POLYNOMIAL TIME GUARANTEES FOR THE BURER-MONTEIRO METHOD

DIEGO CIFUENTES AND ANKUR MOITRA

ABSTRACT. The Burer-Monteiro method is one of the most widely used techniques for solving large-scale semidefinite programs (SDP). The basic idea is to solve a nonconvex program in Y, where Y is an $n\times p$ matrix such that $X=YY^T$. In this paper, we show that this method can solve SDPs in polynomial time in an smoothed analysis setting. More precisely, we consider an SDP whose domain satisfies some compactness and smoothness assumptions, and slightly perturb the cost matrix and the constraints. We show that if $p \gtrsim \sqrt{2(1+\eta)m}$, where m is the number of constraints and $\eta>0$ is any fixed constant, then the Burer-Monteiro method can solve SDPs to any desired accuracy in polynomial time, in the setting of smooth analysis. Our bound on p approaches the celebrated Barvinok-Pataki bound in the limit as η goes to zero, beneath which it is known that the nonconvex program can be suboptimal.

Previous analyses were unable to give polynomial time guarantees for the Burer-Monteiro method, since they either assumed that the criticality conditions are satisfied exactly, or ignored the nontrivial problem of computing an approximately feasible solution. We address the first problem through a novel connection with tubular neighborhoods of algebraic varieties. For the feasibility problem we consider a least squares formulation, and provide the first guarantees that do not rely on the restricted isometry property.

1. INTRODUCTION

Consider a *semidefinite program* (SDP) in the space of $n\times n$ symmetric matrices $\mathbb{S}^n$, involving m equality constraints:

$$(SDP) \qquad \min_{X\in\mathbb{S}^n} \; C\bullet X \quad \text{such that} \quad \mathcal{A}(X)=b, \quad X\succeq 0,$$

where $C\in\mathbb{S}^n$, $b\in\mathbb{R}^m$ and $\mathcal{A}:\mathbb{S}^n\to\mathbb{R}^m$, $X\mapsto(A_1\bullet X,\dots,A_m\bullet X)$ is a linear map. Though interior point methods can solve (SDP) in polynomial time, they typically run into memory problems for large values of n. The Burer-Monteiro method [10,11] is one of the most widely used procedures for large scale problems. Several papers have worked in understanding the practical success of this method, see e.g., [8,9,11]. Although several results have been shown, they all fall short of showing that one can reach an approximately optimal solution of (SDP) in polynomial time. In this paper we prove the

Key words and phrases. Semidefinite programming, Burer-Monteiro, Low rank factorization.

first polynomial time guarantees for the Burer-Monteiro method, under a compactness and smoothness assumption on the domain.

The Burer-Monteiro method consists in writing $X = YY^T$ for some $Y \in \mathbb{R}^{n\times p}$, and solving the following nonconvex optimization problem:

$$(BM) \qquad \min_{Y\in\mathfrak{M}} C \bullet YY^T, \quad \text{where} \quad \mathfrak{M} := \{Y \in \mathbb{R}^{n\times p} : \mathcal{A}(YY^T) = b\}.$$

Let $\tau(k) := \binom{k+1}{2}$ be the k-th triangular number. It is known that problems (SDP) and (BM) have the same optimal value for any p such that $\tau(p) \geq m$; this is known as the *Barvinok-Pataki bound* [3, 29]. But due to nonconvexity, local optimization methods may converge to a critical point of (BM) which is not globally optimal [34]. In this paper we are mainly interested in 2nd-order critical points (abbreviated: 2-critical points).

Boumal et al. [8, 9] showed that (BM) has no spurious 2-critical points when $\tau(p) > m$, assuming that the feasible set $\mathfrak{M}$ is a smooth manifold and that the cost matrix C is *generic*. The result of Boumal et al. gives a strong indication that, for any p above the Barvinok-Pataki bound, a local optimization method for (BM) should lead to the global optimal of (SDP). However, there are two serious technical obstacles to derive polynomial time guarantees.

