A New Perspective on Low-Rank Optimization

Dimitris Bertsimas · Ryan Cory-Wright · Jean Pauphilet

Abstract A key question in many low-rank problems throughout optimization, machine learning, and statistics is to characterize the convex hulls of simple low-rank sets and judiciously apply these convex hulls to obtain strong yet computationally tractable convex relaxations. We invoke the matrix perspective function the matrix analog of the perspective function and characterize explicitly the convex hull of epigraphs of simple matrix convex functions under low-rank constraints. Further, we combine the matrix perspective function with orthogonal projection matrices–the matrix analog of binary variables which capture the row-space of a matrix–to develop a matrix perspective reformulation technique that reliably obtains strong relaxations for a variety of low-rank problems, including reduced rank regression, non-negative matrix factorization, and factor analysis. Moreover, we establish that these relaxations can be modeled via semidefinite constraints and thus optimized over tractably. The proposed approach parallels and generalizes the perspective reformulation technique in mixed-integer optimization and leads to new relaxations for a broad class of problems.

Keywords Low-rank matrix · Semidefinite optimization · Matrix perspective function · Perspective reformulation technique

Mathematics Subject Classification (2010) 90C22 · 90C25 · 90C26 · 15A03 · 26B25

1 Introduction

Over the past decade, a considerable amount of attention has been devoted to low-rank optimization, resulting in theoretically and practically efficient algorithms for problems as disparate as matrix completion, reduced rank regression, or computer vision. In spite of this progress, almost no equivalent progress has been made on developing strong lower bounds for low-rank problems. Accordingly, this paper proposes a procedure for obtaining novel and strong lower bounds.

D. Bertsimas
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139
ORCID: 0000-0002-1985-1003 E-mail: dbertsim@mit.edu

R. Cory-Wright
Imperial College Business School, Imperial College, London, UK
ORCID: 0000-0002-4485-0619 E-mail: r.cory-wright@ic.ac.uk

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

We consider the following low-rank optimization problem:

$$\min_{\boldsymbol{X} \in \mathcal{S}^n_+} \langle \boldsymbol{C}, \boldsymbol{X} \rangle + \Omega(\boldsymbol{X}) + \mu \cdot \mathrm{Rank}(\boldsymbol{X}) \text{ s.t. } \langle \boldsymbol{A}_i, \boldsymbol{X} \rangle = b_i \ \forall i \in [m],\ \boldsymbol{X} \in \mathcal{K},\ \mathrm{Rank}(\boldsymbol{X}) \leq k, \quad (1)$$

where $\boldsymbol{C}, \boldsymbol{A}_1, \ldots \boldsymbol{A}_m \in \mathcal{S}^n$ are $n \times n$ symmetric matrices, $b_1, \ldots b_m \in \mathbb{R}$ are scalars, $[n]$ denotes the set of running indices $\{1, ..., n\}$, $\mathcal{S}^n_+$ denotes the $n \times n$ positive semidefinite cone, and $\mu \in \mathbb{R}_+, k \in \mathbb{N}$ are parameters which controls the complexity of $\boldsymbol{X}$ by respectively penalizing and constraining its rank. The set $\mathcal{K}$ is a properi.e., closed, convex, solid and pointedcone [c.f. 15, Section 2.4.1], and $\Omega(\boldsymbol{X}) = \mathrm{tr}(f(\boldsymbol{X}))$ for some matrix convex function f; see formal definitions and assumptions in Section 3.

For optimization problems with logical constraints, strong relaxations can be obtained by formulating them as mixed-integer optimization (MIO) problems and applying the so-called perspective reformulation technique [see 37, 42]. In this paper, we develop a matrix analog of the perspective reformulation technique to obtain strong yet computationally tractable relaxations of low-rank optimization problems of the form (1).

1.1 Motivating Example

In this section, we illustrate the implications of our results on a statistical learning example. To emphasize the analogy with the perspective reformulation technique in MIO, we first consider the best subset selection problem and review its perspective relaxations. We then consider a reduced-rank regression problem – the rank-analog of best subset selection – and provide new relaxations that naturally arise from our Matrix Perspective Reformulation Technique (MPRT).

Best Subset Selection: Given a data matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ and a response vector $\boldsymbol{y} \in \mathbb{R}^n$, the $\ell_0 - \ell_2$ regularized best subset selection problem is to solve [c.f. 65, 7, 6, 9, 77, 3]:

$$\min_{\boldsymbol{w} \in \mathbb{R}^p} \quad \frac{1}{2n}\|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 + \frac{1}{2\gamma}\|\boldsymbol{w}\|_2^2 + \mu\|\boldsymbol{w}\|_0, \quad (2)$$

where $\mu, \gamma > 0$ are parameters which control $\boldsymbol{w}$'s sparsity and sensitivity to noise respectively.

Early attempts at solving Problem (2) exactly relied upon weak implicit or big-M formulations of logical constraints which supply low-quality relaxations and therefore do not scale well [see 14, 44, for discussions]. However, very similar algorithms now solve these problems to certifiable optimality with millions of features. Perhaps the key ingredient in modernizing these (previously inefficient) algorithms was invoking the perspective reformulation techniquea technique for obtaining high-quality convex relaxations of non-convex setsfirst stated in Stubbs [72] PhD thesis [see also 73, 21] and popularized by Frangioni and Gentile [37], Aktürk et al. [1], Günlük and Linderoth [42] among others.

Relaxation via the Perspective Reformulation Technique: By applying the perspective reformulation technique [37, 1, 42] to the term $\mu\|\boldsymbol{w}\|_0 + \frac{1}{2\gamma}\|\boldsymbol{w}\|_2^2$, we obtain the following reformulation:

$$\min_{\boldsymbol{w}, \boldsymbol{\rho} \in \mathbb{R}^p, \boldsymbol{z} \in \{0,1\}^p} \quad \frac{1}{2n}\|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 + \frac{1}{2\gamma}\boldsymbol{e}^\top \boldsymbol{\rho} + \mu \cdot \boldsymbol{e}^\top \boldsymbol{z} \quad \text{s.t.} \quad z_i \rho_i \geq w_i^2 \quad \forall i \in [p], \quad (3)$$

where $\boldsymbol{e}$ denotes a vector of all ones of appropriate dimension.

Interestingly, this formulation can be represented using second-order cones [42, 65] and optimized over efficiently using projected subgradient descent [9]. Moreover, it reliably supplies near-exact relaxations for most practically relevant cases of best subset selection [65, 6]. In instances where it is not already tight, one can apply a refinement of the perspective reformulation technique to the term $\|\boldsymbol{y}-\boldsymbol{X}\boldsymbol{w}\|_2^2$ and thereby obtain the following (tighter yet more expensive) relaxation [26]:

$$\min_{\boldsymbol{w}\in\mathbb{R}^p,\boldsymbol{z}\in[0,1]^p,\boldsymbol{W}\in S_+^p} \quad \frac{1}{2n}\|\boldsymbol{y}\|_2^2-\frac{1}{n}\langle \boldsymbol{y},\boldsymbol{X}\boldsymbol{w}\rangle+\frac{1}{2}\langle \boldsymbol{W},\frac{1}{\gamma}\mathbb{I}+\frac{1}{n}\boldsymbol{X}^\top\boldsymbol{X}\rangle+\mu\boldsymbol{e}^\top\boldsymbol{z} \qquad (4)$$
$$\text{s.t.} \quad \boldsymbol{W}\succeq \boldsymbol{w}\boldsymbol{w}^\top,\ z_iW_{i,i}\geq w_i^2\ \forall i\in[p].$$

Recently, a class of even tighter relaxations were developed by Atamtürk and Gómez [3], Han et al. [43], Frangioni et al. [39]. As they were developed by considering multiple binary variables simultaneously and therefore do not, to our knowledge, generalize readily to the low-rank case (where we often have one low-rank matrix), we do not discuss (or generalize) them here.

Reduced Rank Regression: Given m observations of a response vector $\boldsymbol{Y}_j\in\mathbb{R}^n$ and a predictor $\boldsymbol{X}_j\in\mathbb{R}^p$, an important problem in high-dimensional statistics is to recover a low-complexity model which relates $\boldsymbol{X},\boldsymbol{Y}$. A popular choice for doing so is to assume that $\boldsymbol{X},\boldsymbol{Y}$ are related via $\boldsymbol{Y}=\boldsymbol{X}\boldsymbol{\beta}+\boldsymbol{E}$, where $\boldsymbol{\beta}\in\mathbb{R}^{p\times n}$ is a coefficient matrix which we assume to be low-rank, $\boldsymbol{E}$ is a matrix of noise and we require that the rank of $\boldsymbol{\beta}$ is small in order that the linear model is parsimonious [57]. Introducing Frobenius regularization gives rise to the problem:

$$\min_{\boldsymbol{\beta}\in\mathbb{R}^{p\times n}} \quad \frac{1}{2m}\|\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{\beta}\|_F^2+\frac{1}{2\gamma}\|\boldsymbol{\beta}\|_F^2+\mu\cdot\mathrm{Rank}(\boldsymbol{\beta}), \qquad (5)$$

where $\gamma,\mu>0$ control the robustness to noise and the complexity of the estimator respectively and we normalize the ordinary least squares loss by dividing by m, the number of observations.

Existing attempts at solving this problem generally involve replacing the low-rank term with a nuclear norm term [57], which succeeds under some strong assumptions on the problem data but not in general. Recently, we proposed a new framework to model rank constraints, using orthogonal projection matrices which satisfy $\boldsymbol{Y}^2=\boldsymbol{Y}$ instead of binary variables which satisfy $z^2=z$ [11]. By building on this work, in this paper we propose a generalization of the perspective function to matrix-valued functions with positive semidefinite arguments and develop a matrix analog of the perspective reformulation technique from MIO which uses projection matrices instead of binary variables.

Relaxations via the Matrix Perspective Reformulation Technique: By applying the matrix perspective reformulation technique (Theorem 1) to the term $\frac{1}{2\gamma}\|\boldsymbol{\beta}\|_F^2+\mu\cdot\mathrm{Rank}(\boldsymbol{\beta})$, we will prove that the following problem is a validand numerically high-qualityrelaxation of (5):

$$\min_{\boldsymbol{\beta}\in\mathbb{R}^{p\times n},\boldsymbol{W}\in\mathcal{S}_+^n,\boldsymbol{\theta}\in S_+^p} \quad \frac{1}{2m}\|\boldsymbol{Y}-\boldsymbol{X}\boldsymbol{\beta}\|_F^2+\frac{1}{2\gamma}\mathrm{tr}(\boldsymbol{\theta})+\mu\cdot\mathrm{tr}(\boldsymbol{W}) \quad \text{s.t.} \quad \boldsymbol{W}\preceq\mathbb{I}, \begin{pmatrix}\boldsymbol{\theta} & \boldsymbol{\beta}\\ \boldsymbol{\beta}^\top & \boldsymbol{W}\end{pmatrix}\succeq\boldsymbol{0}. \qquad (6)$$

The analogy between problems (2)-(5) and their relaxations (3)-(6) is striking. The goal of the present paper is to develop the corresponding theory to support and derive the relaxation (6). Interestingly, the main argument that led [26] to the improved relaxation (4) for (2) can be extended to reduced-rank regression. Combined with our MPRT, it leads to the relaxation:

$$\begin{aligned} \min_{\boldsymbol{\theta}\in\mathcal{S}^n_+,\boldsymbol{\beta}\in\mathbb{R}^{p\times n},\boldsymbol{B}\in\mathcal{S}^n_+,\boldsymbol{W}\in\mathcal{S}^n_+} \quad & \frac{1}{2m}\|\boldsymbol{Y}\|_F^2-\frac{1}{m}\langle \boldsymbol{Y},\boldsymbol{X}\boldsymbol{\beta}\rangle+\frac{1}{2}\langle \boldsymbol{B},\frac{1}{\gamma}\mathbb{I}+\frac{1}{m}\boldsymbol{X}^\top\boldsymbol{X}\rangle+\mu\cdot\mathrm{tr}(\boldsymbol{W}) \qquad (7)\\ \text{s.t.} \quad & \begin{pmatrix}\boldsymbol{B} & \boldsymbol{\beta}\\ \boldsymbol{\beta} & \boldsymbol{W}\end{pmatrix}\succeq \boldsymbol{0}, \boldsymbol{W}\preceq\mathbb{I}. \end{aligned}$$

It is not too hard to see that this is a valid semidefinite relaxation: if $\boldsymbol{W}$ is a rank-k projection matrix then, by the Schur complement lemma [see 16, Equation 2.41], $\boldsymbol{\beta}=\boldsymbol{\beta}\boldsymbol{W}$, and thus the rank of $\boldsymbol{\beta}$ is at most k. Moreover, if we let $\boldsymbol{B}=\boldsymbol{\beta}\boldsymbol{\beta}^\top$ in a solution, we recover a low-rank solution to the original problem[1]. Actually, as we show in Section 3.3, a similar technique can be applied to any instance of Problem (1), for which the applications beyond matrix regression are legion.

1.2 Literature Review

Three classes of approaches have been proposed for solving Problem (1): (a) heuristics, which prioritize computational efficiency and obtain typically high-quality solutions to low-rank problems efficiently but without optimality guarantees [see 59, for a review]; (b) relax-and-round approaches, which balance computational efficiency and accuracy concerns by relaxing the rank constraint and rounding a solution to the relaxation to obtain a provably near-optimal low-rank matrix [11, Section 1.2.2]; and (c) exact approaches, which prioritize accuracy over computational efficiency and solve Problem (1) exactly in exponential time [11, Section 1.2.1].

Of the three classes of approaches, heuristics currently dominate the literature, because their superior runtime and memory usage allows them to address larger-scale problems. However, recent advances in algorithmic theory and computational power have drastically improved the scalability of exact and approximate methods, to the point where they can now solve moderately sized problems which are relevant in practice [11]. Moreover, relaxations of strong exact formulations often give rise to very efficient heuristics (via tight relaxations of the exact formulation) which outperform existing heuristics. This suggests that heuristic approaches may not maintain their dominance going forward, and motivates the exploration of tight yet affordable relaxations of low-rank problems.

1.3 Contributions and Structure

The main contributions of this paper are twofold. First, we propose a general reformulation technique for obtaining high-quality relaxations of low-rank optimization problems: introducing an

[1] Observe that the constraints in Problem (4) are equivalent to the block matrix constraint $\begin{pmatrix}\mathrm{Diag}(\boldsymbol{z}) & \mathrm{Diag}(\boldsymbol{w})\\ \mathrm{Diag}(\boldsymbol{w}) & \boldsymbol{W}\end{pmatrix}\succeq\boldsymbol{0}$. This verifies that the reduced rank regression formulation is indeed a generalization of [26]'s formulation for sparse regression.

orthogonal projection matrix to model a low-rank constraint, and strengthening the formulation by taking the matrix perspective of an appropriate substructure of the problem. This technique can be viewed as a generalization of the perspective reformulation technique for obtaining strong relaxations of sparse or logically constrained problems [37, 42, 10, 43]. Second, by applying this technique, we obtain explicit characterizations of convex hulls of low-rank sets which frequently arise in low-rank problems. As the interplay between convex hulls of indicator sets and perspective functions has engineered algorithms which outperform state-of-the-art heuristics in sparse linear regression [6, 44] and sparse portfolio selection [79, 10], we hope that this work will empower similar developments for low-rank problems.

The rest of the paper is structured as follows: In Section 2 we supply some background on perspective functions and review their role in developing tight formulations of mixed-integer problems. In Section 3, we introduce the matrix perspective function and its properties, extend the function's definition to allow semidefinite in addition to positive definite arguments, and propose a matrix perspective reformulation technique (MPRT) which successfully obtains high-quality relaxations for low-rank problems which commonly arise in the literature.We also connect the matrix perspective function to the convex hulls of epigraphs of simple matrix convex functions under rank constraints. In Section 4, we illustrate the utility of this connection by deriving tighter relaxations of several low-rank problems than are currently available in the literature. Finally, in Section 5, we numerically verify the utility of our approach on reduced rank regression, D-optimal design and non-negative matrix factorization problems.

Notation: We let nonbold face characters such as b denote scalars, lowercase bold faced characters such as $\boldsymbol{x}$ denote vectors, uppercase bold faced characters such as $\boldsymbol{X}$ denote matrices, and calligraphic uppercase characters such as $\mathcal{Z}$ denote sets. We let $[n]$ denote the set of running indices $\{1, ..., n\}$ and $\mathbb{N}$ denote the set of positive integers. 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 let $\mathcal{S}^n$ denote the cone of $n \times n$ symmetric matrices, $\mathcal{S}^n_+$ denote the cone of $n \times n$ positive semidefinite matrices, $\mathcal{S}^n_+ \cap \mathbb{R}^{n \times n}_+$ denote the cone of $n \times n$ doubly non-negative matrices, and $\mathcal{C}^n_+ := \{\boldsymbol{U}\boldsymbol{U}^\top : \boldsymbol{U} \in \mathbb{R}^{n \times n}_+\}$ denote the cone of $n \times n$ completely positive matrices. Finally, we let $\boldsymbol{X}^\dagger$ denote the Moore-Penrose pseudoinverse of a matrix $\boldsymbol{X}$; see Horn and Johnson [47], Bhatia [13] for general theories of matrix operators. Less common matrix operators will be defined as they are needed.

2 Background on Perspective Functions

In this section, we review perspective functions and their interplay with tight formulations of logically constrained problems. This prepares the ground for and motivates our study of matrix perspective functions and their interplay with tight formulations of low-rank problems. Many of our subsequent results can be viewed as (nontrivial) generalizations of the results in this section, since a rank constraint is a cardinality constraint on the singular values.

2.1 Preliminaries

Consider a proper closed convex function $f : \mathcal{X} \to \mathbb{R}$, where $\mathcal{X}$ is a convex subset of $\mathbb{R}^n$. The perspective function of f is commonly defined for any $\boldsymbol{x} \in \mathbb{R}^n$ and any $t > 0$ as $(\boldsymbol{x}, t) \mapsto tf(\boldsymbol{x}/t)$. Its closure is defined by continuity for $t = 0$ and is equal to [c.f. 46, Proposition IV.2.2.2]:

$$g_f(\boldsymbol{x}, t) = \begin{cases} tf(\boldsymbol{x}/t) & \text{if } t > 0, \boldsymbol{x}/t \in \mathcal{X}, \\ 0 & \text{if } t = 0, \boldsymbol{x} = 0, \\ f_\infty(\boldsymbol{x}) & \text{if } t = 0, \boldsymbol{x} \neq 0, \\ +\infty & \text{otherwise,} \end{cases}$$

where f_∞ is the recession function of f, as originally stated in [68, p. 67] which is given by

$$f_\infty(\boldsymbol{x}) = \lim_{t \to 0} tf\left(\boldsymbol{x}_0 - \boldsymbol{x} + \frac{\boldsymbol{x}}{t}\right) = \lim_{t \to +\infty} \frac{f(\boldsymbol{x}_0 + t\boldsymbol{x}) - f(\boldsymbol{x}_0)}{t},$$

for any $\boldsymbol{x}_0$ in the domain of f. That is, $f_\infty(\boldsymbol{x})$ is the asymptotic slope of f in the direction of $\boldsymbol{x}$.

The perspective function was first investigated by Rockafellar [68], who made the important observation that f is convex in $\boldsymbol{x}$ if and only if g_f is convex in $(\boldsymbol{x}, t)$. Among other properties, we have that, for any $t > 0$, $(\boldsymbol{x}, t, s) \in \text{epi}(g_f)$ if and only if $(\boldsymbol{x}/t, s/t) \in \text{epi}(f)$ [46, Proposition IV.2.2.1]. We refer to the review by Combettes [23] for further properties of perspective functions.

Throughout this work, we refer to g_f as the *perspective function* of f –although it technically is the closure of the perspective. We also consider a family of convex functions f which satisfy:

Assumption 1 *The function $f : \mathcal{X} \to \mathbb{R}$ is proper, closed, and convex. $\mathbf{0} \in \mathcal{X}$ and for any $\boldsymbol{x} \neq \mathbf{0}$, $f_\infty(\boldsymbol{x}) = +\infty$.*

The condition $f_\infty(\boldsymbol{x}) = +\infty, \forall \boldsymbol{x} \neq \mathbf{0}$ is equivalent to $\lim_{\boldsymbol{x} \to \infty} f(\boldsymbol{x})/\|\boldsymbol{x}\| = +\infty$, and means that, asymptotically, f increases to infinity faster than any affine function. In particular, it is satisfied if the domain of f is bounded or if f is strictly convex. Under Assumption 1, the definition of the perspective function of f simplifies to

$$g_f(\boldsymbol{x}, t) = \begin{cases} tf(\boldsymbol{x}/t) & \text{if } t > 0, \\ 0 & \text{if } t = 0, \boldsymbol{x} = 0, \\ +\infty & \text{otherwise.} \end{cases} \tag{8}$$

2.2 The Perspective Reformulation Technique

A number of authors have observed that optimization problems over binary and continuous variables admit tight reformulations involving perspective functions of appropriate substructures of the problem, since Ceria and Soares [21], building upon the work of Rockafellar [68, Theorem 9.8], derived the convex hull of a disjunction of convex constraints. To motivate our study of the matrix perspective function in the sequel, we now demonstrate that a class of logically-constrained problems admit reformulations in terms of perspective functions. We remark that this development bears resemblance to other works on perspective reformulations including [10, 43, 39].

