbf{Y}\mathbf{W}_-\mathbf{Y}.$$

The decision variables in the maximization problem decompose the symmetric matrix $\mathbf{Y}\mathbf{A}\mathbf{Y}$ into a positive and negative definite parts, $\mathbf{W}_+$ and $\mathbf{W}_-$ respectively. Without loss of optimality we can restrict our attention to $\mathbf{W}_+ = \mathbf{Y}\mathbf{V}_+\mathbf{Y}$ and $\mathbf{W}_- = \mathbf{Y}\mathbf{V}_-\mathbf{Y}$ where $\mathbf{A} = \mathbf{V}_+ - \mathbf{V}_-$. $\square$

C. A Collection of Useful Matrix Identities

In this work, we have repeatedly used some technical matrix identities. In order to keep this paper self contained, we now state these identities formally.

The following result generalizes the well-known Schur complement lemma to the case where neither on-diagonal block matrix is positive definite (see Boyd et al. 1994, Equation 2.41)

Lemma C.1 (Generalized Schur Complement Lemma) *Let $\mathbf{A}, \mathbf{B}, \mathbf{C}$ be components of*

$$\mathbf{X} := \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{pmatrix}$$

of appropriate dimension. Then, $\mathbf{X}$ is positive semidefinite if and only if the following conditions hold:

- (i) $\mathbf{A} \succeq \mathbf{0}$,
- (ii) $(\mathbb{I} - \mathbf{A}\mathbf{A}^\dagger)\mathbf{B} = \mathbf{0}$,
- (iii) $\mathbf{C} \succeq \mathbf{B}^\top \mathbf{A}^\dagger \mathbf{B}$.

The following result characterizes the rank of a block submatrix in terms of the rank of the entire matrix (see, e.g., Puntanen and Styan 2005, Section 0.9):

Lemma C.2 (Guttman Rank Identity) *Let $\mathbf{A}, \mathbf{B}, \mathbf{C}$ be components of the matrix*

$$\mathbf{X} := \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{pmatrix}$$

of appropriate dimension. Then, we have the identity: $\text{Rank}(\mathbf{X}) = \text{Rank}(\mathbf{A}) + \text{Rank}(\mathbf{C} - \mathbf{B}\mathbf{C}^\dagger \mathbf{B}^\top)$.

In general, a product of positive semidefinite matrices may not be positive semidefinite, indeed, it may not even be symmetric. However, the following result demonstrates that a *symmetric* product of PSD matrices is indeed PSD (see Meenakshi and Rajian 1999, for a proof):

Lemma C.3 (A Symmetric Product of PSD Matrices is PSD) *Let $\mathbf{X}, \mathbf{Y} \in S_+^n$ be matrices of appropriate size, and let $\mathbf{Z} := \mathbf{X}\mathbf{Y}$. Suppose that $\mathbf{Z} = \mathbf{Z}^\top$ is a symmetric matrix. Then, $\mathbf{Z} \succeq \mathbf{0}$.*

The following result demonstrates that evaluating the nuclear norm of a matrix via semidefinite optimization yields a singular value decomposition (see Recht et al. 2010, Proposition 2.1, for a proof)

Lemma C.4 (Nuclear norm minimization and the SVD of a matrix) *An optimal solution to*

$$\min_{\mathbf{W}_1 \in \mathbb{R}^{n \times n}, \mathbf{W}_2 \in \mathbb{R}^{m \times m}} \text{tr}(\mathbf{W}_1) + \text{tr}(\mathbf{W}_2) \quad \text{s.t.} \quad \begin{pmatrix} \mathbf{W}_1 & \mathbf{X} \\ \mathbf{X}^\top & \mathbf{W}_2 \end{pmatrix} \succeq \mathbf{0},$$

is attained by $\mathbf{W}_1 = \mathbf{U}\Sigma\mathbf{U}^\top$ and $\mathbf{W}_2 = \mathbf{V}\Sigma\mathbf{V}^\top$, where $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^\top$ is a singular value decomposition.

D. Additional Results

D.1. Pseudocode for the In-out Method

Our main points of difference from Fischetti et al. (2017)’s implementation are twofold. First, we optimize the outer problem over $\text{Conv}(\mathcal{Y}_n^k)$, rather than the Boolean polytope. Second, we recognize that the purpose of the method is to warm-start Algorithm 1’s lower bound, rather than to solve Problem (21). In this spirit, we accelerate the **in-out** method by first solving Problem (21) in one shot using an interior point method and second using the optimal solution $\mathbf{Y}^*$ as a stabilization point. Note that a similar method was proposed for sparse portfolio selection problems and MICO by Bertsimas and Cory-Wright (2018).

Algorithm 3 The in-out method of Ben-Ameur and Neto (2007).

Require: Stabilization point $\mathbf{Y}^*$, $\epsilon \leftarrow 10^{-10}$, $\lambda \leftarrow 0.1$, $\delta \leftarrow 2\epsilon$, $t \leftarrow 1$

repeat

 Compute $\mathbf{Y}_0, \theta_0$ solution of

$$\min_{\mathbf{Y} \in \text{Conv}(\mathcal{Y}_n^k), \theta} \theta + \lambda \cdot \text{tr}(\mathbf{Y}) \quad \text{s.t.} \quad z_i \theta \geq h_i + \langle H_i, \mathbf{Y} - \mathbf{Y}_i \rangle \quad \forall i \in [t].$$

if $\mathbf{Y}_0$ has not improved for 5 consecutive iterations **then**

 Set $\lambda = 1$

if $\mathbf{Y}_0$ has not improved for 10 consecutive iterations **then**

 Set $\delta = 0$

end if

end if

 Set $\mathbf{Y}_{t+1} \leftarrow \lambda \mathbf{Y}_0 + (1 - \lambda) \mathbf{Y}^* + \delta \mathbb{I}$; project $\mathbf{Y}_{t+1}$ onto $\text{Conv}(\mathcal{Y}_n^k)$.

 Compute $f(\mathbf{Y}_{t+1}), H_{t+1}, z_{t+1}, d_{t+1}$.

until $f(\mathbf{Y}_0) - \theta_0 \leq \epsilon$

return $\mathbf{Y}_t$

D.2. Number of cuts generated vs. γ

Figure 8 Average number of cuts vs. γ for Algorithm 1 single-tree (left) and multi-tree (right) implementations over 20 synthetic matrix completion instances where $p \in \{0.2, 0.3\}$, $r = 1$ and $n \in \{10, 20\}$. The same random seeds were used to generate random matrices completed by single-tree and multi-tree.