The first obstacle is that numerical algorithms must be terminated after finitely many iterations, and hence the criticality conditions are not satisfied exactly. This issue has been addressed by Pumir et al. [30], though for values of p larger than the Barvinok-Pataki bound. They introduce a smoothed analysis [32] setting in which the cost matrix is subject to a small random perturbation of magnitude σ. The purpose of this perturbation is to introduce genericity to the problem. They then defined a notion of *exactly feasible approximately 2-critical* (EFAC) point, and showed that any EFAC point of (BM) is also approximately optimal for (SDP) when $\tau(p) \gtrsim \frac{9}{2} m \log(\Omega(\sqrt{n}/\sigma))$ with high probability. Note that this bound gets worse when the perturbation magnitude σ decreases.

In this paper, we consider the same smoothed analysis setting (the cost C is perturbed), and we improve upon [30] by matching the Barvinok-Pataki bound. To do so, we provide a deterministic characterization of the spurious EFAC points of (BM) in terms of tubular neighborhoods around algebraic varieties. This deterministic characterization, given in Proposition 3.8, is of independent interest. By using effective bounds for the volume of such tubular neighborhoods [12, 22, 27], we derive the following theorem.

Theorem 1.1 (critical $\Rightarrow$ optimal | Informal)**.** *Consider this setting:*

- *The rank p satisfies $\tau(p) > m$.*
- *Apply a perturbation of magnitude σ to the cost matrix C.*

Then, with high probability, any EFAC point for (BM) with bounded norm is also approximately optimal for (SDP), provided that the criticality precision is sufficiently small. The precise statement appears in Theorem 3.5.

The second obstacle we need to overcome is the efficient computation of EFAC points. Feasibility is the main impediment. Indeed, the set $\mathfrak{M}$ is defined by quadratic equations, and solving quadratics is NP-hard. Hence, finding EFAC points is computationally intractable. To address this issue, we relax the feasibility requirement and define a notion of *approximately feasible approximately 2-critical* (AFAC) point. We show that Theorem 1.1 remains valid for AFAC points. Importantly, AFAC points can be computed in polynomial time (see Theorem 2.6), provided that an approximately feasible solution is known. This leads to the following theorem.

Theorem 1.2 (Polytime optimality | Informal)**.** *Consider this setting:*

- *The rank p satisfies $\tau(p) > (1+\eta)m$, for a fixed constant $\eta > 0$.*
- *Apply a perturbation of magnitude σ to the cost matrix C.*
- *Assume that $\mathfrak{M}$ is compact and smooth (LICQ holds).*
- *Let Y_0 be an approximately feasible point, i.e., Y_0 is close to $\mathfrak{M}$.*
- *Solve (BM) using a constrained optimization method with 2nd-order guarantees (e.g., Theorem 2.6) initialized at Y_0.*

Then, after $\mathrm{poly}(n, \sigma^{-1})$ *iterations, the algorithm produces with high probability a point Y such that YY^T is approximately optimal for (SDP).* The precise statement appears in Theorem 5.1.

The above theorem applies to SDPs with compact and smooth domains for which a feasible solution is known. Observe that it requires a slightly larger bound for p, compared to Theorem 1.1. This is needed to ensure that the criticality precision remains polynomially bounded. Also note that the perturbation magnitude σ appears in the complexity of the algorithm (as is usual in smooth analysis), but does not appear in the bound for p.

In order to fully address the computation of AFAC points, we need to find an approximately feasible solution of (SDP). To do so, we consider the least squares problem

$$(SDP_{ls}) \qquad \min_{X \in \mathbb{S}^n} \ \|\mathcal{A}(X) - b\|^2 \quad \text{such that} \quad X \succeq 0,$$

and the corresponding Burer-Monteiro problem

$$(BM_{ls}) \qquad \min_{Y \in \mathbb{R}^{n \times p}} \ \|\mathcal{A}(YY^T) - b\|^2.$$

We can analyze problem (BM_{ls}) using similar techniques as for (BM). We consider a smoothed analysis setting in which the constraint map $\mathcal{A}$ is subject to a small perturbation. We will show in Theorem 4.5 that, for $\tau(p) > m$, any approximately critical point of (BM_{ls}) is approximately optimal for (SDP_{ls}). This leads us to the following theorem.