Consider a logically-constrained problem of the form

$$\min_{\boldsymbol{z}\in\mathcal{Z},\boldsymbol{x}\in\mathbb{R}^n} \quad \boldsymbol{c}^\top\boldsymbol{z}+f(\boldsymbol{x})+\Omega(\boldsymbol{x}) \quad \text{s.t.} \quad x_i=0 \text{ if } z_i=0 \quad \forall i\in[n], \tag{9}$$

where $\mathcal{Z}\subseteq\{0,1\}^n$, $\boldsymbol{c}\in\mathbb{R}^n$ is a cost vector, $f(\cdot)$ is a generic convex function which possibly models convex constraints $\boldsymbol{x}\in\mathcal{X}$ for a convex set $\mathcal{X}\subseteq\mathbb{R}^n$ implicitlyby requiring that $g(\boldsymbol{x})=+\infty$ if $\boldsymbol{x}\notin\mathcal{X}$, and $\Omega(\cdot)$ is a regularization function which satisfies the following assumption:

Assumption 2 (Separability) $\Omega(\boldsymbol{x})=\sum_{i\in[n]}\Omega_i(x_i)$, *where each* Ω_i *satisfies Assumption 1.*

Since z_i is binary, imposing the logical constraint "$x_i=0$ if $z_i=0$" plus the term $\Omega_i(x_i)$ in the objective is equivalent to $g_\Omega(x_i,z_i)+(1-z_i)\Omega_i(0)$ in the objective, where g_{Ω_i} is the perspective function of Ω_i, and thus Problem (9) is equivalent to:

$$\min_{\boldsymbol{z}\in\mathcal{Z},\boldsymbol{x}\in\mathbb{R}^n} \quad \boldsymbol{c}^\top\boldsymbol{z}+f(\boldsymbol{x})+\sum_{i=1}^n\left(g_{\Omega_i}(x_i,z_i)+(1-z_i)\Omega_i(0)\right). \tag{10}$$

Notably, while Problems (9)-(10) have the same feasible regions, (10) often has substantially stronger relaxations, as frequently noted in the perspective reformulation literature [37, 42, 36, 10].

For completeness, we provide a formal proof of equivalence between (9)-(10); note that a related (although dual, and weaker as it requires $\Omega(\mathbf{0})=\mathbf{0}$) result can be found in [10, Thm. 2.5]:

Lemma 1 *Suppose* (9) *attains a finite optimal value. Then,* (10) *attains the same value.*

Proof It suffices to establish that the following equality holds:

$$g_{\Omega_i}(x_i,z_i)+(1-z_i)\Omega_i(0)=\Omega_i(x_i)+\begin{cases}0 & \text{if } x_i=0 \text{ or } z_i=1,\\ +\infty & \text{otherwise.}\end{cases}$$

Indeed, this equality shows that any feasible solution to one problem is a feasible solution to the other with equal cost. We prove this by considering the cases where $z_i=0$, $z_i=1$ separately.

- Suppose $z_i=1$. Then, $g_{\Omega_i}(x_i,z_i)=z_i\Omega_i(x_i/z_i)=\Omega_i(x_i)$ and $x_i=z_i\cdot x_i$, so the result holds.
- Suppose $z_i=0$. If $x_i=0$ we have $g_{\Omega_i}(0,0)+\Omega_i(0)=\Omega_i(0)$, and moreover the right-hand-side of the equality is certainly $\Omega_i(0)$. Alternatively, if $x_i\neq 0$ then both sides equal $+\infty$. □

In Table 1, we present examples of penalties Ω for which Assumption 1 holds and the perspective reformulation technique is applicable. We remind the reader that the exponential cone is [c.f. 22]:

$$\mathcal{K}_{\exp}=\{\boldsymbol{x}\in\mathbb{R}^3: x_1\geq x_2\exp(x_2/x_3), x_2>0\}\cup\{(x_1,0,x_3)\in\mathbb{R}^3: x_1\geq 0, x_3\leq 0\},$$

while the power cone is defined for any $\alpha\in(0,1)$ as [c.f. 22]:

$$\mathcal{K}_{\text{pow}}^{\alpha}=\{\boldsymbol{x}\in\mathbb{R}^3: x_1^{\alpha}x_2^{1-\alpha}\geq|x_3|\}.$$

Table 1: Convex substructures which frequently arise in MIOs and their perspective reformulations. For conciseness, we give $g_\Omega(x,z)$ for $z>0$ only, i.e., the first case in (8), $g_\Omega(x,z)$ for $z=0$ being defined as in Equation (8).

Penalty	$\Omega(x)$	$g_\Omega(x,z)$ if $z>0$	Formulation
Big-M	$\begin{cases}0 & \text{if } \lvert x\rvert \le M, \\ +\infty & \text{otherwise}\end{cases}$	$\begin{cases}0 & \text{if } \lvert x\rvert \le Mz \\ +\infty & \text{otherwise}\end{cases}$	$\lvert x\rvert \le Mz$
Ridge	$\frac{1}{2\gamma}x^2$	$x^2/2\gamma z$	$\min\ \theta$ s.t. $\theta z \ge \frac{1}{2\gamma}x^2$
Ridge + Big-M	$\frac{1}{2\gamma}x^2,\ \text{if } \lvert x\rvert \le M$	$x^2/2\gamma z,\ \text{if } \lvert x\rvert \le Mz$	$\min\ \theta$ s.t. $\theta z \ge \frac{1}{2\gamma}x^2,\ \lvert x\rvert \le Mz$
Power	$\lvert x\rvert^p,\ p{>}1$	$\lvert x\rvert^p z^{1-p}$	$\min\ \theta$ s.t. $(\theta, z, x) \in \mathcal{K}_{\text{pow}}^{1/p}$
Log_ϵ + Big-M	$-\log(x+\epsilon),\ \text{if } 0 \le x \le M$	$-z\log(x/z+\epsilon),\ \text{if } x \le Mz$	$\min\ \theta$ s.t. $(x+z\epsilon, z, -\theta) \in \mathcal{K}_{\exp}$, $x \le Mz$
Entropy	$x\log x$	$x\log(x/z),\ \text{if } x>0$	$\min\ \theta$ s.t. $(z, x, -\theta) \in \mathcal{K}_{\exp}$, $x \le Mz$
Softplus+Big-M	$\log(1+\exp(x)),\ \text{if } \lvert x\rvert \le M$	$z\log(1+\exp(x/z)),\ \text{if } \lvert x\rvert \le Mz$	$\min\ \theta$ s.t. $z \ge u+v, \lvert x\rvert \le Mz$, $(u, z, -\theta) \in \mathcal{K}_{\exp}$, $(v, z, x-\theta) \in \mathcal{K}_{\exp}$

2.3 Perspective Cuts

Another computationally useful application of the perspective reformulation technique has been to derive a class of cutting-planes for MIOs with logical constraints [37]. To motivate our generalization of these cuts to low-rank problems, we now briefly summarize their main result.

Consider the following problem:

$$\begin{aligned} \min_{\boldsymbol{z}\in\mathcal{Z}} \min_{\boldsymbol{x}\in\mathbb{R}^n} \quad & \boldsymbol{c}^\top \boldsymbol{z} + f(\boldsymbol{x}) + \sum_{i=1}^n \Omega_i(x_i) \\ \text{s.t.} \quad & \boldsymbol{A}^i x_i \le b_i z_i \quad \forall i \in [n], \end{aligned} \tag{11}$$

where $\{x_i : \boldsymbol{A}^i x_i \le 0\} = \{0\}$, which implies the set of feasible $\boldsymbol{x}$ is bounded, $\Omega_i(x_i)$ is a closed convex function, we take $\Omega_i(0) = 0$ as in [37] for simplicity, and $f(\boldsymbol{x})$ is a convex function. Then, letting ρ_i model the epigraph of $\Omega_i(x_i) + c_i z_i$ and s_i be a subgradient of Ω_i at $\bar{x}_i$, i.e., $s_i \in \partial\Omega_i(\bar{x}_i)$, we have the following result [37, 42]:

Proposition 1 *The following cut*

$$\rho_i \geq (c_i + \Omega_i(\bar{x}_i))z_i + s_i(x_i - \bar{x}_i z_i) \tag{12}$$

is valid for the equivalent MINLO:

$$\begin{aligned} \min_{\boldsymbol{z} \in \mathcal{Z}} \min_{\boldsymbol{x}, \boldsymbol{\rho} \in \mathbb{R}^n} \quad & f(\boldsymbol{x}) + \sum_{i=1}^n \rho_i \\ s.t. \quad & \boldsymbol{A}^i x_i \leq b_i z_i \quad \forall i \in [n], \\ & \rho_i \geq \Omega_i(x_i) + c_i z_i \quad \forall i \in [n]. \end{aligned}$$

Remark 1 In the special case where $\Omega_i(x_i) = x_i^2$, the cut reduces to:

$$\rho_i \geq 2x_i\bar{x}_i - \bar{x}_i^2 z_i + c_i z_i \quad \forall \bar{x}_i. \tag{13}$$

The class of cutting planes defined in Proposition 1 are commonly referred to as perspective cuts, because they define a linear lower approximation of the perspective function of $\Omega_i(x_i)$, $g_{\Omega_i}(x_i, z_i)$. Consequently, Proposition 1 implies that a perspective reformulation of (11) is equivalent to adding all (infinitely many) perspective cuts (12). This may be helpful where the original problem is nonlinear, as a sequence of linear MIOs can be easier to solve than one nonlinear MIO [see 38, for a comparison].

3 The Matrix Perspective Function and Its Applications

In this section, we generalize the perspective function from vectors to matrices, and invoke the matrix perspective function to propose a new technique for generating strong yet efficient relaxations of a diverse family of low-rank problems, which we call the Matrix Perspective Reformulation Technique (MPRT). Selected background on matrix analysis [see 13, for a general theory] and semidefinite optimization [see 76, for a general theory] which we use throughout this section can be found in Appendix A.

3.1 A Matrix Perspective Function

To generalize the ideas from the previous section to low-rank constraints, we require a more expressive transform than the perspective transform, which introduces a single (scalar) additional degree of freedom and cannot control the eigenvalues of a matrix. Therefore, we invoke a generalization from quantum mechanicsthe matrix perspective function defined in [27, 28], building upon the work of [29]; see also [54, 55, 56, 24] for a related generalization of perspective functions to perspective functionals.

Definition 1 For a matrix-valued function $f : \mathcal{X} \to \mathcal{S}_+^n$ where $\mathcal{X} \subseteq \mathcal{S}^n$ is a convex set, the matrix perspective function of f, g_f, is defined as

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{Y}^{\frac{1}{2}} f\left(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}\right)\boldsymbol{Y}^{\frac{1}{2}} & \text{if } \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}} \in \mathcal{X}, \boldsymbol{Y} \succ \boldsymbol{0}, \\ \infty & \text{otherwise.} \end{cases}$$

Remark 2 If $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute and f is analytic, then Definition 1 simplifies into $\boldsymbol{Y} f\left(\boldsymbol{Y}^{-1}\boldsymbol{X}\right)$, which is the analog of the usual definition of the perspective function originally stated in [29]. Definition 1, however, generalizes this definition to the case where $\boldsymbol{X}$ and $\boldsymbol{Y}$ do not commute by ensuring that $\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}$ is nonetheless symmetric, in a manner reminiscent of the development of interior point methods [see, e.g., 2]. In particular, if $\boldsymbol{Y}$ is a projection matrix such that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$as occurs for the exact formulations of the low-rank problems we consider in this paperthen it is safe to assume that $\boldsymbol{X}, \boldsymbol{Y}$ commute. However, when $\boldsymbol{Y}$ is not a projection matrix, this cannot be assumed in general.

The matrix perspective function generalizes the definition of the perspective transformation to matrix-valued functions and satisfies analogous properties:

Proposition 2 *Let f be a matrix-valued function and g_f its matrix perspective function. Then:*

(a) f is matrix convex, i.e.,

$$tf(\boldsymbol{X}) + (1-t)f(\boldsymbol{W}) \succeq f(t\boldsymbol{X} + (1-t)\boldsymbol{W}) \quad \forall \boldsymbol{X}, \boldsymbol{W} \in \mathcal{S}^n,\ t \in [0,1], \tag{14}$$

if and only if g_f is matrix convex in $(\boldsymbol{X}, \boldsymbol{Y})$.

(b) g_f is a positive homogeneous function, i.e., for any $\mu>0$ we have

$$g_f(\mu\boldsymbol{X}, \mu\boldsymbol{Y}) = \mu g_f(\boldsymbol{X}, \boldsymbol{Y}). \tag{15}$$

(c) Let $\boldsymbol{Y} \succ \boldsymbol{0}$ be a positive definite matrix. Then, letting the epigraph of f be denoted by

$$\mathrm{epi}(f) := \{(\boldsymbol{X}, \boldsymbol{\theta}) : \boldsymbol{X} \in \mathrm{dom}(f), f(\boldsymbol{X}) \preceq \boldsymbol{\theta}\}, \tag{16}$$

we have $(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\theta}) \in \mathrm{epi}(g_f)$ if and only if $(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}, \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{\theta}\boldsymbol{Y}^{-\frac{1}{2}}) \in \mathrm{epi}(f)$.

Proof We prove the claims successively:

(a) This is precisely the main result of Ebadian et al. [27, Theorem 2.2].
(b) For $\mu > 0$, $g_f(\mu\boldsymbol{X}, \mu\boldsymbol{Y}) = \mu\boldsymbol{Y}^{\frac{1}{2}} f\left((\mu\boldsymbol{Y})^{-\frac{1}{2}}\mu\boldsymbol{X}(\mu\boldsymbol{Y})^{-\frac{1}{2}}\right)\boldsymbol{Y}^{\frac{1}{2}} = \mu g_f(\boldsymbol{X}, \boldsymbol{Y})$.
(c) By generalizing the main result in [15, Chapter 3.2.6], for any $\boldsymbol{Y} \succ \boldsymbol{0}$ we have that

$$\begin{aligned}
(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\theta}) \in \mathrm{epi}(g_f) \iff &\quad \boldsymbol{Y}^{\frac{1}{2}} f(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}})\boldsymbol{Y}^{\frac{1}{2}} \preceq \boldsymbol{\theta}, \\
\iff &\quad f(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}) \preceq \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{\theta}\boldsymbol{Y}^{-\frac{1}{2}}, \\
\iff &\quad (\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}, \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{\theta}\boldsymbol{Y}^{-\frac{1}{2}}) \in \mathrm{epi}(f). \quad \square
\end{aligned}$$

We now specialize our attention to matrix-valued functions defined by a scalar convex function, as suggested in the introduction.

3.2 Matrix Perspectives of Operator Functions

From any function $\omega : \mathbb{R} \to \mathbb{R}$, we can define its extension to the set of symmetric matrices, $f_\omega : \mathcal{S}^n \to \mathcal{S}^n$ as

$$f_\omega(\boldsymbol{X}) = \boldsymbol{U}\,\mathrm{Diag}(\omega(\lambda_1^x), \ldots, \omega(\lambda_n^x))\boldsymbol{U}^\top, \tag{17}$$

where $\boldsymbol{X} = \boldsymbol{U}\,\mathrm{Diag}(\lambda_1^x, \ldots, \lambda_n^x)\boldsymbol{U}^\top$ is an eigendecomposition of $\boldsymbol{X}$. Functions of this form are called *operator functions* [see 13, for a general theory]. In particular, one can show that $f_\omega(\boldsymbol{X})$ is well-defined (does not depend explicitly on the eigenbasis of $\boldsymbol{X}$, $\boldsymbol{U}$). Among other examples, taking $\omega(x) = \exp(x)$ (resp. $\log(x)$) provides a matrix generalization of the exponential (resp. logarithm) function; see Appendix A.1.

Central to our analysis is that we can explicitly characterize the closure of the matrix perspective of f_ω under some assumptions on ω, i.e., define by continuity $g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})$ for rank-deficient matrices $\boldsymbol{Y}$:

Proposition 3 *Consider a function $\omega : \mathbb{R} \to \mathbb{R}$ satisfying Assumption 1. Then, the closure of the matrix perspective of f_ω is, for any $\boldsymbol{X} \in \mathcal{S}^n$, $\boldsymbol{Y} \in \mathcal{S}^n_+$,*

$$g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{Y}^{\frac{1}{2}} f_\omega(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}})\boldsymbol{Y}^{\frac{1}{2}} & \textit{if } \mathrm{Span}(\boldsymbol{X}) \subseteq \mathrm{Span}(\boldsymbol{Y}), \boldsymbol{Y} \succeq \boldsymbol{0}, \\ \infty & \textit{otherwise,} \end{cases}$$

where $\boldsymbol{Y}^{-\frac{1}{2}}$ denotes the pseudo-inverse of the square root of $\boldsymbol{Y}$.

Remark 3 Note that in the expression of g_{f_ω} above, the matrix $\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}$ is unambiguously defined if and only if $\mathrm{Span}(\boldsymbol{X}) \subseteq \mathrm{Span}(\boldsymbol{Y})$ (otherwise, its value depends on how we define the pseudo-inverse of $\boldsymbol{Y}^{\frac{1}{2}}$ outside of its range). Accordingly, in the remainder of the paper, we omit the condition $\mathrm{Span}(\boldsymbol{X}) \subseteq \mathrm{Span}(\boldsymbol{Y})$ whenever the analytic expression for g_{f_ω} explicitly involves $\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}$.

The proof of Proposition 3 is deferred to Appendix B.1. In the appendix, we also present an immediate extension where additional constraints, $\boldsymbol{X} \in \mathcal{X}$, are imposed on the argument of f_ω. As in our prior work Bertsimas et al. [11], we reformulate the rank constraints in (1) by introducing a projection matrix $\boldsymbol{Y}$ to encode for the span of $\boldsymbol{X}$. Naturally, $\boldsymbol{Y}$ should be rank-deficient. Hence, Proposition 3 ensures that having $\mathrm{tr}(g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})) < \infty$ is a sufficient condition for $\boldsymbol{Y}$ to indeed control $\mathrm{Span}(\boldsymbol{X})$.

To gain intuition on how the matrix perspective function transforms $\boldsymbol{X}$ and $\boldsymbol{Y}$, we now provide an interesting connection between the matrix perspective of f_ω and the perspective of ω in the case where $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute.

Proposition 4 *Consider two matrices $\boldsymbol{X} \in \mathcal{S}^n, \boldsymbol{Y} \in \mathcal{S}^n_+$ that commute and such that* $\mathrm{Span}(\boldsymbol{X}) \subseteq \mathrm{Span}(\boldsymbol{Y})$. *Hence, there exists an orthogonal matrix $\boldsymbol{U}$ which jointly diagonalizes $\boldsymbol{X}$ and $\boldsymbol{Y}$. Let $\lambda_1^x, \ldots, \lambda_n^x$ and $\lambda_1^y, \ldots, \lambda_n^y$ denote the eigenvalues of $\boldsymbol{X}$ and $\boldsymbol{Y}$ respectively, ordered according to this basis $\boldsymbol{U}$. Consider an operator function f_ω with ω satisfying Assumption 1. Then, we have that:*

$$g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y}) = \boldsymbol{U}\,\mathrm{Diag}\left(g_\omega(\lambda_1^x, \lambda_1^y), \ldots, g_\omega(\lambda_n^x, \lambda_n^y)\right)\boldsymbol{U}^\top$$

Proof By simultaneously diagonalizing $\boldsymbol{X}$ and $\boldsymbol{Y}$, we get

$$\begin{aligned}
\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}} &= \boldsymbol{U}\mathrm{Diag}\left(\lambda_1^x/\lambda_1^y, \ldots, \lambda_n^x/\lambda_n^y\right)\boldsymbol{U}^\top,\\
f_\omega\left(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}\right) &= \boldsymbol{U}\mathrm{Diag}\left(\omega(\lambda_1^x/\lambda_1^y), \ldots, \omega(\lambda_n^x/\lambda_n^y)\right)\boldsymbol{U}^\top,\\
\boldsymbol{Y}^{\frac{1}{2}} f_\omega\left(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}\right)\boldsymbol{Y}^{\frac{1}{2}} &= \boldsymbol{U}\mathrm{Diag}\left(\lambda_1^y\omega(\lambda_1^x/\lambda_1^y), \ldots, \lambda_n^y\omega(\lambda_n^x/\lambda_n^y)\right)\boldsymbol{U}^\top. \quad \square
\end{aligned}$$

Note that if $\boldsymbol{Y}$ is a projection matrix such that $\mathrm{Span}(\boldsymbol{X}) \subseteq \mathrm{Span}(\boldsymbol{Y})$ then we necessarily have that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X} = \boldsymbol{X}\boldsymbol{Y}$ and the assumptions of Proposition 4 hold.

In the general case where $\boldsymbol{X}$ and $\boldsymbol{Y}$ do not commute, we cannot simultaneously diagonalize them and connect g_{f_ω} with g_ω. However, we can still project $\boldsymbol{Y}$ onto the space of matrices that commute with $\boldsymbol{X}$ and obtain the following result when g_{f_ω} is matrix convex (proof deferred to Appendix B.2):

Lemma 2 *Let $\boldsymbol{X} \in \mathcal{S}^n$ and $\boldsymbol{Y} \in \mathcal{S}^n_+$ be matrices, and denote the commutant of $\boldsymbol{X}$, $\mathcal{C}_{\boldsymbol{X}} := \{\boldsymbol{M} : \boldsymbol{M}\boldsymbol{X} = \boldsymbol{X}\boldsymbol{M}\}$, i.e., the set of matrices which commute with $\boldsymbol{X}$. For any matrix $\boldsymbol{M}$, denote $\boldsymbol{M}_{|\boldsymbol{X}}$ the orthogonal projection of $\boldsymbol{M}$ onto $\mathcal{C}_{\boldsymbol{X}}$. Then, since $\boldsymbol{M} \mapsto \boldsymbol{M}_{|\boldsymbol{X}}$ is a projection operator, we have that*

$$\boldsymbol{Y}_{|\boldsymbol{X}} \in \mathcal{S}^n_+, \quad \text{and} \quad \mathrm{tr}\left(\boldsymbol{Y}_{|\boldsymbol{X}}\right) = \mathrm{tr}\left(\boldsymbol{Y}\right).$$

Moreover, if $\boldsymbol{Y} \mapsto g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})$ is matrix convex, then we have

$$\mathrm{tr}\left[g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y}_{|\boldsymbol{X}})\right] \leq \mathrm{tr}\left[g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})\right].$$

3.3 The Matrix Perspective Reformulation Technique

Definition 1 and Proposition 3 supply the necessary language to lay out our Matrix Perspective Reformulation Technique (MPRT). Therefore, we now state the technique; details regarding its implementation will become clearer throughout the paper.

Let us revisit Problem (1), and assume that the term $\Omega(\boldsymbol{X})$ satisfies the following properties:

Assumption 3 $\Omega(\boldsymbol{X}) = \mathrm{tr}\left(f_\omega(\boldsymbol{X})\right)$*, where ω is a function satisfying Assumption 1 and whose associated operator function, f_ω, is matrix convex.*

Assumption 3 implies that the regularizer can be rewritten as operating on the eigenvalues of $\boldsymbol{X}$, $\lambda_i(\boldsymbol{X})$, directly: $\Omega(\boldsymbol{X}) = \sum_{i\in[n]} \omega(\lambda_i(\boldsymbol{X}))$. As we discuss in the next section, a broad class of functions satisfy this property. For ease of notation, we refer to f_ω as f in the remainder of the paper (and accordingly denote by g_f its matrix perspective function).

After letting an orthogonal projection matrix $\boldsymbol{Y}$ model the rank of $\boldsymbol{X}$as per [11]Problem (1) admits the equivalent mixed-projection reformulation:

$$\begin{aligned}
\min_{\boldsymbol{Y}\in\mathcal{Y}^k_n} \min_{\boldsymbol{X}\in\mathcal{S}^n} \quad & \langle \boldsymbol{C}, \boldsymbol{X}\rangle + \mu \cdot \mathrm{tr}(\boldsymbol{Y}) + \mathrm{tr}(f(\boldsymbol{X})) \\
\text{s.t.} \quad & \langle \boldsymbol{A}_i, \boldsymbol{X}\rangle = b_i \quad \forall i \in [m], \ \boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}, \ \boldsymbol{X} \in \mathcal{K},
\end{aligned} \qquad (18)$$

where $\boldsymbol{Y} \in \mathcal{Y}_n^k$ is the set of $n \times n$ orthogonal projection matrices with trace at most k:

$$\mathcal{Y}_n^k := \left\{\boldsymbol{Y} \in \mathcal{S}_+^n : \boldsymbol{Y}^2 = \boldsymbol{Y},\ \operatorname{tr}(\boldsymbol{Y}) \le k\right\}.$$

Note that for $k \in \mathbb{N}$, the convex hull of $\mathcal{Y}_n^k$ is given by $\operatorname{Conv}(\mathcal{Y}_n^k) = \{\boldsymbol{Y} \in \mathcal{S}_+^n : \boldsymbol{Y} \preceq \mathbb{I}, \operatorname{tr}(\boldsymbol{Y}) \le k\}$, which is a well-studied object in its own right [60, 61, 52, 62].

Since $\boldsymbol{Y}$ is an orthogonal projection matrix, imposing the nonlinear constraint $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$ and introducing the term $\Omega(\boldsymbol{X}) = \operatorname{tr}(f(\boldsymbol{X}))$ in the objective is equivalent to introducing the following term in the objective:

$$\operatorname{tr}(g_f(\boldsymbol{X}, \boldsymbol{Y})) + (n - \operatorname{tr}(\boldsymbol{Y}))\omega(0),$$

where g_f is the matrix perspective of f, and thus Problem (18) is equivalent to:

$$\begin{aligned} \min_{\boldsymbol{Y} \in \mathcal{Y}_n^k} \min_{\boldsymbol{X} \in \mathcal{S}^n} \quad & \langle \boldsymbol{C}, \boldsymbol{X}\rangle + \mu \cdot \operatorname{tr}(\boldsymbol{Y}) + \operatorname{tr}(g_f(\boldsymbol{X}, \boldsymbol{Y})) + (n - \operatorname{tr}(\boldsymbol{Y}))\omega(0) \\ \text{s.t.} \quad & \langle \boldsymbol{A}_i, \boldsymbol{X}\rangle = b_i \quad \forall i \in [m],\ \boldsymbol{X} \in \mathcal{K}, \end{aligned} \tag{19}$$

Let us formally state and verify the equivalence between Problems (18)-(19) via:

Theorem 1 *Problems* (18)-(19) *attain the same optimal objective value.*

Proof It suffices to show that for any feasible solution to (18) we can construct a feasible solution to (19) with an equal or lower cost, and vice versa:

– Let $(\boldsymbol{X}, \boldsymbol{Y})$ be a feasible solution to (18). Since $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X} \in \mathcal{S}^n$, $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute. Hence, by Proposition 4, we have (using the same notation as in Proposition 4):

$$\operatorname{tr}\left(g_f(\boldsymbol{X}, \boldsymbol{Y})\right) = \sum_{i \in [n]} g_\omega\left(\lambda_i^x, \lambda_i^y\right) = \sum_{i \in [n]} 1\{\lambda_i^y > 0\}\omega(\lambda_i^x),$$

where $1\{\lambda_i^y > 0\}$ is an indicator function which denotes whether the ith eigenvalue of $\boldsymbol{Y}$ (which is either 0 or 1) is strictly positive. Moreover, since $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$, $\lambda_i^y = 0 \implies \lambda_i^x = 0$ and

$$\begin{aligned} \operatorname{tr}\left(f(\boldsymbol{X})\right) &= \sum_{i \in [n]} \omega(\lambda_i^x) = \operatorname{tr}\left(g_f(\boldsymbol{X}, \boldsymbol{Y})\right) + \sum_{i \in [n]} 1\{\lambda_i^y = 0\}\omega(0) \\ &= \operatorname{tr}\left(g_f(\boldsymbol{X}, \boldsymbol{Y})\right) + (n - \operatorname{tr}(\boldsymbol{Y}))\omega(0). \end{aligned} \tag{20}$$

This establishes that $(\boldsymbol{X}, \boldsymbol{Y})$ is feasible in (19) with the same cost.

– Let $(\boldsymbol{X}, \boldsymbol{Y})$ be a feasible solution to (19). Then, it follows that $\boldsymbol{X} \in \operatorname{Span}(\boldsymbol{Y})$, which implies that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$ since $\boldsymbol{Y}$ is a projection matrix. Therefore, (20) holds, which establishes that $(\boldsymbol{X}, \boldsymbol{Y})$ is feasible in (18) with the same cost. □

Eventually, relaxing $\boldsymbol{Y} \in \mathcal{Y}_n^k$ in Problem (19) supplies as strongand sometimes significantly strongerrelaxations than by any other technique we are aware of, as we explore in Section 4.

Remark 4 Note that, based on the proof of Theorem 1, we could replace $g_f(\boldsymbol{X}, \boldsymbol{Y})$ in (19) by any function $\tilde{g}(\boldsymbol{X}, \boldsymbol{Y})$ such that $g_f(\boldsymbol{X}, \boldsymbol{Y}) = \tilde{g}(\boldsymbol{X}, \boldsymbol{Y})$ for $\boldsymbol{X}, \boldsymbol{Y}$ that commute, with no impact on the objective value. However, it might impact tractability if $\tilde{g}(\boldsymbol{X}, \boldsymbol{Y})$ is not convex in $(\boldsymbol{X}, \boldsymbol{Y})$.

Remark 5 Under Assumption 3, the regularization term $\Omega(\boldsymbol{X})$ penalizes all eigenvalues of $f_\omega(\boldsymbol{X})$ equally. The MPRT can be extended to a wider class of regularization functions that penalize the largest eigenvalues more heavily, at the price of (a significant amount of) additional notation. For brevity, we lay out this extension in Appendix C.

Theorem 1 only uses the fact that f is an operator function with ω satisfying Assumption 1, not the fact that f is matrix convex. In other words, (19) is always an equivalent reformulation of (18). An interesting question is to identify the set of necessary conditions for the objective of (19) to be convex in $(\boldsymbol{X}, \boldsymbol{Y})$–$f$ being matrix convex is clearly sufficient. The objective in (19) is convex only as long as $\operatorname{tr}(g_f)$ is. Interestingly, this is not equivalent to the convexity of $\operatorname{tr}(f)$. See Appendix B.3 for a counter-example. It is, however, an open question whether a weaker notion than matrix convexity could ensure the joint convexity of $\operatorname{tr}(g_f)$. It would also be interesting to investigate the benefits and the tractability of non-convex penalties (either by having f not matrix convex or ω non-convex), given the successes of non-convex penalty functions in sparse regression problems [78, 30].

3.4 Convex Hulls of Low-Rank Sets and the MPRT

We now show that, for a general class of low-rank sets, applying the MPRT is equivalent to taking the convex hull of the set. This is significant, because we are not aware of any general-purpose techniques for taking convex hulls of low-rank sets. Formally, we have the following result:

Theorem 2 *Consider an operator function $f = f_\omega$ satisfying Assumption 3. Let*

$$\mathcal{T} = \{\boldsymbol{X} \in \mathcal{S}^n : \operatorname{tr}(f(\boldsymbol{X})) + \mu \cdot \operatorname{Rank}(\boldsymbol{X}) \le t, \operatorname{Rank}(\boldsymbol{X}) \le k\} \tag{21}$$

be a set where $t \in \mathbb{R}, k \in \mathbb{N}$ are fixed. Then, an extended formulation of the convex hull of $\mathcal{T}$ is given by:

$$\mathcal{T}^c = \left\{(\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{S}^n \times \operatorname{Conv}(\mathcal{Y}_n^k) : \operatorname{tr}(g_f(\boldsymbol{X}, \boldsymbol{Y})) + \mu \cdot \operatorname{tr}(\boldsymbol{Y}) + (n - \operatorname{tr}(\boldsymbol{Y}))\omega(0) \le t\right\}. \tag{22}$$

Where $\operatorname{Conv}(\mathcal{Y}_n^k) = \{\boldsymbol{Y} \in \mathcal{S}_+^n : \boldsymbol{Y} \preceq \mathbb{I}, \operatorname{tr}(\boldsymbol{Y}) \le k\}$ *is the convex hull of trace-k projection matrices, and g_f is the matrix perspective function of f.*

Proof We prove the two directions sequentially:

- $\operatorname{Conv}(\mathcal{T}) \subseteq \mathcal{T}^c$: let $\boldsymbol{X} \in \mathcal{T}$. Then, since the rank of $\boldsymbol{X}$ is at most k, there exists some $\boldsymbol{Y} \in \mathcal{Y}_n^k$ such that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$ and $\operatorname{tr}(\boldsymbol{Y}) = \operatorname{Rank}(\boldsymbol{X})$. Moreover, by the same argument as in the proof of Theorem 1, it follows that (20) holds and $\operatorname{tr}(g_f(\boldsymbol{X}, \boldsymbol{Y})) + \mu \cdot \operatorname{tr}(\boldsymbol{Y}) + (n - \operatorname{tr}(\boldsymbol{Y}))\omega(0) \le t$, which confirms that $(\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{T}^c$. Since $\mathcal{T}^c$ is a convex set, we therefore have $\operatorname{Conv}(\mathcal{T}) \subseteq \mathcal{T}^c$.
- $\mathcal{T}^c \subseteq \operatorname{Conv}(\mathcal{T})$: let $(\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{T}^c$. Denote $\boldsymbol{Y}_{|\boldsymbol{X}}$ the projection of $\boldsymbol{Y}$ onto the set of matrices that commute with $\boldsymbol{X}$: $\{\boldsymbol{M} : \boldsymbol{X}\boldsymbol{M} = \boldsymbol{M}\boldsymbol{X}\}$. By Lemma 2, we have that $\boldsymbol{Y}_{|\boldsymbol{X}} \in \operatorname{Conv}(\mathcal{Y}_n^k)$, and $\operatorname{tr}\left(g_f(\boldsymbol{X}, \boldsymbol{Y}_{|\boldsymbol{X}})\right) \le \operatorname{tr}(g_f(\boldsymbol{X}, \boldsymbol{Y})) < \infty$ so $(\boldsymbol{X}, \boldsymbol{Y}_{|\boldsymbol{X}}) \in \mathcal{T}^c$ as well. Hence, without loss of generality, by renaming $\boldsymbol{Y} \leftarrow \boldsymbol{Y}_{|\boldsymbol{X}}$, we can assume that $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute. Then, it follows

from Proposition 4 that the vectors of eigenvalues of $\boldsymbol{X}$ and $\boldsymbol{Y}$ (ordered according to a shared eigenbasis $\boldsymbol{U}$), $(\boldsymbol{\lambda}(\boldsymbol{X}), \boldsymbol{\lambda}(\boldsymbol{Y}))$ belong to the set

$$\left\{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^n \times [0,1]^n : \sum_i y_i \leq k, \sum_{i=1}^n y_i \omega\left(\frac{x_i}{y_i}\right) + \mu \sum_i y_i + (n - \sum_i y_i)\omega(0) \leq t\right\},$$

which, via a straightforward generalization of [75, Theorem 3] to consider $\omega(0) \neq 0$ as in Section 2.2 of this paper, is the closure of the convex hull of

$$\mathcal{U}^c := \left\{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^n \times \{0,1\}^n : \sum_i y_i \leq k, \sum_{i=1}^n \omega(x_i) + \mu \sum_i y_i \leq t, x_i = 0 \text{ if } y_i = 0\ \forall i \in [n]\right\}.$$

Let us decompose $(\boldsymbol{\lambda}(\boldsymbol{X}), \boldsymbol{\lambda}(\boldsymbol{Y}))$ into $\boldsymbol{\lambda}(\boldsymbol{X}) = \sum_k \alpha_k \boldsymbol{x}^{(k)}$, $\lambda(\boldsymbol{Y}) = \sum_k \alpha_k \boldsymbol{y}^{(k)}$, with $\alpha_k \geq 0$, $\sum_k \alpha_k = 1$, and $(\boldsymbol{x}^{(k)}, \boldsymbol{y}^{(k)}) \in \mathcal{U}^c$. By definition,

$$\boldsymbol{T}^{(k)} := \boldsymbol{U}\mathrm{Diag}(\boldsymbol{x}^{(k)})\boldsymbol{U}^\top \in \mathcal{T}$$

and $\boldsymbol{X} = \sum_k \alpha_k \boldsymbol{T}^{(k)}$. Therefore, we have that $\boldsymbol{X} \in \mathrm{Conv}(\mathcal{T})$, as required. □

Remark 6 Since linear optimization problems over convex sets admit extremal optima, Theorem 2 demonstrates that *unconstrained* low-rank problems with spectral objectives can be recast as linear semidefinite problems, where the rank constraint is dropped without loss of optimality. This suggests that work on hidden convexity in low-rank optimization, i.e., deriving conditions under which low-rank linear optimization problems admit exact relaxations where the rank constraint is omitted [see, e.g., 62, 74, 12], could be extended to incorporate spectral functions.

3.5 Examples of the Matrix Perspective Function

Theorem 2 demonstrates that, for spectral functions under low-rank constraints, taking the matrix perspective is equivalent to taking the convex hull. To highlight the utility of Theorems 1-2, we therefore supply the perspective functions of some spectral regularization functions which frequently arise in the low-rank matrix literature, and summarize them in Table 2. We also discuss how these functions and their perspectives can be efficiently optimized over. Note that all functions introduced in this section are either matrix convex or the trace of a matrix convex function, and thus supply valid convex relaxations when used as regularizers for the MPRT.

Spectral constraint: Let $\omega(x) = 0$ if $|x| \leq M$, $+\infty$ otherwise. Then,

$$f(\boldsymbol{X}) = \begin{cases} \boldsymbol{0} & \text{if } \|\boldsymbol{X}\|_\sigma \leq M, \\ +\infty & \text{otherwise,} \end{cases}$$

for $\boldsymbol{X} \in \mathcal{S}^n$, where $\|\cdot\|_\sigma$ denotes the spectral norm, i.e., the largest eigenvalue in absolute magnitude of $\boldsymbol{X}$. Observe that the condition $\|\boldsymbol{X}\|_\sigma \leq M$ can be expressed via semidefinite constraints$-M\mathbb{I} \preceq \boldsymbol{X} \preceq M\mathbb{I}$. The perspective function g_f can then be expressed as

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{0} & \text{if } -M\boldsymbol{Y} \preceq \boldsymbol{X} \preceq M\boldsymbol{Y}, \\ +\infty & \text{otherwise.} \end{cases}$$

If $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute, $g_f(\boldsymbol{X}, \boldsymbol{Y})$ requires that $|\lambda_j(\boldsymbol{X})| \leq M\lambda_j(\boldsymbol{Y})\ \forall j \in [n]$–the spectral analog of a big-M constraint. This constraint can be modeled using two semidefinite cones, and thus handled by semidefinite solvers.

Convex quadratic: For $\omega(x) = x^2$, $f(\boldsymbol{X}) = \boldsymbol{X}^\top \boldsymbol{X}$. Then, the perspective function g_f is

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{X}^\top \boldsymbol{Y}^\dagger \boldsymbol{X} & \text{if } \boldsymbol{Y} \succeq \boldsymbol{0}, \\ +\infty & \text{otherwise.} \end{cases}$$

Observe that this function's epigraph is semidefinite-representable. Indeed, by the Schur complement lemma [16, Equation 2.41], minimizing the trace of $g_f(\boldsymbol{X}, \boldsymbol{Y})$ is equivalent to solving

$$\min_{\boldsymbol{\theta} \in \mathcal{S}^n, \boldsymbol{Y} \in \mathcal{S}^n, \boldsymbol{X} \in \mathcal{S}^n} \quad \mathrm{tr}(\boldsymbol{\theta}) \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{\theta} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{Y} \end{pmatrix} \succeq \boldsymbol{0}.$$

Interestingly, this perspective function allows us to rewrite the rank-k SVD problem

$$\min_{\boldsymbol{X} \in \mathbb{R}^{n \times m}} \quad \|\boldsymbol{X} - \boldsymbol{A}\|_F^2 : \ \mathrm{Rank}(\boldsymbol{X}) \leq k$$

as a linear optimization problem over the set of orthogonal projection matrices, which implies that the orthogonal projection constraint can be relaxed to its convex hull without loss of optimality (since some extremal solution will be optimal for the relaxation). This is significant, because while rank-k SVD is commonly thought of as a non-convex problem which "surprisingly" admits a closed-form solution, the MPRT shows that it actually admits an *exact* convex reformulation:

$$\min_{\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\theta}} \quad \frac{1}{2}\mathrm{tr}(\boldsymbol{\theta}) - \langle \boldsymbol{A}, \boldsymbol{X} \rangle + \frac{1}{2}\|\boldsymbol{A}\|_F^2 \text{ s.t. } \boldsymbol{Y} \preceq \mathbb{I}, \ \mathrm{tr}(\boldsymbol{Y}) \leq k, \begin{pmatrix} \boldsymbol{\theta} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{Y} \end{pmatrix} \succeq \boldsymbol{0}.$$

Note that, in the above formulation, we extended our results for symmetric matrices to rectangular matrices $\boldsymbol{X} \in \mathbb{R}^{n \times m}$ without justification. We rigorously derive this extension for $f(\boldsymbol{X}) = \boldsymbol{X}^\top \boldsymbol{X}$ in Appendix D and defer the study of the general case to future research.

Spectral plus convex quadratic: Let

$$f(\boldsymbol{X}) = \begin{cases} \boldsymbol{X}^\top \boldsymbol{X} & \text{if} \quad \|\boldsymbol{X}\|_\sigma \leq M, \\ +\infty & \text{otherwise,} \end{cases}$$

for $\boldsymbol{X} \in \mathcal{S}^n$. Then, the perspective function g_f is

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{X}^\top \boldsymbol{Y}^\dagger \boldsymbol{X} & \text{if} \quad -M\boldsymbol{Y} \preceq \boldsymbol{X} \preceq M\boldsymbol{Y}, \\ +\infty & \text{otherwise.} \end{cases}$$

This can be interpreted as the spectral analog of combining a big-M and a ridge penalty.

Convex quadratic over completely positive cone: Consider the following optimization problem

$$\min_{\boldsymbol{X}\in\mathcal{S}^n} \boldsymbol{X}^\top \boldsymbol{X} \text{ s.t. } \boldsymbol{X}\in\mathcal{C}^n_+,$$

where $\mathcal{C}^n_+ = \{\boldsymbol{X} : \boldsymbol{X} = \boldsymbol{U}\boldsymbol{U}^\top, \boldsymbol{U}\in\mathbb{R}^{n\times n}_+\} \subseteq \mathcal{S}^n_+$ denotes the completely positive cone. Then, by denoting $f(\boldsymbol{X}) = \boldsymbol{X}^\top\boldsymbol{X}$ and g_f its perspective function we obtain a valid relaxation by minimizing $\mathrm{tr}(g_f)$, which, by the Schur complement lemma [see 16, Equation 2.41], can be reformulated as

$$\min_{\boldsymbol{\theta}\in\mathcal{S}^n, \boldsymbol{Y}\in\mathcal{S}^n, \boldsymbol{X}\in\mathcal{S}^n} \quad \mathrm{tr}(\boldsymbol{\theta}) \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{\theta} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{Y} \end{pmatrix} \in \mathcal{S}^{2n}_+, \boldsymbol{X}\in\mathcal{C}^n_+.$$

Unfortunately, this formulation cannot be tractably optimized over, since separating over the completely positive cone is NP-hard. However, by relaxing the completely positive cone to the doubly non-negative cone$\mathcal{S}^n_+ \cap \mathbb{R}^{n\times n}_+$we obtain a tractable and near-exact relaxation. Indeed, as we shall see in our numerical experiments, combining this relaxation with a state-of-the-art heuristic supplies certifiably near-optimal solutions in both theory and practice.

Note that we could have obtained an alternative relaxation by instead considering the perspective of

$$f(\boldsymbol{X}) = \begin{cases} \boldsymbol{X}^\top\boldsymbol{X} & \text{if} \quad \boldsymbol{X}\in\mathcal{C}^n_+, \\ +\infty & \text{otherwise.} \end{cases}$$

Remark 7 One can obtain a nearly identical formulation over the copositive cone [c.f. 17].

Power: Let[2] $f(\boldsymbol{X}) = \boldsymbol{X}^\alpha$ for $\alpha\in[0,1]$ and $\boldsymbol{X}\in\mathcal{S}^n_+$. The matrix perspective function is[3]

$$g_f(\boldsymbol{X},\boldsymbol{Y}) = \begin{cases} \boldsymbol{Y}^{1/2}\left(\boldsymbol{Y}^{-1/2}\boldsymbol{X}\boldsymbol{Y}^{-1/2}\right)^\alpha \boldsymbol{Y}^{1/2} & \text{if } \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}} \in\mathcal{S}^n_+, \boldsymbol{Y}\succeq \boldsymbol{0}, \\ +\infty & \text{otherwise.} \end{cases}$$

The expression above can be simplified into $\boldsymbol{Y}^{\frac{1-\alpha}{2}}\boldsymbol{X}^\alpha\boldsymbol{Y}^{\frac{1-\alpha}{2}}$ when $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute and, per Remark 4, the former expression can be used equivalently for optimization purposes, even when $\boldsymbol{X}$ and $\boldsymbol{Y}$ do not commute.

Remark 8 (Matrix Power Cone) This function's epigraph, the matrix power cone, i.e.,

$$\mathcal{K}^{\mathrm{pow},\alpha}_{\mathrm{mat}} = \{(\boldsymbol{X}_1,\boldsymbol{X}_2,\boldsymbol{X}_3)\in\mathcal{S}^n_+\times\mathcal{S}^n_+\times\mathcal{S}^n : \boldsymbol{X}_2^{\frac{1-\alpha}{2}}\boldsymbol{X}_1^\alpha\boldsymbol{X}_2^{\frac{1-\alpha}{2}} \succeq \boldsymbol{X}_{3,+}+\boldsymbol{X}_{3,-}\}$$

is a closed convex cone which is semidefinite representable for any rational α [32]. Consequently, it is a tractable object which successfully models the matrix power function (and its perspective) and we shall make repeated use of it when we apply the MPRT to several important low-rank problems in Section 3.5.

[2] Note that $f(\boldsymbol{X})$ and its perspective are concave functions; hence we model their hypographs, not epigraphs.

[3] We only consider the PSD case for notational convenience. However, the symmetric case follows in much the same manner, after splitting $\boldsymbol{X} = \boldsymbol{X}_+ - \boldsymbol{X}_- : \boldsymbol{X}_+, \boldsymbol{X}_- \succeq \boldsymbol{0}, \langle \boldsymbol{X}_+, \boldsymbol{X}_-\rangle = 0$ and replacing $\boldsymbol{X}$ with $\boldsymbol{X}_+ + \boldsymbol{X}_-$.

Logarithm: Let $f(\boldsymbol{X}) = -\log(\boldsymbol{X})$ be the matrix logarithm function. We have that

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} -\boldsymbol{Y}^{\frac{1}{2}} \log\left(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}\right) \boldsymbol{Y}^{\frac{1}{2}} & \text{if } \boldsymbol{X}, \boldsymbol{Y} \succ \boldsymbol{0}, \\ +\infty & \text{otherwise.} \end{cases}$$

Observe that when $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute, $g_f(\boldsymbol{X}, \boldsymbol{Y})$ can be rewritten as $\boldsymbol{Y}(\log(\boldsymbol{Y}) - \log(\boldsymbol{X}))$, whose trace is the (Umegaki) quantum relative entropy function [see 33, for a general theory]. We remark that the domain of $\log(\boldsymbol{X})$ requires that $\boldsymbol{X}$ is full-rank, which at a first glance makes the use of this function problematic for low-rank optimization. Accordingly, we consider the $\epsilon-$logarithm function, i.e., $\log_\epsilon(\boldsymbol{X}) = \log(\boldsymbol{X} + \epsilon \mathbb{I})$ for $\epsilon > 0$, as advocated by Fazel et al. [35] in a different context. Background on matrix exponential and logarithm functions can be found in Appendix A.

Observe that $\mathrm{tr}(\log(\boldsymbol{X})) = \log\det(\boldsymbol{X})$ while $\mathrm{tr}(g_f) = \mathrm{tr}(\boldsymbol{X}(\log(\boldsymbol{X}) - \log(\boldsymbol{Y}))$. Thus, the matrix logarithm and its trace verify the concavity of the logdet functionwhich has numerous applications in low-rank problems [35] and interior point methods [67] among otherswhile the perspective of the matrix logarithm provides an elementary proof of the convexity of the quantum relative entropy: a task for which perspective-free proofs are technically demanding [29].

Von Neumann entropy: Let $f(\boldsymbol{X}) = \boldsymbol{X}\log(\boldsymbol{X})$ denote the von Neumann quantum entropy of a density matrix $\boldsymbol{X}$. Then, its perspective function is $g_f(\boldsymbol{X}, \boldsymbol{Y}) = \boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}} \log(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}) \boldsymbol{Y}^{\frac{1}{2}}$. When $\boldsymbol{X}$ and $\boldsymbol{Y}$ commute, this perspective can be equivalently written as

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{X}^{\frac{1}{2}} \log(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}) \boldsymbol{X}^{\frac{1}{2}} & \text{if } \boldsymbol{X}, \boldsymbol{Y} \succ \boldsymbol{0}, \\ +\infty & \text{otherwise.} \end{cases}$$

Note that various generalizations of the relative entropy for matrices have been proposed in the quantum physics literature [45]. However, these different definitions agree on the set of commuting matrices, hence can be used interchangeably for optimization purposes (see Remark 4).

Remark 9 (Quantum relative entropy cone) Note the epigraph of g_f, namely,

$$\mathcal{K}_{\text{mat}}^{\text{op, rel}} = \{(\boldsymbol{X}_1, \boldsymbol{X}_2, \boldsymbol{X}_3) \in \mathcal{S}^n \times \mathcal{S}^n_{++} \times \mathcal{S}^n_{++} : \boldsymbol{X}_1 \succeq -\boldsymbol{X}_2^{\frac{1}{2}} \log(\boldsymbol{X}_2^{-\frac{1}{2}} \boldsymbol{X}_3 \boldsymbol{X}_2^{-\frac{1}{2}}) \boldsymbol{X}_2^{\frac{1}{2}}\},$$

is a convex cone which can be approximated using semidefinite cones and optimized over using either the `Matlab` package `CVXQuad` (see [33]), or optimized over directly using an interior point method for asymmetric cones [49][4]. Consequently, this is a tractable object which models the matrix logarithm and Von Neumann entropy (and their perspectives).

Finally, Table 2 relates the matrix perspectives discussed above with their scalar analogs.

[4] Specifically, if we are interested in quantum relative entropy problems where we minimize the trace of $\boldsymbol{X}_1$, as occurs in the context of the MPRT, we may achieve this using the domain-driven solver developed by [49]. However, we are not aware of any IPMs which can currently optimize over the full quantum relative entropy cone.

Table 2: Analogy between perspectives of scalars and perspectives of matrix convex functions.

	Perspective of function			Matrix perspective of function		
Type	$f(x):\mathbb{R}\to\mathbb{R}$	$g_f(\boldsymbol{x},t)$	Ref.	f	g_f	Ref.
Quadratic	x^2	x^2/t	[4]	$\boldsymbol{X}^\top \boldsymbol{X}$	$\boldsymbol{X}^\top \boldsymbol{Y}^\dagger \boldsymbol{X}$	[11]
Power	$-x^\alpha : 0<\alpha<1$	$-x^\alpha t^{1-\alpha}$	[15]	$-\boldsymbol{X}^\alpha$	$-\boldsymbol{Y}^{\frac{1-\alpha}{2}}\boldsymbol{X}^\alpha \boldsymbol{Y}^{\frac{1-\alpha}{2}}$	Prop. 3
Log	$-\log(x)$	$-t\log(\frac{x}{t})$	[15]	$-\log(\boldsymbol{X})$	$-\boldsymbol{Y}^{\frac{1}{2}}\log\left(\boldsymbol{X}^{-\frac{1}{2}}\boldsymbol{Y}\boldsymbol{X}^{-\frac{1}{2}}\right)\boldsymbol{Y}^{\frac{1}{2}}$	[33]
Entropy	$x\log(x)$	$x\log(\frac{x}{t})$	[15]	$\boldsymbol{X}\log(\boldsymbol{X})$	$\boldsymbol{X}^{\frac{1}{2}}\log(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}})\boldsymbol{X}^{\frac{1}{2}}$	[53, 29]

3.6 Matrix Perspective Cuts

We now generalize the perspective cuts of [37, 42] from vectors to matrices and cardinality to rank constraints. Let us reconsider the previously defined mixed-projection optimization problem:

$$\min_{\boldsymbol{Y}\in\mathcal{Y}_n^k}\min_{\boldsymbol{X}\in\mathcal{S}_+^n} \quad \langle \boldsymbol{C},\boldsymbol{X}\rangle+\mu\cdot\mathrm{tr}(\boldsymbol{Y})+\mathrm{tr}(f(\boldsymbol{X})) \text{ s.t. } \langle \boldsymbol{A}_i,\boldsymbol{X}\rangle=b_i \quad \forall i\in[m],\ \boldsymbol{X}=\boldsymbol{Y}\boldsymbol{X},\boldsymbol{X}\in\mathcal{K},$$

where similarly to [37] we assume that $f(\mathbf{0})=\mathbf{0}$ to simplify the cut derivation procedure. Letting $\boldsymbol{\theta}$ model the epigraph of f via $\boldsymbol{\theta}\succeq f(\boldsymbol{X})$ and $\boldsymbol{S}$ be a subgradient of f at $\bar{\boldsymbol{X}}$, we have:

$$\boldsymbol{\theta}\succeq f(\bar{\boldsymbol{X}})\boldsymbol{Y}+\boldsymbol{S}^\top(\boldsymbol{X}-\bar{\boldsymbol{X}}\boldsymbol{Y}), \tag{23}$$

which if $f(\boldsymbol{X})=\boldsymbol{X}^2$ as discussed previouslyreduces to

$$\boldsymbol{\theta}^i\succeq \bar{\boldsymbol{X}}(2\boldsymbol{X}-\bar{\boldsymbol{X}}\boldsymbol{Y}),$$

which is precisely the analog of perspective cuts in the vector case. Note however that these cuts require semidefinite constraints to impose, which suggests they may not be as practically useful. For instance, our prior work [11]'s outer-approximation scheme for low-rank problems has a non-convex QCQOP master problem, which can only be currently solved using `Gurobi`, while `Gurobi` currently does not support semidefinite constraints.

We remark however that the inner product of Equation (23) with an arbitrary PSD matrix supplies a valid linear inequality. Two interesting cases of this observation arise when we take the inner product of the cut with either a rank-one matrix or the identity matrix.

Taking an inner product with the identity matrix supplies the inequality:

$$\mathrm{tr}(\boldsymbol{\theta})\geq\langle f(\bar{\boldsymbol{X}}),\boldsymbol{Y}\rangle+\langle \boldsymbol{S},\boldsymbol{X}-\bar{\boldsymbol{X}}\boldsymbol{Y}\rangle \quad \forall \boldsymbol{Y}\in\mathcal{Y}_n^k. \tag{24}$$

Moreover, by analogy to [10, Section 3.4], if we "project out" the $\boldsymbol{X}$ variables by decomposing the problem into a master problem in $\boldsymbol{Y}$ and subproblems in $\boldsymbol{X}$ then this cut becomes the Generalized Benders Decomposition cuts derived in our prior work [11, Equation (17)].

Alternatively, taking the inner product of the cut with a rank-one matrix $\boldsymbol{b}\boldsymbol{b}^\top$ gives:

$$\boldsymbol{b}^\top\boldsymbol{\theta}\boldsymbol{b}\geq\boldsymbol{b}^\top\left(f(\bar{\boldsymbol{X}})\boldsymbol{Y}+\boldsymbol{S}^\top(\boldsymbol{X}-\bar{\boldsymbol{X}}\boldsymbol{Y})\right)\boldsymbol{b}.$$

A further improvement is actually possible: rather than requiring that the semidefinite inequality is non-negative with respect to one rank-one matrix, we can require that it is simultaneously non-negative in the directions $\boldsymbol{v}^1$ and $\boldsymbol{v}^2$. This supplies the second-order cone [64, Eqn. (8)] cut:

$$\begin{pmatrix} \boldsymbol{v}^1 \\ \boldsymbol{v}^2 \end{pmatrix}^\top \left(\boldsymbol{\theta} - f(\bar{\boldsymbol{X}})\boldsymbol{Y} - \boldsymbol{S}^\top(\boldsymbol{X} - \bar{\boldsymbol{X}}\boldsymbol{Y})\right) \begin{pmatrix} \boldsymbol{v}^1 \\ \boldsymbol{v}^2 \end{pmatrix}^\top \succeq \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

The analysis in this section suggests that applying a perspective cut decomposition scheme out-of-the-box may be impractical, but leaves the door open to adaptations of the scheme which account for the projection matrix structure.

4 Examples and Perspective Relaxations

In this section, we apply the MRPT to several important low-rank problems, in addition to the previously discussed reduced-rank regression problem (Section 1.1). We also recall Theorem 2 to demonstrate that applying the MPRT to spectral functions which feature in these problems actually gives the convex hull of relevant substructures.

4.1 Matrix Completion

Given a sample $(A_{i,j} : (i,j) \in \mathcal{I} \subseteq [n] \times [n])$ of a matrix $\boldsymbol{A} \in \mathcal{S}^n_+$, the matrix completion problem is to reconstruct the entire matrix, by assuming $\boldsymbol{A}$ is approximately low-rank [19]. Letting $\mu, \gamma > 0$ be penalty multipliers, this problem admits the formulation:

$$\min_{\boldsymbol{X} \in \mathcal{S}^n_+} \quad \sum_{(i,j) \in \mathcal{I}} (X_{i,j} - A_{i,j})^2 + \frac{1}{2\gamma}\|\boldsymbol{X}\|_F^2 + \mu \cdot \mathrm{Rank}(\boldsymbol{X}). \tag{25}$$

Applying the MPRT to the $\|\boldsymbol{X}\|_F^2 = \mathrm{tr}(\boldsymbol{X}^\top \boldsymbol{X})$ term demonstrates that this problem is equivalent to the mixed-projection problem:

$$\min_{\boldsymbol{X}, \boldsymbol{\theta} \in \mathcal{S}^n_+, \boldsymbol{Y} \in \mathcal{Y}^n_n} \quad \sum_{(i,j) \in \mathcal{I}} (X_{i,j} - A_{i,j})^2 + \frac{1}{2\gamma}\mathrm{tr}(\boldsymbol{\theta}) + \mu \cdot \mathrm{tr}(\boldsymbol{Y}) \quad \text{s.t.} \quad \begin{pmatrix} \boldsymbol{Y} & \boldsymbol{X} \\ \boldsymbol{X} & \boldsymbol{\theta} \end{pmatrix} \succeq \boldsymbol{0},$$

and relaxing $\boldsymbol{Y} \in \mathcal{Y}^n_n$ to $\boldsymbol{Y} \in \mathrm{Conv}(\mathcal{Y}^n_n) = \{\boldsymbol{Y} \in \mathcal{S}^n : \boldsymbol{0} \preceq \boldsymbol{Y} \preceq \mathbb{I}\}$ supplies a valid relaxation. We now argue that this relaxation is often high-quality, by demonstrating that the MPRT supplies the convex envelope of $t \geq \frac{1}{2\gamma}\|\boldsymbol{X}\|_F^2 + \mu \cdot \mathrm{Rank}(\boldsymbol{X})$, via the following corollary to Theorem 2:

Corollary 1

$$\textit{Let} \quad \mathcal{S} = \left\{(\boldsymbol{Y}, \boldsymbol{X}, \boldsymbol{\theta}) \in \mathcal{Y}^k_n \times \mathcal{S}^n_+ \times \mathcal{S}^n : \boldsymbol{\theta} \succeq \boldsymbol{X}^\top \boldsymbol{X}, u\boldsymbol{Y} \succeq \boldsymbol{X} \succeq \ell \boldsymbol{Y}\right\}$$

be a set where $\ell, u \in \mathbb{R}_+$. Then, this set's convex hull is given by:

$$\mathcal{S}^c = \left\{(\boldsymbol{Y}, \boldsymbol{X}, \boldsymbol{\theta}) \in \mathcal{S}^n_+ \times \mathcal{S}^n_+ \times \mathcal{S}^n : \boldsymbol{Y} \preceq \mathbb{I}, \mathrm{tr}(\boldsymbol{Y}) \leq k, u\boldsymbol{Y} \succeq \boldsymbol{X} \succeq \ell \boldsymbol{Y}, \begin{pmatrix} \boldsymbol{Y} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{\theta} \end{pmatrix} \succeq \boldsymbol{0}\right\}.$$

4.2 Tensor Completion

A central problem in machine learning is to reconstruct a d-tensor $\boldsymbol{\mathcal{X}}$ given a subsample of its entries $(A_{i_1,\dots i_d} : (i_1, \dots i_d) \in \mathcal{I} \subseteq [n_1] \times [n_2] \times \dots \times [n_d])$, by assuming that the tensor is low-rank. Since even evaluating the rank of a tensor is NP-hard [50], a popular approach for solving this problem is to minimize the reconstruction error while constraining the ranks of different unfoldings of the tensor [see, e.g., 40]. After imposing Frobenius norm regularization and letting $\|\cdot\|_{HS} = \sqrt{\sum_{i_1=1}^{n_1} \dots \sum_{i_d=1}^{n_d} X_{i_1,\dots,i_d}^2}$ denote the (second-order cone representable) Hilbert-Schmidt norm of a tensor, this leads to optimization problems of the form:

$$\min_{\mathcal{X} \in \mathbb{R}^{n_1 \times \dots \times n_d}} \quad \sum_{(i_1,\dots i_d) \in \mathcal{I}} (\mathcal{A}_{i_1,\dots i_d} - \mathcal{X}_{i_1,\dots i_d})^2 + \sum_{i=1}^{n} \|\mathcal{X}_{(i)}\|_F^2 \quad \text{s.t.} \quad \text{Rank}(\mathcal{X}_{(i)}) \le k \quad \forall i \in [n]. \tag{26}$$

Similarly to low-rank matrix completion, it is tempting to apply the MRPT to model the $\boldsymbol{X}_{(i)}^\top \boldsymbol{X}_{(i)}$ term for each mode-n unfolding. We now demonstrate this supplies a tight approximation of the convex hull of the sum of the regularizers, via the following lemma (proof omitted, follows in the spirit of [42, Lemma 4]):

Lemma 3

Let $\mathcal{Q} = \left\{ (\rho, \boldsymbol{Y}_1, \dots, \boldsymbol{Y}_m, \boldsymbol{X}_1, \dots, \boldsymbol{X}_m, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_m) : \rho \ge \sum_{i=1}^{m} q_i \text{tr}(\boldsymbol{\theta}_i), (\boldsymbol{X}_i, \boldsymbol{Y}_i, \boldsymbol{\theta}_i) \in \mathcal{S}^i \ \forall i \in [m] \right\}$

be a set where $l_i, u_i, q_i \in \mathbb{R}_+^n \ \forall i \in [m]$*, and* $\mathcal{S}_i$ *is a set of the same form as* $\mathcal{S}$*, but* l, u *are replaced by* l_i, u_i*. Then, an extended formulation of this set's convex hull is given by:*

$$\mathcal{Q}^c = \left\{ (\rho, \boldsymbol{Y}_1, \dots, \boldsymbol{Y}_m, \boldsymbol{X}_1, \dots, \boldsymbol{X}_m, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_m) : \rho \ge \sum_{i=1}^{m} q_i \text{tr}(\boldsymbol{\theta}_i), (\boldsymbol{X}_i, \boldsymbol{Y}_i, \boldsymbol{\theta}_i) \in \mathcal{S}_i^c \ \forall i \in [m] \right\}.$$

Lemma 3 suggests that the MPRT may improve algorithms which aim to recover tensors of low slice rank. For instance, in low-rank tensor problems where (26) admits multiple local solutions, solving the convex relaxation coming from $\mathcal{Q}^c$ and greedily rounding may give a high-quality initial point for an alternating minimization method such as the method of [31], and indeed allow such a strategy to return better solutions than if it were initialized at a random point.

Note however that Lemma 3 does not necessarily give the convex hull of the sum of the regularizers, since the regularization terms involve different slices of the same tensor and thus interact; see also [69] for a related proof that the tensor trace norm does not give the convex envelope of the sum of ranks of slices.

4.3 Low-Rank Factor Analysis

An important problem in statistics, psychometrics and economics is to decompose a covariance matrix $\boldsymbol{\Sigma} \in \mathcal{S}_+^n$ into a low-rank matrix $\boldsymbol{X} \in \mathcal{S}_+^n$ plus a diagonal matrix $\boldsymbol{\Phi} \in \mathcal{S}_+^n$, as explored by Bertsimas et al. [8] and references therein. This corresponds to solving:

$$\min_{\boldsymbol{X}, \boldsymbol{\Phi} \in \mathcal{S}_+^n} \ \|\boldsymbol{\Sigma} - \boldsymbol{\Phi} - \boldsymbol{X}\|_q^q \text{ s.t. } \text{Rank}(\boldsymbol{X}) \le k, \ \Phi_{i,j} = 0, \forall i, j \in [n] : i \ne j, \ \|\boldsymbol{X}\|_\sigma \le M \tag{27}$$

where $q \geq 1$, $\|\boldsymbol{X}\|_q = (\sum_{i=1}^n \lambda_i(\boldsymbol{X})^q)^{\frac{1}{q}}$ denotes the matrix q-norm, and we constrain the spectral norm of $\boldsymbol{X}$ via a big-M constraint for the sake of tractability.

This problem's objective involves minimizing $\mathrm{tr}\,(\boldsymbol{\Sigma} - \boldsymbol{\Phi} - \boldsymbol{X})^q$, and it is not immediately obvious how to either apply the technique in the presence of the $\boldsymbol{\Phi}$ variables or alternatively seperate out the $\boldsymbol{\Phi}$ term and apply the MPRT to an appropriate ($\boldsymbol{\Phi}$-free) substructure. To proceed, let us therefore first consider its scalar analog, obtaining the convex closure of the following set:

$$\mathcal{T} = \{(x, y, z, t) \in \mathbb{R} \times \mathbb{R} \times \{0, 1\} \times \mathbb{R}^+ : t \geq |x + y - d|^q, |x| \leq M, x = 0 \text{ if } z = 0\},$$

where $d \in \mathbb{R}$ and $q \geq 1$ are fixed constants, and we require that $|x| \leq M$ for the sake of tractability. We obtain the convex closure via the following proposition (proof deferred to Appendix B):

Proposition 5 *The convex closure of the set $\mathcal{T}$, $\mathcal{T}^c$, is given by:*

$$\mathcal{T}^c = \left\{(x, y, z, t) \in \mathbb{R} \times \mathbb{R} \times [0, 1] \times \mathbb{R}^+ : \exists \beta \geq 0 : \; t \geq \frac{|y - \beta - d(1-z)|^q}{(1-z)^{q-1}} + \frac{|x + \beta - dz|^q}{z^{q-1}}, \; |x| \leq Mz\right\}.$$

Remark 10 To check that this set is indeed a valid convex relaxation, observe that if $z = 0$ then $x = 0$ and $x = -\beta \implies \beta = 0$ and $t \geq |y - d|^q$, while if $z = 1$ then $y = \beta$ and $t \geq |x + y - d|^q$.

Observe that $\mathcal{T}^c$ can be modeled using two power cones and one inequality constraint.

Proposition 5 suggests that we can obtain high-quality convex relaxations for low-rank factor analysis problems via a judicious use of the matrix power cone. Namely, introduce an epigraph matrix $\boldsymbol{\theta}$ to model the eigenvalues of $(\boldsymbol{\Sigma} - \boldsymbol{\Phi} - \boldsymbol{X})^q$ and an orthogonal projection matrix $\boldsymbol{Y}_2$ to model the span of $\boldsymbol{X}$. This then leads to the following matrix power cone representable relaxation:

$$\begin{aligned}
\min_{\boldsymbol{X}, \boldsymbol{\Phi}, \boldsymbol{\theta}, \boldsymbol{Y}_1, \boldsymbol{Y}_2 \in \mathcal{S}^n_+, \boldsymbol{\beta} \in \mathcal{S}^n} \quad & \mathrm{tr}(\boldsymbol{\theta}) \\
\text{s.t.} \quad & \boldsymbol{\theta} \succeq \boldsymbol{Y}_1^{\frac{1-q}{2}} (\boldsymbol{Y}_1^{\frac{1}{2}} \boldsymbol{\Sigma} \boldsymbol{Y}_1^{\frac{1}{2}} - \boldsymbol{\beta} - \boldsymbol{\Phi}) \boldsymbol{Y}_1^{\frac{1-q}{2}} + \boldsymbol{Y}_2^{\frac{1-q}{2}} (\boldsymbol{Y}_2^{\frac{1}{2}} \boldsymbol{\Sigma} \boldsymbol{Y}_2^{\frac{1}{2}} + \boldsymbol{\beta} - \boldsymbol{X}) \boldsymbol{Y}_2^{\frac{1-q}{2}}, \\
& \boldsymbol{Y}_1 + \boldsymbol{Y}_2 = \mathbb{I}, \mathrm{tr}(\boldsymbol{Y}) \leq k, \Phi_{i,j} = 0, \forall i, j \in [n] : i \neq j, \\
& \boldsymbol{\Phi} \preceq \boldsymbol{X}, \boldsymbol{X} \preceq M\boldsymbol{Y}_2, -\boldsymbol{X} \preceq M\boldsymbol{Y}_2.
\end{aligned}$$

4.4 Optimal Experimental Design

Letting $\boldsymbol{A} \in \mathbb{R}^{n \times m}$ where $m \geq n$ be a matrix of linear measurements of the form $y_i = \boldsymbol{a}_i^\top \boldsymbol{\beta} + \epsilon_i$ from an experimental setting, the D-optimal experimental design problem (a.k.a. the sensor selection problem) is to pick $k \leq m$ of these experiments in order to make the most accurate estimate of $\boldsymbol{\beta}$ possible, by solving [see 48, 70, for a modern approach]:

$$\max_{\boldsymbol{z} \in \{0,1\}^n : \boldsymbol{e}^\top \boldsymbol{z} \leq k} \quad \log\det_\epsilon \left(\sum_{i \in [n]} z_i \boldsymbol{a}_i \boldsymbol{a}_i^\top \right), \tag{28}$$

where we define $\log\det_\epsilon(\boldsymbol{X}) = \log\det(\boldsymbol{X} + \epsilon\mathbb{I})$ for $\epsilon > 0$ to be the pseudo log-determinant of a rank-deficient PSD matrix, which can be thought of as imposing an uninformative prior of importance ϵ on the experimental design process. Since $\log\det(\boldsymbol{X}) = \mathrm{tr}(\log(\boldsymbol{X}))$, a valid convex relaxation is given by:

$$\max_{\boldsymbol{z}\in[0,1]^n, \boldsymbol{\theta}\in\mathcal{S}^n_+} \quad \mathrm{tr}(\boldsymbol{\theta}) \quad \text{s.t.} \quad \log\left(\boldsymbol{A}\mathrm{Diag}(\boldsymbol{z})\boldsymbol{A}^\top + \epsilon\mathbb{I}\right) \succeq \boldsymbol{\theta},$$

which can be modeled using the quantum relative entropy cone, via $(-\boldsymbol{\theta}, \mathbb{I}, \boldsymbol{A}\mathrm{Diag}(\boldsymbol{z})\boldsymbol{A}^\top + \epsilon\mathbb{I}) \in \mathcal{K}^{\text{rel, op}}_{\text{mat}}$. This is equivalent to perhaps the most common relaxation of D-optimal design, as proposed by Boyd and Vandenberghe [15, Eqn. 7.2.6]. By formulating in terms of the quantum relative entropy cone, the identity term suggests this relaxation leaves something "on the table".

In this direction, let us apply the MPRT. Observe that $\boldsymbol{X} := \sum_{i\in[n]} z_i \boldsymbol{a}_i\boldsymbol{a}_i^\top$ is a rank-k matrix and thus at an optimal solution to the original problem there is some orthogonal projection matrix $\boldsymbol{Y}$ such that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$. Therefore, we can take the perspective function of $f(\boldsymbol{X}) = \log(\boldsymbol{X}+\epsilon\mathbb{I})$, and thereby obtain the following validand potentially much tighter when $k < n$convex relaxation:

$$\begin{aligned} \max_{\boldsymbol{z}\in[0,1]^n, \boldsymbol{\theta}, \boldsymbol{Y}\in\mathcal{S}^n_+} \quad & \mathrm{tr}(\boldsymbol{\theta}) + (n - \mathrm{tr}(\boldsymbol{Y}))\log(\epsilon) \\ \text{s.t.} \quad & \boldsymbol{Y}^{\frac{1}{2}} \log\left(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{A}\mathrm{Diag}(\boldsymbol{z})\boldsymbol{A}^\top\boldsymbol{Y}^{-\frac{1}{2}} + \epsilon\mathbb{I}\right)\boldsymbol{Y}^{\frac{1}{2}} \succeq \boldsymbol{\theta}, \boldsymbol{Y} \preceq \mathbb{I}, \mathrm{tr}(\boldsymbol{Y}) \leq k, \end{aligned} \tag{29}$$

which can be modeled via the quantum relative entropy cone: $(-\boldsymbol{\theta}, \boldsymbol{Y}, \boldsymbol{A}\mathrm{Diag}(\boldsymbol{z})\boldsymbol{A}^\top + \epsilon\boldsymbol{Y}) \in \mathcal{K}^{\text{rel, op}}_{\text{mat}}$. We now argue that this relaxation is high-quality, by demonstrating that the MPRT supplies the convex envelope of $t \geq -\log\det_\epsilon(\boldsymbol{X})$ under a low-rank constraint, via the following corollary to Theorem 2:

Corollary 2

$$\textit{Let} \quad \mathcal{S} = \left\{\boldsymbol{X} \in \mathcal{S}^n_+ : t \geq -\log\det_\epsilon(\boldsymbol{X}), \mathrm{Rank}(\boldsymbol{X}) \leq k\right\}$$

be a set where ϵ, k, t are fixed. Then, this set's convex hull is:

$$\begin{aligned} \mathcal{S}^c = \Big\{ & (\boldsymbol{Y}, \boldsymbol{X}) \in \mathcal{S}^n_+ \times \mathcal{S}^n_+ : \boldsymbol{0} \preceq \boldsymbol{Y} \preceq \mathbb{I}, \mathrm{tr}(\boldsymbol{Y}) \leq k, \\ & t \geq -\mathrm{tr}(\boldsymbol{Y}^{\frac{1}{2}} \log_\epsilon(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}})\boldsymbol{Y}^{\frac{1}{2}}) - (n - \mathrm{tr}(\boldsymbol{Y}))\log(\epsilon) \Big\}. \end{aligned}$$

Remark 11 Observe that (29)'s relaxation is not useful in the over-determined regime where $k \geq n$, since setting $\boldsymbol{Y} = \mathbb{I}$ recovers (28)'s Boolean relaxation, which is considerably cheaper to optimize over. Accordingly, we only consider the under-determined regime in our experiments.

4.5 Non-Negative Matrix Optimization

Many important problems in combinatorial optimization, statistics and computer vision [see, e.g., 17] reduce to optimizing over the space of low-rank matrices with non-negative factors. An important special case is when we would like to find the low-rank completely positive matrix $\boldsymbol{X}$ which best approximates (in a least-squares sense) a given matrix $\boldsymbol{A} \in \mathcal{S}^n_+$, i.e., perform non-negative principal component analysis. Formally, we have the problem:

$$\min_{\boldsymbol{X}\in\mathcal{C}^n_+ : \mathrm{Rank}(\boldsymbol{X}) \leq k} \quad \|\boldsymbol{X} - \boldsymbol{A}\|_F^2, \tag{30}$$

where $\mathcal{C}^n_+ := \{\boldsymbol{U}\boldsymbol{U}^\top : \boldsymbol{U} \in \mathbb{R}^{n\times n}_+\}$ denotes the cone of $n \times n$ completely positive matrices.

Applying the MPRT to the strongly convex $\frac{1}{2}\|\boldsymbol{X}\|_F^2$ term in the objective therefore yields the following completely positive program:

$$\min_{\boldsymbol{X}\in\mathcal{C}^n_+, \boldsymbol{Y}, \boldsymbol{\theta}\in\mathcal{S}^n} \quad \frac{1}{2}\mathrm{tr}(\boldsymbol{\theta}) - \langle \boldsymbol{X}, \boldsymbol{A}\rangle + \frac{1}{2}\|\boldsymbol{A}\|_F^2 \text{ s.t. } \boldsymbol{Y} \preceq \mathbb{I},\ \mathrm{tr}(\boldsymbol{Y}) \leq k,\ \begin{pmatrix} \boldsymbol{Y} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{\theta}\end{pmatrix} \in S^{2n}_+. \quad (31)$$

Interestingly, since (31)'s reformulation has a linear objective, some extreme point in its relaxation is optimal, which means we can relax the requirement that $\boldsymbol{Y}$ is a projection matrix without loss of optimality and the computational complexity of the problem is entirely concentrated in the completely positive cone. Unfortunately however, completely positive optimization itself is intractable. Nonetheless, it can be approximated by replacing the completely positive cone with the doubly non-negative cone, $\mathcal{S}^n_+ \cap \mathbb{R}^{n\times n}_+$. Namely, we instead solve

$$\min_{\boldsymbol{X}\in\mathcal{S}^n_+\cap\mathbb{R}^{n\times n}_+, \boldsymbol{Y}, \boldsymbol{\theta}\in\mathcal{S}^n} \frac{1}{2}\mathrm{tr}(\boldsymbol{\theta}) - \langle \boldsymbol{X}, \boldsymbol{A}\rangle + \frac{1}{2}\|\boldsymbol{A}\|_F^2 \text{ s.t. } \begin{pmatrix} \boldsymbol{Y} & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{\theta}\end{pmatrix} \in S^{2n}_+,\ \boldsymbol{Y} \preceq \mathbb{I},\ \mathrm{tr}(\boldsymbol{Y}) \leq k. \quad (32)$$

Unfortunately, rounding a solution to (32) to obtain a completely positive $\boldsymbol{X}$ is non-trivial. Indeed, according to Ge and Ye [41], there is currently no effective mechanism for rounding doubly non-negative programs. Nonetheless, as we shall see in our numerical results, there are already highly effective heuristic methods for completely positive matrix factorization, and combining our relaxation with such a procedure offers certificates of near optimality in a tractable fashion.

Remark 12 If $\boldsymbol{X} = \boldsymbol{D}\boldsymbol{\Pi}$ is a monomial matrix, i.e., decomposable as the product of a diagonal matrix $\boldsymbol{D}$ and a permutation matrix $\boldsymbol{\Pi}$, as occurs in binary optimization problems such as k-means clustering problems among others [c.f. 63], then it follows that $(\boldsymbol{X}^\top\boldsymbol{X})^\dagger \geq \boldsymbol{0}$ [see 66] and thus $\boldsymbol{Y} := \boldsymbol{X}(\boldsymbol{X}^\top\boldsymbol{X})^\dagger\boldsymbol{X}^\top$ is elementwise non-negative. In this case, the doubly non-negative relaxation (32) should be strengthened by requiring that $\boldsymbol{Y} \geq \boldsymbol{0}$.

5 Numerical Results

In this section, we evaluate the algorithmic strategies derived in the previous section, implemented in `Julia` 1.5 using `JuMP.jl` 0.21.6 and `Mosek` 9.1 to solve the conic problems considered here. 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/MatrixPerspectiveSoftware`.

5.1 Reduced Rank Regression

In this section, we compare our convex relaxations for reduced rank regression developed in the introduction and laid out in (6)-(7)which we refer to as "Persp" and "DCL" respectivelyagainst the nuclear norm estimator proposed by [57] ("NN"), who solve

$$\min_{\boldsymbol{\beta}\in\mathbb{R}^{p\times n}} \quad \frac{1}{2m}\|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta}\|_F^2 + \frac{1}{2\gamma}\|\boldsymbol{\beta}\|_F^2 + \mu\|\boldsymbol{\beta}\|_*. \quad (33)$$

Similarly to [57], we attempt to recover rank$-k_{true}$ estimators $\boldsymbol{\beta}_{\text{true}} = \boldsymbol{U}\boldsymbol{V}^\top$, where each entry of $\boldsymbol{U} \in \mathbb{R}^{p\times k_{true}}, \boldsymbol{V} \in \mathbb{R}^{n\times k_{true}}$ is i.i.d. standard Gaussian $\mathcal{N}(0,1)$, the matrix $\boldsymbol{X} \in \mathbb{R}^{m\times p}$ contains i.i.d. standard Gaussian $\mathcal{N}(0,1)$ entries, $\boldsymbol{Y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{E}$, and $E_{i,j} \sim \mathcal{N}(0,\sigma)$ injects a small amount of i.i.d. noise. We set $n = p = 50, k = 10$, $\gamma = 10^6$, $\sigma = 0.05$ and vary m. To ensure a fair comparison, we cross-validate μ for both of our relaxations and [57]'s approach so as to minimize the MSE on a validation set. For each m, we evaluate 20 different values of μ which are distributed uniformly in logspace between 10^{-4} and 10^4 across 50 random instances for our convex relaxations and report on 100 different random instances with the "best" μ for each method and each p.

Rank recovery and statistical accuracy: Figures 1a-1c report the relative accuracy ($\|\boldsymbol{\beta}_{\text{est}} - \boldsymbol{\beta}_{\text{true}}\|_F/\|\boldsymbol{\beta}_{\text{true}}\|_F$), the rank (i.e., number of singular values of $\boldsymbol{\beta}_{\text{est}}$ which exceed 10^{-4}), and the out-of-sample MSE[5] $\|\boldsymbol{X}_{\text{new}}\boldsymbol{\beta}_{\text{est}} - \boldsymbol{y}_{\text{new}}\|_F^2$ (normalized by the out-of-sample MSE of the ground truth $\|\boldsymbol{X}_{\text{new}}\boldsymbol{\beta}_{\text{true}} - \boldsymbol{y}_{\text{new}}\|_F^2$). Results are averaged over 100 random instances per value of m. We observe thateven though we did not supply the true rank of the optimal solution in our formulationProblem (7)'s relaxation returns solutions of the correct rank ($k_{true} = 10$) and better MSE/accuracy, while our more "naive" perspective relaxation (6) and the nuclear norm approach (33) return solutions of a higher rank and lower accuracy. This suggests that (7)'s formulation should be considered as a more accurate estimator for reduced rank problems, and empirically confirms that the MPRT can lead to significant improvements in statistical accuracy.

Scalability w.r.t. m: Figure 1d reports the average time for `Mosek` to converge[6] to an optimal solution (over 100 random instances per m). Surprisingly, although (7) is a stronger relaxation than (6), it is one to two orders of magnitude faster than (6) and (33)'s formulations. The relative scalability of (7)'s formulation as $m$the number of observation increases can be explained by the fact that (7) considers a linear inner product of the Gram matrix $\boldsymbol{X}^\top\boldsymbol{X}$ with a semidefinite matrix $\boldsymbol{B}$ (the size of which does not vary with m) while Problems (6) and (33) have a quadratic inner product $\langle \boldsymbol{\beta}\boldsymbol{\beta}^\top, \boldsymbol{X}^\top\boldsymbol{X}\rangle$ which must be modeled using a rotated second-order cone constraint (the size of which depends on m), since modern conic solvers such as `Mosek` do not allow quadratic objective terms and semidefinite constraints to be simultaneously present (if they did, we believe all three formulations would scale similarly).

Scalability w.r.t p: Next, we evaluate the scalability of all three approaches in terms of their solve times and peak memory usage (measured using the `slurm` command `MaxRSS`), as $n = p$ increases. Fig. 2 depicts the average time to converge to an optimal solution (a) and peak memory consumption (b) by each method as we vary $n = p$ with $m = n$, $k = 10$, $\gamma = 10^6$, each μ fixed to the average cross-validated value found in the previous experiment, a peak memory budget of 120GB, a runtime budget of 12 hours, and otherwise the same experimental setup as previously

[5] Evaluated on $m = 1000$ new observations of $\boldsymbol{X}_j, \boldsymbol{Y}_k$ generated from the same distribution.

[6] We model the convex quadratic $\|\boldsymbol{X}\boldsymbol{\beta} - \boldsymbol{Y}\|_F^2$ using a rotated second order cone for formulations (6) and (33) (the quadratic term doesn't appear directly in (7)), model the nuclear norm term in (33) by introducing matrices $\boldsymbol{U}, \boldsymbol{V}$ such that $\begin{pmatrix} \boldsymbol{U} & \boldsymbol{\beta} \\ \boldsymbol{\beta}^\top & \boldsymbol{V} \end{pmatrix} \succeq \boldsymbol{0}$ and minimizing $\text{tr}(\boldsymbol{U}) + \text{tr}(\boldsymbol{V})$, use default Mosek parameters for all approaches.

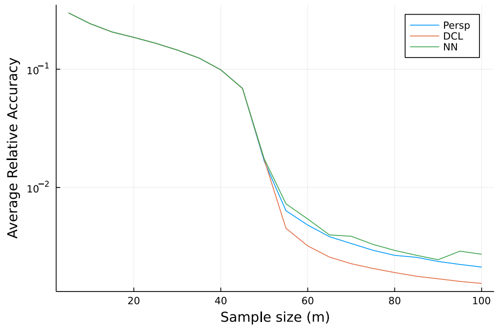

(a) Accuracy

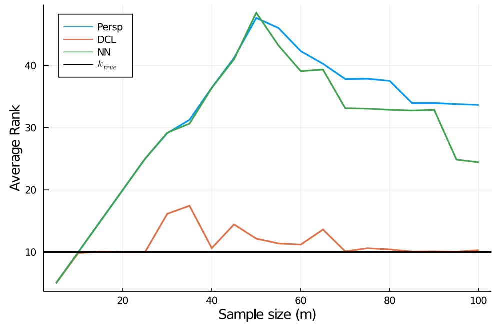

(b) Rank

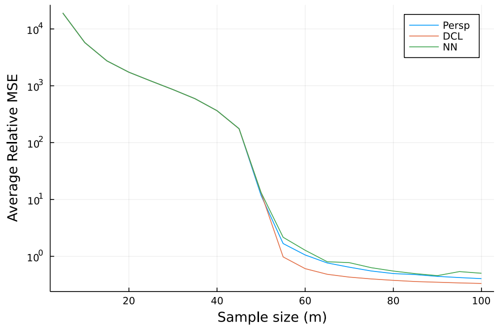

(c) Relative MSE

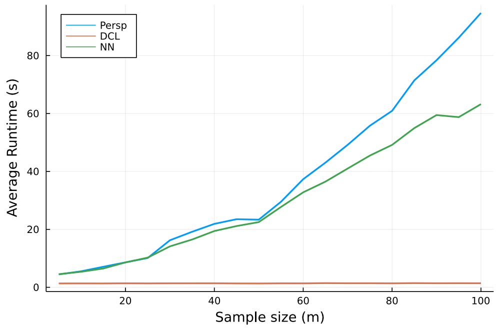

(d) Runtime

Fig. 1: Comparative performance, as the number of samples m increases, of formulations (6) (Persp, in blue), (7) (DCL, in orange) and (33) (NN, in green), averaged over 100 synthetic reduced rank regression instances where $n = p = 50$, $k_{true} = 10$. The hyperparameter μ was first cross-validated for all approaches separately.

(averaged over 20 random instances per n). We observe (7)'s relaxation is dramatically more scalable than the other two approaches considered, and can solve problems of nearly twice the size (4 times as many variables), and solves problems of a similar size in substantially less time and with substantially less peak memory consumption (40s vs. 1000s when $n = 100$). All in all, the proposed relaxation (7) seems to be the best method of the three considered.

5.2 Non-Negative Matrix Factorization

In this section, we benchmark the quality of our dual bound for non-negative matrix factorization laid out in Section 4.5 by using the non-linear reformulation strategy proposed by [18] (alternating least squares or ALS) to obtain upper bounds. Namely, we obtain upper bounds by solving for local minima of the problem

$$\min_{\boldsymbol{U} \in \mathbb{R}_+^{n \times k}} \quad \|\boldsymbol{U}\boldsymbol{U}^\top - \boldsymbol{A}\|_F^2. \tag{34}$$

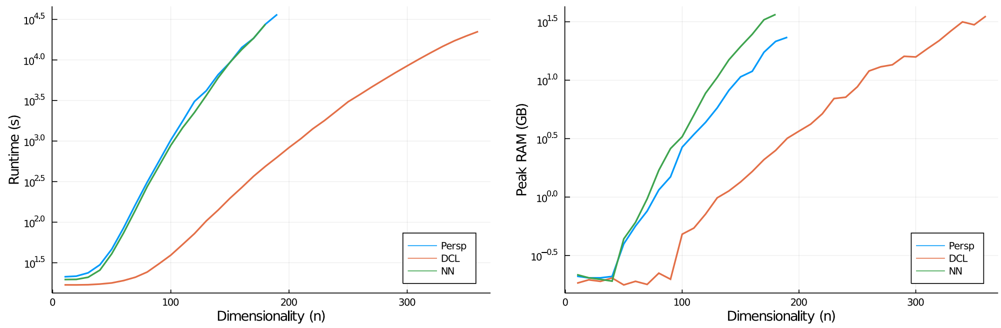

(a) Runtime (b) Peak Memory

Fig. 2: Average time to compute an optimal solution (left panel) and peak memory usage (right panel) vs. dimensionality $n = p$ for Problems (6) (Persp, in blue), (7) (DCL. in orange) and (33) (NN, in green) over 20 synthetic reduced rank regression instances where $k_{true} = 10$.

In our implementation of ALS, we obtain a local minimum by introducing a dummy variable $\boldsymbol{V}$ which equals $\boldsymbol{U}$ at optimality and alternating between solving the following two problems

$$\boldsymbol{U}_{t+1} = \arg \min_{\boldsymbol{U} \in \mathbb{R}_+^{n \times k}} \quad \|\boldsymbol{U}\boldsymbol{V}_t^\top - \boldsymbol{A}\|_F^2 + \rho_t \|\boldsymbol{U} - \boldsymbol{V}_t\|_F^2, \tag{35}$$

$$\boldsymbol{V}_{t+1} = \arg \min_{\boldsymbol{V} \in \mathbb{R}_+^{n \times k}} \quad \|\boldsymbol{U}_t\boldsymbol{V}^\top - \boldsymbol{A}\|_F^2 + \rho_t \|\boldsymbol{U}_t - \boldsymbol{V}\|_F^2, \tag{36}$$

where we set $\rho_t = \min(10^{-4} \times 2^{t-1}, 10^5)$ at the tth iteration in order that the final matrix is positive semidefinite, as advocated in [5, Section 5.2.3] (we cap ρ_t to avoid numerical instability). We iterate over solving these two problems from a random initialization point $\boldsymbol{V}_0$where each $V_{0,i,j}$ is i.i.d. standard uniformuntil either the objective value between iterations does not change by 10^{-4} or we exceed the maximum number of allowable iterations, which we set to 100.

To generate problem instances, we let $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{U}^\top + \boldsymbol{E}$ where $\boldsymbol{U} \in \mathbb{R}^{n \times k_{true}}$, each $U_{i,j}$ is uniform on $[0, 1]$, $E_{i,j} \sim \mathcal{N}(0, 0.0125k_{true})$, and set $A_{i,j} = 0$ if $A_{i,j} < 0$. We set $n = 50, k_{true} = 10$. We use the ALS heuristic to compute a feasible solution $\boldsymbol{X}$ and an upper-bound on the problem's objective value. By comparing it with the lower bound derived from our MPRT, we can assess the sub-optimality of the heuristic solution, which previously lacked optimality guarantees.

Figure 3 depicts the average relative in-sample MSE of the heuristic ($\|\boldsymbol{X} - \boldsymbol{A}\|_F / \|\boldsymbol{A}\|_F$) and the relative bound gap(UB-LB)/UB as we vary the target rank, averaged over 100 random synthetic instances. We observe that the method is most accurate and has the lowest MSE when k is set to $k_{true} = 10$, which confirms that the method can recover solutions of the correct rank. In addition, by combining the solution from OLS with our lower-bound, we can compute a duality gap and assert that the heuristic solution is $0\% - 3\%$-optimal, with the gap peaking at $k = k_{true}$ and stabilizing as $k \to n$. This echoes similar findings in k-means clustering and alternating current optimal power flow problems, where the SDO relaxation need not be near-tight in theory but nonetheless is nearly exact in practice [63, 51]. Further, this suggests our convex relaxation may be a powerful weapon for providing gaps for heuristics for non-negative matrix factorization, and particularly detecting when they are performing well or can be further improved.

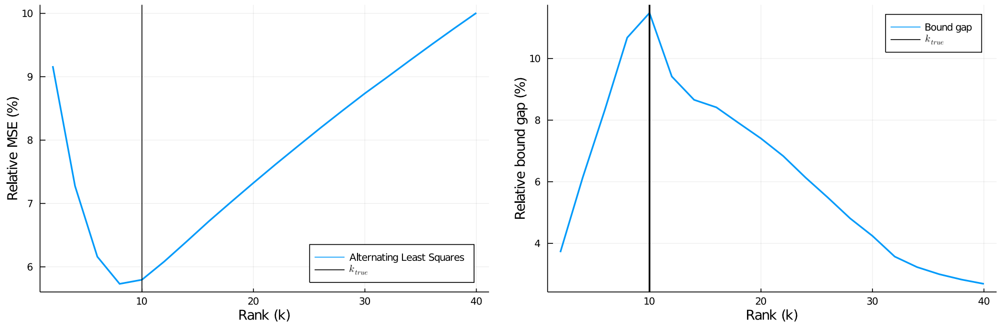

(a) Relative MSE

(b) Bound gap

Fig. 3: Average relative MSE and duality gap vs. target rank k using the ALS heuristic (UB) and the MPRT relaxation (LB). Results are averaged over 100 synthetic completely positive matrix factorization instances where $n = 50$, $k_{true} = 10$.

Figure 4 reports the time needed to compute both the upper bound and a lower bound solution as we vary the target rank.

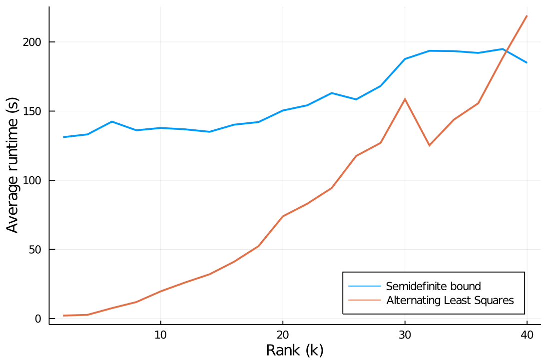

Fig. 4: Computational time to compute a feasible solution (ALS) and solve the relaxation (Semidefinite bound) vs. target rank k, averaged over 100 synthetic completely positive matrix factorization instances where $n = 50$, $k_{true} = 10$.

5.3 Optimal Experimental Design

In this section, we benchmark our dual bound for D-optimal experimental design (29) against the convex relaxation (28) and a greedy submodular maximization approach, in terms of both bound quality and the ability of all three approaches to generate high-quality feasible solutions. We round both relaxations to generate feasible solutions greedily, by setting the k largest z_i's in a continuous relaxation to 1, while for the submodular maximization approach we iteratively set the jth index of $\boldsymbol{z}$ to 1, where $\mathcal{S}$ is initially an empty set and we iteratively take

$$\mathcal{S} \leftarrow \mathcal{S} \cup \{j\} : j \in \arg\max_{i\in[n]\setminus\mathcal{S}} \left\{ \log\det_{\epsilon} \left(\sum_{l\in\mathcal{S}} z_l \boldsymbol{a}_l \boldsymbol{a}_l^\top + \boldsymbol{a}_i \boldsymbol{a}_i^\top \right) \right\}.$$

Interestingly, the greedy rounding approach enjoys rigorous approximation guarantees [see 48, 70], while the submodular maximization approach also enjoys strong guarantees [see 58].

We benchmark all methods in terms of their performance on synthetic D-optimal experimental design problems, where we let $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ be a matrix with i.i.d. $\mathcal{N}(0, \frac{1}{\sqrt{n}})$ entries. We set $n = 20, m = 10, \epsilon = 10^{-6}$ and vary $k < m$ over 20 random instances. Table 3 depicts the average relative bound gap, objective values, and runtimes for all 3 methods (we use the lower bound from (28)'s relaxation to compute the submodular bound gap). Note that all results for this experiment were generated on a standard Macbook pro laptop with a 2.9GHZ 6-core Intel i9 CPU using 16GB DDR4 RAM, CVX version 1.22, Matlab R2021a, and `Mosek` 9.1. Moreover, we optimize over (29)'s relaxation using the `CVXQuad` package developed by [33].

Table 3: Average runtime in seconds and relative bound gap per approach, over 20 random instances where $n = 10, m = 20$.

	Problem (28)+round		Submodular		Problem (29)+round	
k	Time(s)	Gap (%)	Time(s)	Gap (%)	Time(s)	Gap (%)
1	0.52	88.8	0.00	88.9	347.0	0.00
2	0.63	93.7	0.00	93.7	338.5	0.01
3	0.59	97.1	0.00	97.0	320.8	0.06
4	0.63	100.2	0.00	100.2	338.7	0.18
5	0.53	103.8	0.00	103.9	331.1	0.37
6	0.53	109.0	0.00	109.0	287.5	1.40
7	0.55	117.7	0.00	117.7	255.1	2.39
8	0.60	136.9	0.00	138.5	236.1	5.25
9	0.54	260.9	0.00	287.5	235.9	28.43

Relaxation quality: We observe that (29)'s relaxation is dramatically stronger than (28), offering bound gaps on the order of $0\% - 3\%$ when $k \leq 7$, rather than gaps of 90% or more. This confirms the efficacy of the MPRT, and demonstrates the value of taking low-rank constraints into account when designing convex relaxations, even when not obviously present.

Scalability: We observe that (29)'s relaxation is around two orders of magnitude slower than the other proposed approaches, largely because semidefinite approximations of quantum relative entropy are expensive, but is still tractable for moderate sizes. We believe, however, that the relaxation would scale significantly better if it were optimized over using an interior point method for non-symmetric cones [see, e.g., 71, 49], or an alternating minimization approach [see 34]. As such, (29)'s relaxation is potentially useful at moderate problem sizes with off-the-shelf software, or at larger problem sizes with problem-specific techniques such as alternating minimization.

6 Conclusion

In this paper, we introduced the Matrix Perspective Reformulation Technique (MPRT), a new technique for deriving tractable and often high-quality relaxations of a wide variety of low-rank problems. We also invoked the technique to derive the convex hulls of some frequently-studied low-rank sets, and provided examples where the technique proves useful in practice. This is significant and potentially useful to the community, because substantial progress on producing tractable upper bounds for low-rank problems has been made over the past decade, but until now almost no progress on tractable lower bounds has followed.

Future work could take three directions: (1) automatically detecting structures where the MPRT could be applied, as is already done for perspective reformulations in the MIO case by `CPLEX` and `Gurobi`, (2) developing scalable semidefinite-free techniques for solving the semidefinite relaxations proposed in this paper, and (3) combining the ideas in this paper and in our prior work [11] with custom branching strategies to solve low-rank problems to optimality at scale.

Acknowledgments: We are very grateful to two anonymous referees for useful and constructive comments. In particular, we would like to thank reviewer #1 for a very helpful refinement of our definition of the matrix perspective function, and reviewer #2 for suggesting the name *matrix perspective function* and supplying some new references on perspective operator functions.

References

1. M. S. Aktürk, A. Atamtürk, and S. Gürel. A strong conic quadratic reformulation for machine-job assignment with controllable processing times. *Operations Research Letters*, 37 (3):187–191, 2009.
2. F. Alizadeh. Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization*, 5(1):13–51, 1995.
3. A. Atamtürk and A. Gómez. Rank-one convexification for sparse regression. *arXiv:1901.10334*, 2019.
4. A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization: Analysis, algorithms, and engineering applications*, volume 2. SIAM Philadelphia, PA, 2001.
5. D. P. Bertsekas. *Nonlinear programming.* Athena Scientific Belmont MA, 3rd edition, 2016.
6. D. Bertsimas and B. Van Parys. Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. *The Annals of Statistics*, 48(1):300–323, 2020.
7. D. Bertsimas, A. King, and R. Mazumder. Best subset selection via a modern optimization lens. *The Annals of Statistics*, pages 813–852, 2016.
8. D. Bertsimas, M. S. Copenhaver, and R. Mazumder. Certifiably optimal low rank factor analysis. *Journal of Machine Learning Research*, 18(1):907–959, 2017.
9. D. Bertsimas, J. Pauphilet, and B. Van Parys. Sparse regression: Scalable algorithms and empirical performance. *Statistical Science*, 35(4):555–578, 2020.
10. D. Bertsimas, R. Cory-Wright, and J. Pauphilet. A unified approach to mixed-integer optimization problems with logical constraints. *SIAM Journal on Optimization*, 31(3): 2340–2367, 2021.

11. D. Bertsimas, R. Cory-Wright, and J. Pauphilet. Mixed-projection conic optimization: A new paradigm for modeling rank constraints. *Operations Research, Articles in Advance*, 2021.
12. D. Bertsimas, R. Cory-Wright, and J. Pauphilet. Solving large-scale sparse PCA to certifiable (near) optimality. *Journal of Machine Learning Research*, 23(13):1–35, 2022.
13. R. Bhatia. *Matrix analysis*, volume 169. Springer Science & Business Media New York, 2013.
14. D. Bienstock. Eigenvalue techniques for convex objective, nonconvex optimization problems. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 29–42. Springer, 2010.
15. S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge University Press, Cambridge, UK, 2004.
16. S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. Studies in Applied Mathematics, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1994.
17. S. Burer. On the copositive representation of binary and continuous nonconvex quadratic programs. *Mathematical Programming*, 120(2):479–495, 2009.
18. S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
19. E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717, 2009.
20. E. Carlen. Trace inequalities and quantum entropy: An introductory course. *Entropy and the quantum*, 529:73–140, 2010.
21. S. Ceria and J. Soares. Convex programming for disjunctive convex optimization. *Mathematical Programming*, 86(3):595–614, 1999.
22. R. Chares. *Cones and interior-point algorithms for structured convex optimization involving powers and exponentials.* PhD thesis, UCL-Université Catholique de Louvain, 2009.
23. P. L. Combettes. Perspective functions: Properties, constructions, and examples. *Set-Valued and Variational Analysis*, 26(2):247–264, 2018.
24. B. Dacorogna and P. Maréchal. The role of perspective functions in convexity, polyconvexity, rank-one convexity and separate convexity. *Journal of Convex Analysis*, 15(2):271–284, 2008.
25. C. Davis. Various averaging operations onto subalgebras. *Illinois Journal of Mathematics*, 3 (4):538–553, 1959.
26. H. Dong, K. Chen, and J. Linderoth. Regularization vs. relaxation: A conic optimization perspective of statistical variable selection. *arXiv:1510.06083*, 2015.
27. A. Ebadian, I. Nikoufar, and M. E. Gordji. Perspectives of matrix convex functions. *Proceedings of the National Academy of Sciences*, 108(18):7313–7314, 2011.
28. E. Effros and F. Hansen. Non-commutative perspectives. *Annals of Functional Analysis*, 5 (2):74–79, 2014.
29. E. G. Effros. A matrix convexity approach to some celebrated quantum inequalities. *Proceedings of the National Academy of Sciences*, 106(4):1006–1008, 2009.
30. 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.
31. V. F. Farias and A. A. Li. Learning preferences with side information. *Management Science*, 65(7):3131–3149, 2019.

32. H. Fawzi and J. Saunderson. Lieb's concavity theorem, matrix geometric means, and semidefinite optimization. *Linear Algebra and its Applications*, 513:240–263, 2017.
33. H. Fawzi, J. Saunderson, and P. A. Parrilo. Semidefinite approximations of the matrix logarithm. *Foundations of Computational Mathematics*, 19(2):259–296, 2019.
34. L. Faybusovich and C. Zhou. Self-concordance and matrix monotonicity with applications to quantum entanglement problems. *Applied Mathematics and Computation*, 375:125071, 2020.
35. M. Fazel, H. Hindi, and S. P. Boyd. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003.*, volume 3, pages 2156–2162. IEEE, 2003.
36. M. Fischetti, I. Ljubić, and M. Sinnl. Redesigning Benders decomposition for large-scale facility location. *Management Science*, 63(7):2146–2162, 2016.
37. A. Frangioni and C. Gentile. Perspective cuts for a class of convex 0–1 mixed integer programs. *Mathematical Programming*, 106(2):225–236, 2006.
38. A. Frangioni and C. Gentile. A computational comparison of reformulations of the perspective relaxation: SOCP vs. cutting planes. *Operations Research Letters*, 37(3):206–210, 2009.
39. A. Frangioni, C. Gentile, and J. Hungerford. Decompositions of semidefinite matrices and the perspective reformulation of nonseparable quadratic programs. *Mathematics of Operations Research*, 45(1):15–33, 2020.
40. S. Gandy, B. Recht, and I. Yamada. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Problems*, 27(2):025010, 2011.
41. D. Ge and Y. Ye. On doubly positive semidefinite programming relaxations. *Optimization Online*, 2010.
42. O. Günlük and J. Linderoth. Perspective reformulations of mixed integer nonlinear programs with indicator variables. *Mathematical Programming*, 124(1-2):183–205, 2010.
43. S. Han, A. Gómez, and A. Atamtürk. 2x2 convexifications for convex quadratic optimization with indicator variables. *arXiv:2004.07448*, 2020.
44. H. Hazimeh, R. Mazumder, and A. Saab. Sparse regression at scale: Branch-and-bound rooted in first-order optimization. *Mathematical Programming, articles in advance*, pages 1–42, 2021.
45. F. Hiai and D. Petz. The proper formula for relative entropy and its asymptotics in quantum probability. *Communications in Mathematical Physics*, 143(1):99–114, 1991.
46. J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex analysis and minimization algorithms I: Fundamentals*, volume 305. Springer Science & Business Media Berlin, 2013.
47. R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, New York, 1985.
48. S. Joshi and S. Boyd. Sensor selection via convex optimization. *IEEE Transactions on Signal Processing*, 57(2):451–462, 2008.
49. M. Karimi and L. Tunçel. Domain-driven solver (DDS): a MATLAB-based software package for convex optimization problems in domain-driven form. *arXiv preprint arXiv:1908.03075*, 2019.
50. T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3): 455–500, 2009.

51. J. Lavaei and S. H. Low. Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems*, 27(1):92–107, 2011.
52. A. S. Lewis. Convex analysis on the Hermitian matrices. *SIAM Journal on Optimization*, 6 (1):164–177, 1996.
53. E. H. Lieb and M. B. Ruskai. Proof of the strong subadditivity of quantum-mechanical entropy. with an appendix by B. Simon. *Journal of Mathematical Physics*, 14:1938–1941, 1973.
54. P. Maréchal. On the convexity of the multiplicative potential and penalty functions and related topics. *Mathematical Programming*, 89(3):505–516, 2001.
55. P. Maréchal. On a functional operation generating convex functions, part 1: duality. *Journal of Optimization Theory and Applications*, 126(1):175–189, 2005.
56. P. Maréchal. On a functional operation generating convex functions, part 2: algebraic properties. *Journal of Optimization Theory and Applications*, 126(2):357–366, 2005.
57. S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, pages 1069–1097, 2011.
58. G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functionsi. *Mathematical Programming*, 14(1):265–294, 1978.
59. L. T. Nguyen, J. Kim, and B. Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.
60. M. L. Overton and R. S. Womersley. On the sum of the largest eigenvalues of a symmetric matrix. *SIAM Journal on Matrix Analysis and Applications*, 13(1):41–45, 1992.
61. M. L. Overton and R. S. Womersley. Optimality conditions and duality theory for minimizing sums of the largest eigenvalues of symmetric matrices. *Mathematical Programming*, 62(1-3): 321–357, 1993.
62. G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of Operations Research*, 23(2):339–358, 1998.
63. J. Peng and Y. Wei. Approximating K-means-type clustering via semidefinite programming. *SIAM Journal on Optimization*, 18(1):186–205, 2007.
64. F. Permenter and P. Parrilo. Partial facial reduction: simplified, equivalent SDPs via approximations of the PSD cone. *Mathematical Programming*, 171(1-2):1–54, 2018.
65. M. Pilanci, M. J. Wainwright, and L. El Ghaoui. Sparse learning via Boolean relaxations. *Mathematical Programming*, 151(1):63–87, 2015.
66. R. Plemmons and R. Cline. The generalized inverse of a nonnegative matrix. *Proceedings of the American Mathematical Society*, pages 46–50, 1972.
67. J. Renegar. *A mathematical view of interior-point methods in convex optimization*, volume 3. Society for Industrial and Applied Mathematics, 2001.
68. R. T. Rockafellar. *Convex analysis*. Number 28. Princeton university press, 1970.
69. B. Romera-Paredes and M. Pontil. A new convex relaxation for tensor completion. *arXiv preprint arXiv:1307.4653*, 2013.
70. M. Singh and W. Xie. Approximation algorithms for D-optimal design. *Mathematics of Operations Research*, 45:1193–1620, 2020.
71. A. Skajaa and Y. Ye. A homogeneous interior-point algorithm for nonsymmetric convex conic optimization. *Mathematical Programming*, 150(2):391–422, 2015.

72. R. A. Stubbs. *Branch-and-cut methods for mixed 0-1 convex programming.* PhD thesis, Northwestern University, 1996.
73. R. A. Stubbs and S. Mehrotra. A branch-and-cut method for 0-1 mixed convex programming. *Mathematical Programming*, 86(3):515–532, 1999.
74. A. L. Wang and F. Kılınç-Karzan. On the tightness of SDP relaxations of QCQPs. *Mathematical Programming, Articles in Advance*, pages 1–41, 2021.
75. L. Wei, A. Gómez, and S. Küçükyavuz. Ideal formulations for constrained convex optimization problems with indicator variables. *Mathematical Programming*, 192(1):57–88, 2022.
76. H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of semidefinite programming: theory, algorithms, and applications*, volume 27. Springer Science & Business Media, 2012.
77. W. Xie and X. Deng. Scalable algorithms for the sparse ridge regression. *SIAM Journal on Optimization*, 30(4):3359–3386, 2020.
78. C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.
79. 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 Journal on Computing*, 26(4):690–703, 2014.

A Background on Operator Functions

In this work, we make repeated use of operator functions, i.e., functions defined from the spectral decomposition of a matrix. Namely, for any function $\omega : \mathbb{R} \to \mathbb{R}$, its corresponding operator function $f_\omega : \mathcal{S}^n \to \mathcal{S}^n$ is defined as

$$f_\omega(\boldsymbol{X}) = \boldsymbol{U}\,\mathrm{Diag}(\omega(\lambda_1^x), \dots, \omega(\lambda_n^x))\boldsymbol{U}^\top$$

where $\boldsymbol{X} = \boldsymbol{U}\,\mathrm{Diag}(\lambda_1^x, \dots, \lambda_n^x)\boldsymbol{U}^\top$ is an eigendecomposition of $\boldsymbol{X}$. In this appendix, we present some common examples and useful properties of operator functions.

A.1 Examples: Matrix exponential and logarithm

For self-consistency of the paper, we now define the matrix exponential and logarithm functions and summarize their properties. These results are well known and can be found in modern matrix analysis textbooks [see, e.g., 13]

Definition 2 (Matrix exponential) Let $\boldsymbol{X} \in \mathcal{S}^n$ be a symmetric matrix with eigendecomposition $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^\top$. Letting $\exp(\boldsymbol{\Lambda}) = \mathrm{diag}(e^{\lambda_1}, e^{\lambda_2}, \dots, e^{\lambda_n})$, we define $\exp(\boldsymbol{X}) := \boldsymbol{U}\exp(\boldsymbol{\Lambda})\boldsymbol{U}^\top$.

Proposition 6 *The matrix exponential,* $\exp : \mathcal{S}^n \to \mathcal{S}^n_+$*, satisfies the following properties:*

– *Power series expansion:* $\exp(\boldsymbol{X}) = \mathbb{I} + \sum_{i=1}^{\infty} \frac{1}{i!}\boldsymbol{X}^i$.
– *Trace monotonicity:* $\boldsymbol{X} \preceq \boldsymbol{Y} \implies \mathrm{tr}(\exp(\boldsymbol{X})) \leq \mathrm{tr}(\exp(\boldsymbol{Y}))$.
– *Golden-Thompson-inequality:* $\mathrm{tr}(\exp(\boldsymbol{X}+\boldsymbol{Y})) \leq \mathrm{tr}(\exp(\boldsymbol{X})) + \mathrm{tr}(\exp(\boldsymbol{Y}))$.

Remark 13 The matrix exponential is not monotone: $\boldsymbol{X} \preceq \boldsymbol{Y} \;\;\not\Longrightarrow\;\; \exp(\boldsymbol{X}) \preceq \exp(\boldsymbol{Y})$ [13, Ch.V].

Definition 3 (Matrix logarithm) Let $\boldsymbol{X} \in \mathcal{S}^n$ be a symmetric matrix with eigendecomposition $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^\top$. Letting $\log(\boldsymbol{\Lambda}) = \operatorname{diag}(\log(\lambda_1), \log(\lambda_2), \dots, \log(\lambda_n))$, we have $\log(\boldsymbol{X}) := \boldsymbol{U}\log(\boldsymbol{\Lambda})\boldsymbol{U}^\top$.

Proposition 7 *The matrix logarithm,* $\log(\boldsymbol{X}) : \mathcal{S}^n_{++} \to \mathcal{S}^n$*, satisfies the following properties:*

– *Operator monotonicity:* $\boldsymbol{X} \preceq \boldsymbol{Y} \implies \log(\boldsymbol{X}) \preceq \log(\boldsymbol{Y})$.
– *Functional inversion:* $\log(\exp(\boldsymbol{X})) = \boldsymbol{X} \quad \forall \boldsymbol{X} \in \mathcal{S}^n$.
– *Jacobi formula I:* $\operatorname{tr}(\log(\boldsymbol{X})) = \log\det(\boldsymbol{X})$.
– *Jacobi formula II:* $\exp\left(\frac{1}{n}\operatorname{tr}\log(\boldsymbol{X})\right) = \det(\boldsymbol{X})^{\frac{1}{n}}$.

A.2 Properties of operator functions

Among other properties, one can show that the trace of operator functions is invariant under an orthogonal rotation, i.e., $\operatorname{tr}(f_\omega(\boldsymbol{X})) = \operatorname{tr}(f_\omega(\boldsymbol{U}^\top\boldsymbol{X}\boldsymbol{U}))$ for any orthogonal rotation $\boldsymbol{U}$. Also, if ω is analytical, then f_ω is also analytical with the same Taylor expansion.

In our analysis (in particular the proof of Proposition 3), we will use this simple bound on $\boldsymbol{v}^\top f_\omega(\boldsymbol{A})\boldsymbol{v}$ in the case where ω is convex:

Lemma 4 *Consider a convex function* $\omega : \mathbb{R} \to \mathbb{R}$ *and a symmetric matrix* $\boldsymbol{A} \in \mathcal{S}^n$*. Consider a unit vector* $\boldsymbol{v}$*. Then,*

$$\boldsymbol{v}^\top f_\omega(\boldsymbol{A})\boldsymbol{v} \geq \omega\left(\boldsymbol{v}^\top\boldsymbol{A}\boldsymbol{v}\right).$$

Proof Consider a spectral decomposition of $\boldsymbol{A}$, $\boldsymbol{A} = \sum_{i=1}^n \lambda_i \boldsymbol{u}_i\boldsymbol{u}_i^\top$. Then, $f_\omega(\boldsymbol{A}) = \sum_{i=1}^n \omega(\lambda_i)\boldsymbol{u}_i\boldsymbol{u}_i^\top$ and

$$\boldsymbol{v}^\top f_\omega(\boldsymbol{A})\boldsymbol{v} = \sum_{i=1}^n \omega(\lambda_i)\boldsymbol{v}^\top\boldsymbol{u}_i\boldsymbol{u}_i^\top\boldsymbol{v} \geq \omega\left(\sum_{i=1}^n \lambda_i\boldsymbol{v}^\top\boldsymbol{u}_i\boldsymbol{u}_i^\top\boldsymbol{v}\right) = \omega\left(\boldsymbol{v}^\top\boldsymbol{A}\boldsymbol{v}\right),$$

where the inequality comes from the convexity of ω since $\boldsymbol{v}^\top\boldsymbol{u}_i\boldsymbol{u}_i^\top\boldsymbol{v} = (\boldsymbol{u}_i^\top\boldsymbol{v})^2 \geq 0$ and $\sum_{i=1}^n \boldsymbol{v}^\top\boldsymbol{u}_i\boldsymbol{u}_i^\top\boldsymbol{v} = \boldsymbol{v}^\top\left(\sum_{i=1}^n \boldsymbol{u}_i\boldsymbol{u}_i^\top\right)\boldsymbol{v} = \|\boldsymbol{v}\|^2 = 1$.

B Omitted Proofs

In this section, we supply all omitted proofs, in the order the results were stated.

B.1 Proof of Proposition 3

Proof Fix $\boldsymbol{X} \in \mathcal{S}^n$. For $\boldsymbol{Y} \succ \boldsymbol{0}$, the perspective of f_ω is well-defined according to Definition 1. Now, consider an arbitrary $\boldsymbol{Y} \succeq \boldsymbol{0}$ and define $\boldsymbol{P}$ as the orthogonal projection onto the kernel of $\boldsymbol{Y}$, which is orthogonal to $\operatorname{Span}(\boldsymbol{Y})$. Then, $\boldsymbol{Y}_\varepsilon := \boldsymbol{Y} + \varepsilon\boldsymbol{P}$ for $\varepsilon > 0$ is invertible. The closure of the matrix perspective of f_ω is defined by continuity as the limit of $\boldsymbol{M}_\varepsilon := \boldsymbol{Y}_\varepsilon^{\frac{1}{2}} f_\omega \left(\boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \right) \boldsymbol{Y}_\varepsilon^{\frac{1}{2}}$ for $\varepsilon \to 0$.

Since the ranges of $\boldsymbol{Y}$ and $\boldsymbol{P}$ are orthogonal ($\boldsymbol{YP} = \boldsymbol{PY} = \boldsymbol{0}$), we have $\boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} = \boldsymbol{Y}^{-\frac{1}{2}} + \varepsilon^{-\frac{1}{2}} \boldsymbol{P}$, and

$$\boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} = \boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}} + \varepsilon^{-\frac{1}{2}} \boldsymbol{P} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}} + \varepsilon^{-\frac{1}{2}} \boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{P} + \varepsilon^{-1} \boldsymbol{P} \boldsymbol{X} \boldsymbol{P}.$$

Note that $\lim_{\varepsilon \to 0} \boldsymbol{Y}_\varepsilon^{\frac{1}{2}} = \boldsymbol{Y}^{\frac{1}{2}}$ but $\lim_{\varepsilon \to 0} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \neq \boldsymbol{Y}^{-\frac{1}{2}}$. We now distinguish two cases.

Case 1: If $\operatorname{span}(\boldsymbol{X}) \subseteq \operatorname{span}(\boldsymbol{Y})$, $\boldsymbol{XP} = \boldsymbol{PX} = \boldsymbol{0}$ so

$$\begin{aligned} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} &= \boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}, \\ \boldsymbol{M}_\varepsilon = \boldsymbol{Y}_\varepsilon^{\frac{1}{2}} f_\omega \left(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}} \right) \boldsymbol{Y}_\varepsilon^{\frac{1}{2}} &\quad \to_{\varepsilon \to 0} \boldsymbol{Y}^{\frac{1}{2}} f_\omega \left(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}} \right) \boldsymbol{Y}^{\frac{1}{2}}. \end{aligned}$$

Case 2: If $\operatorname{span}(\boldsymbol{X}) \not\subseteq \operatorname{span}(\boldsymbol{Y})$, consider an orthonormal basis of $\mathbb{R}^n$ such that $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k$ is an eigenbasis of $\operatorname{Span}(\boldsymbol{Y})$ (with respective eigenvalues $\lambda_1^y, \ldots \lambda_k^y$) and $\boldsymbol{u}_{k+1}, \ldots, \boldsymbol{u}_n$ is a basis of $\operatorname{Span}(\boldsymbol{Y})^\perp = \operatorname{Ker}(\boldsymbol{Y})$. By assumption, $k < n$ and there exists $j > k$ such that $\boldsymbol{u}_j^\top \boldsymbol{X} \boldsymbol{u}_j \neq 0$. Without loss of generality, we shall assume $\boldsymbol{u}_n^\top \boldsymbol{X} \boldsymbol{u}_n \neq 0$. We show that the matrix $\boldsymbol{M}_\varepsilon$ goes to infinity as $\varepsilon \to 0$ by showing that $\boldsymbol{u}_n^\top \boldsymbol{M}_\varepsilon \boldsymbol{u}_n$ diverges.

Since $\boldsymbol{Y}_\varepsilon^{\pm\frac{1}{2}} \boldsymbol{u}_n = \varepsilon^{\pm\frac{1}{2}} \boldsymbol{u}_n$, we have

$$\boldsymbol{u}_n^\top \boldsymbol{M}_\varepsilon \boldsymbol{u}_n = \varepsilon\, \boldsymbol{u}_n^\top f_\omega \left(\boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \right) \boldsymbol{u}_n \quad \geq \varepsilon\, \omega \left(\boldsymbol{u}_n^\top \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{u}_n \right) = \varepsilon\, \omega \left(\varepsilon^{-1} \boldsymbol{u}_n^\top \boldsymbol{X} \boldsymbol{u}_n \right),$$

where the inequality follows from the convexity of ω and Lemma 4. By Assumption 1,

$$\lim_{\varepsilon \to 0} \varepsilon \omega \left(\varepsilon^{-1} \boldsymbol{u}_n^\top \boldsymbol{X} \boldsymbol{u}_n \right) = \omega_\infty (\boldsymbol{u}_n^\top \boldsymbol{X} \boldsymbol{u}_n) = +\infty,$$

because $\boldsymbol{u}_n^\top \boldsymbol{X} \boldsymbol{u}_n \neq 0$ and ω is coercive. □

We now provide a simple extension of Proposition 3 that will prove useful later in our exposition.

Corollary 3 *Consider a function $\omega : \mathbb{R} \to \mathbb{R}$ satisfying Assumption 1 and denote its associated operator function f_ω. Consider a closed set $\mathcal{X} \subseteq \mathcal{S}^n$ and define*

$$f(\boldsymbol{X}) = \begin{cases} f_\omega(\boldsymbol{X}) & \text{if } \boldsymbol{X} \in \mathcal{X}, \\ +\infty & \text{otherwise.} \end{cases}$$

Then, the closure of the matrix perspective of f is, for any $\boldsymbol{X} \in \mathcal{S}^n$, $\boldsymbol{Y} \in \mathcal{S}_+^n$,

$$g_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{Y}^{\frac{1}{2}} f_\omega(\boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}) \boldsymbol{Y}^{\frac{1}{2}} & \text{if } \operatorname{Span}(\boldsymbol{X}) \subseteq \operatorname{Span}(\boldsymbol{Y}), \boldsymbol{Y} \succeq \boldsymbol{0}, \boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}} \in \mathcal{X}, \\ \infty & \text{otherwise,} \end{cases}$$

where $\boldsymbol{Y}^{-\frac{1}{2}}$ denotes the pseudo-inverse of the square root of $\boldsymbol{Y}$.

Proof Fix $\boldsymbol{X} \in \mathcal{S}^n$ and $\boldsymbol{Y} \in \mathcal{S}_+^n$. From Proposition 3, we know that $g_f(\boldsymbol{X}, \boldsymbol{Y}) = +\infty$ if $\operatorname{Span}(\boldsymbol{X}) \not\subseteq \operatorname{Span}(\boldsymbol{Y})$. Let us assume that $\operatorname{Span}(\boldsymbol{X}) \subseteq \operatorname{Span}(\boldsymbol{Y})$. Following the same construction as in the proof of Proposition 3, we obtain a sequence $\boldsymbol{Y}_\varepsilon$ that converges to $\boldsymbol{Y}$ as $\varepsilon \to 0$ and such that $\boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}_\varepsilon^{-\frac{1}{2}} = \boldsymbol{Y}^{-\frac{1}{2}} \boldsymbol{X} \boldsymbol{Y}^{-\frac{1}{2}}$, which concludes the proof. □

B.2 Proof of Lemma 2

Proof First, let us observe that $\mathcal{C}_{\boldsymbol{X}} = \{\boldsymbol{M} : \boldsymbol{MX} = \boldsymbol{XM}\}$ is a closed subset of $\mathcal{S}^n$, contains the identity, and is closed under multiplication and transposition, also know as a Von Neumann subalgebra [see 20, Section 4 for a detailed treatment of projections onto subalgebras]. The orthogonal projection of a semidefinite matrix onto $\mathcal{C}_{\boldsymbol{X}}$ is also semidefinite and has the same trace [20, Theorem. 4.13], so

$$\operatorname{tr}\left(\boldsymbol{Y}_{|\boldsymbol{X}}\right) = \operatorname{tr}\left(\boldsymbol{Y}\right).$$

Furthermore, since $\boldsymbol{Y} \mapsto g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})$ is matrix convex, Carlen [20, Theorem 4.16] yields

$$g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y}_{|\boldsymbol{X}}) \preceq g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})_{|\boldsymbol{X}}.$$

Taking the trace on both sides and using that $\operatorname{tr}\left(g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})_{|\boldsymbol{X}}\right) = \operatorname{tr}\left(g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})\right)$ concludes the proof. □

Remark 14 If $\boldsymbol{X} = \sum_{i\in[n]} \lambda_i \boldsymbol{u}_i \boldsymbol{u}_i^\top$ is a spectral decomposition of $\boldsymbol{X}$, then the projection of any matrix $\boldsymbol{Y}$ onto the commutant of $\boldsymbol{X}$ can be computed as $Y_{|\boldsymbol{X}} = \sum_{i\in[n]} \lambda_i \boldsymbol{u}_i \boldsymbol{u}_i^\top \boldsymbol{Y} \boldsymbol{u}_i \boldsymbol{u}_i^\top$. This operation is known in the literature as pinching [25, 20].

In other words, taking the projection of $\boldsymbol{Y}$ onto the commutant of $\boldsymbol{X}$ is a trace preserving operation that can only reduce the value of $\operatorname{tr}\left(g_{f_\omega}(\boldsymbol{X}, \cdot)\right)$. In this paper, we invoke the projection onto the commutant of $\boldsymbol{X}$ for theoretical purposes, not computational ones. So we are not interested in how to compute $\boldsymbol{Y}_{|\boldsymbol{X}}$ in practice. Note that, according to Proposition 2(a), Lemma 2 holds if f_ω is matrix convex.

B.3 Counterexample to joint convexity of trace of matrix perspective of cube

In this section, we demonstrate by counterexample that if ω is a convex and continuous function then, even though the trace of its matrix extension, $\operatorname{tr}(f_\omega)$, is convex [c.f. 20, Theorem 2.10], the trace of its matrix perspective need not be convex.

Specifically, let us consider $\omega(x) = x^3$. In this case, ω is convex on $\mathbb{R}_+$, f_ω is not matrix convex, but $\operatorname{tr}(f_\omega)$ is matrix convex. We have that

$$\operatorname{tr}(g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y})) = \operatorname{tr}\left(\boldsymbol{X}\boldsymbol{Y}^\dagger \boldsymbol{X}\boldsymbol{Y}^\dagger \boldsymbol{X}\right)$$

for $\boldsymbol{X} \in \operatorname{Span}(\boldsymbol{Y}), \boldsymbol{X}, \boldsymbol{Y} \in \mathcal{S}^n_+$. Let us now consider

$$\boldsymbol{Y}_1 = \begin{pmatrix} 0.160378 & 0.343004 \\ 0.343004 & 0.764592 \end{pmatrix}, \quad \boldsymbol{Y}_2 = \begin{pmatrix} 0.0859208 & 0.181976 \\ 0.181976 & 0.52666 \end{pmatrix},$$

$$\boldsymbol{X}_1 = \begin{pmatrix} 0.242865 & 0.543321 \\ 0.543321 & 1.26604 \end{pmatrix}, \quad \boldsymbol{X}_2 = \begin{pmatrix} 0.0595215 & 0.241702 \\ 0.241702 & 1.0596 \end{pmatrix}.$$

Then, some elementary algebra reveals that

$$\mathrm{tr}\left[g_{f_\omega}\left(\tfrac{1}{2}\boldsymbol{X}_1+\tfrac{1}{2}\boldsymbol{X}_2, \tfrac{1}{2}\boldsymbol{Y}_1+\tfrac{1}{2}\boldsymbol{Y}_2\right)\right]=6.248327,$$

while

$$\tfrac{1}{2}\mathrm{tr}\left[g_{f_\omega}\left(\boldsymbol{X}_1,\boldsymbol{Y}_1\right)\right]+\tfrac{1}{2}\mathrm{tr}\left[g_{f_\omega}\left(\boldsymbol{X}_2,\boldsymbol{Y}_2\right)\right]=6.23977,$$

which verifies that $\mathrm{tr}(g_{f_\omega}(\boldsymbol{X},\boldsymbol{Y}))$ is not midpoint convex in $(\boldsymbol{X},\boldsymbol{Y})$, despite $\mathrm{tr}(f_\omega)$ being convex.

B.4 Proof of Proposition 5

Proof We use the proof technique laid out in [43, Section 3.1], namely writing $\mathcal{T}$ as the disjunction of two convex sets driven by whether z is active and applying Fourier-Motzkin elimination. That is, we have $\mathcal{T}=\mathcal{T}^1\cup\mathcal{T}^2$ where:

$$\begin{aligned}\mathcal{T}^1&=\left\{(0,y_1,0,t_1): t_1\geq |y_1-d|^q\right\},\\ \mathcal{T}^2&=\left\{(x_2,y_2,1,t_2): t_2\geq |x_2-y_2-d|^q, |x_2|\leq M\right\}.\end{aligned}$$

Moreover, a point (x,y,z,t) is in the convex hull $\mathcal{T}^c$ if and only if it can be written as a convex combination of points in $\mathcal{T}^1,\mathcal{T}^2$. Letting λ_1,λ_2 denote the weight of points in this system, we then have that $(x,y,z,t)\in\mathcal{T}^c$ if and only if the following system admits a solution:

$$\begin{aligned}&\lambda_1+\lambda_2=1,\\ &x=\lambda_2 x_2,\\ &y=\lambda_1 y_1+\lambda_2 y_2,\\ &t=\lambda_1 t_1+\lambda_2 t_2,\\ &z=\lambda_2,\\ &t_1\geq |y_1-d|^q,\\ &t_2\geq |x_2+y_2-d|^q,\\ &\lambda_1,\lambda_2\geq 0,\\ &|x_2|\leq M.\end{aligned}\tag{37}$$

For ease of computation, we now eliminate variables. First, one can substitute t_1,t_2 for their lower bounds in the definition of t and replace λ_2 with z to obtain

$$\begin{aligned}&\lambda_1+z=1,\\ &x=zx_2,\\ &y=\lambda_1 y_1+zy_2,\\ &t\geq \lambda_1|y_1-d|^q+z|x_2+y_2-d|^q,\\ &\lambda_1,z\geq 0,\\ &|x_2|\leq M.\end{aligned}\tag{38}$$

Next, we substitute x/z for x_2 and $(y - zy_2)/\lambda_1$ for y_1 to obtain

$$\begin{aligned} &\lambda_1 + z = 1,\ \lambda_1, z \geq 0,\ |x| \leq Mz \\ &t \geq \frac{1}{\lambda_1^{q-1}}|y - y_2 z - d(1-z)|^q + \frac{1}{z^{q-1}}|x + y_2 z - dz|^q. \end{aligned} \tag{39}$$

Finally, we let zy_2 be the free variable β and set $\lambda_1 = 1 - z$ to obtain the required convex set. □

C Generalizing the Matrix Perspective Reformulation Technique to Functions

We now demonstrate the MPRT can be extended to incorporate a different separability of eigenvalues assumption, at the price of (a possibly significant amount of) additional notations. For any symmetric matrix $\boldsymbol{X}$, let us denote $\lambda_i^{\downarrow}(\boldsymbol{X})$ the ith largest eigenvalue of $\boldsymbol{X}$. Before proceeding any further, we recall the following result, due to [4, Example 18.c], which provides a semidefinite representation of the sum of the k largest eigenvalues:

Lemma 5 (Representability of sums of largest eigenvalues) *Let $S_k(\boldsymbol{X}) := \sum_{i=1}^k \lambda_i^{\downarrow}(\boldsymbol{X})$ denote the sum of the k largest eigenvalues of a symmetric matrix $\boldsymbol{X} \in \mathcal{S}^n$. Then, the epigraph of S_k, $S_k(\boldsymbol{X}) \leq t_k$, admits the following semidefinite representation:*

$$t_k \geq ks_k + \mathrm{tr}(\boldsymbol{Z}_k),\ \boldsymbol{Z}_k + s_k\mathbb{I} \succeq \boldsymbol{X}, \boldsymbol{Z}_k \succeq \boldsymbol{0}.$$

Based on this result, we can relax the assumption that the penalty term $\Omega(\boldsymbol{X})$ corresponds to the trace of an operator function. Instead, we can assume:

Assumption 4 $\Omega(\boldsymbol{X}) = \sum_{i\in[n]} p_i \lambda_i^{\downarrow}(f_\omega(\boldsymbol{X}))$, *where $p_1 \geq \ldots \geq p_n \geq 0$ and where ω is a function satisfying Assumption 1 and whose associated operator function, f_ω, is matrix convex.*

This assumption is particularly suitable for Markov Chain problems [see, e.g., 15, Chapter 4.6], where we are interested in controlling the behaviour of the largest eigenvalue (which always equals 1) plus the second largest eigenvalue of a matrix. However, it might appear to be challenging to model, since, e.g., $\lambda_2^{\downarrow}(\boldsymbol{X})$ is a non-convex function. By applying a telescoping sum argument reminiscent of the one in [4, Prop. 4.2.1], namely

$$\Omega(\boldsymbol{X}) = \sum_{i=1}^n p_i \lambda_i^{\downarrow}(f(\boldsymbol{X})) = \sum_{i=1}^n (p_i - p_{i+1}) S_i(f(\boldsymbol{X}))$$

with the convention $p_{n+1} = 0$, Lemma 5 allows us to rewrite low-rank problems where $\Omega(\boldsymbol{X})$ satisfies Assumption 4 in the form:

$$\begin{aligned} \min_{\boldsymbol{Y}\in\mathcal{Y}_n^k}\ \min_{\substack{\boldsymbol{X}\in\mathcal{S}_+^n, \\ \boldsymbol{Z}_i\in\mathcal{S}_+^n, s_i, t_i\in\mathbb{R}_+\ \forall i\in[n]}}\ & \langle \boldsymbol{C}, \boldsymbol{X}\rangle + \mu \cdot \mathrm{tr}(\boldsymbol{Y}) + \sum_{i=1}^n (p_i - p_{i+1}) t_i \\ \text{s.t.}\ & \langle \boldsymbol{A}_i, \boldsymbol{X}\rangle = b_i\ \forall i \in [m],\ \boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X},\ \boldsymbol{X} \in \mathcal{K}, \\ & t_i \geq is_i + \mathrm{tr}(\boldsymbol{Z}_i),\ \boldsymbol{Z}_i + s_i\mathbb{I} \succeq f(\boldsymbol{X}), \boldsymbol{Z}_i \succeq \boldsymbol{0}\ \forall i \in [n], \end{aligned} \tag{40}$$

where t_i models the sum of the i largest eigenvalues of $f(\boldsymbol{X})$. Applying the MPRT then yields the following extension to Theorem 1:

Proposition 8 *Suppose Problem* (40) *attains a finite optimal value. Then, the following problem attains the same value:*

$$\min_{\boldsymbol{Y}\in\mathcal{Y}_n^k} \min_{\substack{\boldsymbol{X}\in\mathcal{S}_+^n,\\ \boldsymbol{Z}_i\in\mathcal{S}_+^n, s_i, t_i\in\mathbb{R}_+\ \forall i\in[n]}} \langle \boldsymbol{C}, \boldsymbol{X}\rangle + \mu\cdot\mathrm{tr}(\boldsymbol{Y}) + \sum_{i=1}^n (p_i - p_{i+1}) t_i \tag{41}$$
$$\begin{aligned} s.t.\quad & \langle \boldsymbol{A}_i, \boldsymbol{X}\rangle = b_i\ \forall i\in[m],\ \boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Y}^{-\frac{1}{2}}\in\mathcal{K},\\ & t_i \geq i s_i + i - \mathrm{tr}(\boldsymbol{Y}) + \mathrm{tr}(\boldsymbol{Z}_i)\ \forall i\in[n],\\ & \boldsymbol{Z}_i + s_i\mathbb{I} \succeq g_f(\boldsymbol{X},\boldsymbol{Y}) + \omega(0)(\mathbb{I}-\boldsymbol{Y}), \boldsymbol{Z}_i\succeq \boldsymbol{0}\ \forall i\in[n]. \end{aligned}$$

The proof of this reformulation is almost identical to the proof of Theorem 1, after observing that (20) holds not only for the traces but for the matrices directly, i.e., if $\boldsymbol{X}$ and $\boldsymbol{Y}\in\mathcal{Y}_n^k$ commute, we have

$$f(\boldsymbol{X}) = g_f(\boldsymbol{X},\boldsymbol{Y}) + \omega(0)(\mathbb{I}-\boldsymbol{Y}).$$

Problem (41) involves n times as many variables as Problem (18) and therefore supplies substantially less tractable relaxations. Nonetheless, it could be useful in specific instances. In the aforementioned Markov Chain mixing problem, $p_i - p_{i+1} = 0\ \forall i\geq k$ with $k=2$, so we can omit the variables which model the eigenvalues larger than 2 .

D Extension to the rectangular case

In this section, we extend the MPRT to the case where $\boldsymbol{X}$ is a generic $n\times m$ matrix and $f(\boldsymbol{X})$ is the convex quadratic penalty $f(\boldsymbol{X}) = \boldsymbol{X}^\top\boldsymbol{X}$. In this case, $\mathrm{tr}(f(\boldsymbol{X})) = \|\boldsymbol{X}\|_F^2$ is the squared Frobenius norm of $\boldsymbol{X}$.

First, observe that $f:\mathbb{R}^{n\times m}\to\mathcal{S}_+^m$. Alternatively, one could have considered $g(\boldsymbol{X}) = \boldsymbol{X}\boldsymbol{X}^\top\in\mathcal{S}_+^n$ and obtain the same penalty, i.e., $\mathrm{tr}(f(\boldsymbol{X})) = \mathrm{tr}(g(\boldsymbol{X}))$. In other words, one can arbitrarily choose whether f preserves the row or the column space of $\boldsymbol{X}$. By the Schur complement lemma, the epigraph is semidefinite representable via

$$\mathrm{epi}(f) := \left\{(\boldsymbol{X},\boldsymbol{\theta})\in\mathbb{R}^{n\times m}\times\mathcal{S}_+^m\ :\ \begin{pmatrix}\boldsymbol{\theta} & \boldsymbol{X}^\top\\ \boldsymbol{X} & \mathbb{I}\end{pmatrix}\succeq\boldsymbol{0}\right\},$$

so f is matrix convex.

In the symmetric case, we considered the matrix perspective of f at $(\boldsymbol{X},\boldsymbol{Y})$, where $\boldsymbol{Y}\succeq\boldsymbol{0}$ is a matrix controlling the range of $\boldsymbol{X}$. When $\boldsymbol{X}$ is no longer symmetric, it is natural to consider a matrix perspective function which involves two projection matrices, one of which models the row space and one which models the column space, as proposed in our prior work [11]. More precisely, for $\boldsymbol{Y},\boldsymbol{Z}\succ\boldsymbol{0}$ we define a perspective of f as

$$g_f(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{Z}) = \boldsymbol{Z}^{\frac{1}{2}} f(\boldsymbol{Y}^{-\frac{1}{2}}\boldsymbol{X}\boldsymbol{Z}^{-\frac{1}{2}})\boldsymbol{Z}^{\frac{1}{2}}. \tag{42}$$

For $f(\boldsymbol{X}) = \boldsymbol{X}^\top\boldsymbol{X}$, this function actually does not depend on $\boldsymbol{Z}$. Hence, we consider

$$\tilde{g}_f(\boldsymbol{X},\boldsymbol{Y}) = g_f(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{Z}) = \boldsymbol{X}^\top\boldsymbol{Y}^{-1}\boldsymbol{X}.$$

Extending this function to positive semidefinite $\boldsymbol{Y}$ using the same proof technique as in Proposition 3, we then obtain

$$\tilde{g}_f(\boldsymbol{X}, \boldsymbol{Y}) = \begin{cases} \boldsymbol{X}^\top \boldsymbol{Y}^\dagger \boldsymbol{X} & \text{if } \boldsymbol{Y} \succeq \boldsymbol{0},\ \text{Span}(\boldsymbol{X}) \subseteq \text{Span}(\boldsymbol{Y}), \\ \infty & \text{otherwise.} \end{cases}$$

Proof Fix $\boldsymbol{X} \in \mathcal{S}^n$ and $\boldsymbol{Y} \succeq \boldsymbol{0}$. As in the proof of Proposition 3 denote $\boldsymbol{P}$ the orthogonal projection onto the kernel of $\boldsymbol{Y}$, and define $\boldsymbol{Y}_\varepsilon := \boldsymbol{Y} + \varepsilon \boldsymbol{P}$ for $\varepsilon > 0$. Hence,

$$\boldsymbol{X}^\top \boldsymbol{Y}_\varepsilon^{-1} \boldsymbol{X} = \boldsymbol{X}^\top \boldsymbol{Y}^\dagger \boldsymbol{X} + \varepsilon^{-1} \boldsymbol{X}^\top \boldsymbol{P} \boldsymbol{X}.$$

The right-hand side admits a finite limit if and only if

$$\boldsymbol{X}^\top \boldsymbol{P} \boldsymbol{X} = \boldsymbol{0} \iff \text{Span}(\boldsymbol{X}) \subseteq \text{Ker}(\boldsymbol{P}) = \text{Span}(\boldsymbol{Y}). \quad \square$$

Furthermore, using the Schur complement lemma as in [11], one can show that $\tilde{g}_f$ is SDP-representable:

$$\text{epi}(\tilde{g}_f) = \left\{ (\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\theta}) \in \mathbb{R}^{n \times m} \times \mathcal{S}^n_+ \times \mathcal{S}^m \ : \ \begin{pmatrix} \boldsymbol{\theta} & \boldsymbol{X}^\top \\ \boldsymbol{X} & \boldsymbol{Y} \end{pmatrix} \succeq \boldsymbol{0} \right\},$$

and hence matrix convex.

Finally, we can easily check that Theorem 1 still holds in the symmetric case because (20) –which simplifies to $\text{tr}(f(\boldsymbol{X})) = \text{tr}(\tilde{g}_f(\boldsymbol{X}))$ in this case– holds for any $\boldsymbol{Y} \in \mathcal{Y}^k_n$ such that $\boldsymbol{X} = \boldsymbol{Y}\boldsymbol{X}$.

Remark 15 We believe the approach outlined above could be generalized to a broader class of function that generalizes operator functions to the non-symmetric case. Namely, we could consider functions of the form

$$f_\omega(\boldsymbol{X}) = \boldsymbol{V} \text{Diag}\left(\omega(\sigma^x_1), \ldots, \omega(\sigma^x_m)\right) \boldsymbol{V}^\top$$

where $\boldsymbol{X} = \boldsymbol{U} \text{Diag}\left(\sigma^x_1, \ldots, \sigma^x_m\right) \boldsymbol{V}^\top$ is a singular value decomposition of $\boldsymbol{X}$ and ω is a convex function satisfying Assumption 1. Again, f_ω could arbitrarily be defined as preserving $\boldsymbol{U}$ or $\boldsymbol{V}$. For these functions, the perspective $g_{f_\omega}(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{X})$ is well defined for $\boldsymbol{Y}, \boldsymbol{Z} \succ \boldsymbol{0}$. Unlike in the quadratic case, however, its value will depend on both $\boldsymbol{Y}$ and $\boldsymbol{Z}$. Developing the theoretical tools necessary to extend the MPRT to rectangular matrices, is therefore a question for future research.