Theorem 1.3 (Polytime feasibility | Informal)**.** *Consider this setting:*

- *The rank p satisfies $\tau(p) > (1+\eta)m$, for a fixed constant $\eta > 0$.*
- *Apply a perturbation of magnitude σ to the constraint map $\mathcal{A}$.*
- *Assume that the perturbed set $\mathfrak{M}$ is nonempty and compact.*

- *Solve* (BM_{ls}) *using an unconstrained optimization method with 2nd-order guarantees (e.g., [13, 14]).*

Then, after $\mathrm{poly}(n, \sigma^{-1})$ *iterations, the algorithm produces with high probability a point* Y *such that* YY^T *is approximately feasible for* (SDP). The precise statement appears in Theorem 5.3.

Theorems 1.2 and 1.3 together provide polynomial time guarantees for the Burer-Monteiro method, assuming that both $C, \mathcal{A}$ are slightly perturbed and that the domain is compact and smooth.

The structure of this paper is as follows. In Section 2 we introduce the notion of AFAC points in nonlinear programming, and we discuss how to compute them in polynomial time. In Section 3 we prove Theorem 1.1. In Section 4 we analyze the least squares problem (SDP_{ls}). In Section 5 we put together the results from the paper, and show Theorems 1.2 and 1.3. We conclude with some experimental results in Section 6.

Related work. The Burer-Monteiro method applied to problems (SDP) and (SDP_{ls}) has attracted much research in past years. Several papers have tried to explain the practical success of the method, as we elaborate now.

For problem (SDP), early work by Burer and Monteiro [11] and by Journée et al. [24] gave strong indications that (BM) has no spurious critical points above the Barvinok-Pataki bound. A formal proof was given recently by Boumal et al. [4, 8]. Subsequent work by Bhojanapalli et al. [4] and Pumir et al. [30] investigated the case of approximately critical points in a smoothed analysis setting. The work in [5] did not focus on (SDP), but in a penalized version. The Barvinok-Pataki bound was recently shown to be optimal up to lower order terms for general SDPs [34]. However for structured families of SDPs a smaller rank might suffice [2, 28]. As discussed above, previous work has not yet shown polynomial time guarantees for problem (BM).

Problem (SDP_{ls}) has been well studied in the matrix sensing community. Bhojanapalli et al. [5] showed that (BM_{ls}) has no spurious local minima under the restricted isometry property (RIP), and they also provided polynomial time guarantees. Similar results have been derived later, e.g., [21, 26]. Note that RIP is a very strong assumption about the condition number of the linear map $\mathcal{A}$, particularly since the RIP constant is small [35]. In contrast, our result in Theorem 1.3 simply assumes a small perturbation around a worst-case instance $\mathcal{A}$. This is a nondegeneracy condition, much weaker than RIP. To the best of our knowledge, this is the first paper that shows that (BM_{ls}) has no spurious local minima without relying on RIP.

2. Critical points in nonlinear programming

In this section we revisit the notion of critical points in unconstrained and constrained optimization. We also discuss notions of approximately critical points, and review methods with finite-time complexity guarantees.

2.1. **Unconstrained case.** Consider the optimization problem

$$(P_{\text{un}}) \qquad \min_{y \in \mathbb{R}^n} \; f(y),$$

with $f : \mathbb{R}^n \to \mathbb{R}$ twice continuously differentiable. A vector $y \in \mathbb{R}^n$ is a *2nd-order critical* point for (P_{un}), abbreviated 2-critical, if it satisfies:

$$\nabla f(y) = 0, \qquad \nabla^2 f(y) \succeq 0.$$

In the unconstrained case any local minimum of f is also a critical point.

Practical optimization algorithms cannot obtain a solution satisfying the above equations exactly. Hence, we consider a relaxation of these conditions.

Definition 2.1. Given $\varepsilon_1, \varepsilon_2 \in \mathbb{R}_+$, we say that y is $(\varepsilon_1, \varepsilon_2)$-*approximately 2-critical* (AC) for (P_{un}) if:

$$\|\nabla f(y)\| \leq \varepsilon_1, \qquad \nabla^2 f(y) \succeq -\varepsilon_2 I_n. \tag{1}$$

Several algorithms for unconstrained optimization with provable convergence guarantees are known. Recent work has focused on deriving algorithms with finite-time guarantees. In particular, the trust region method computes an $(\varepsilon_1, \varepsilon_2)$-AC point in $O(\max\{\varepsilon_1^{-2}\varepsilon_2^{-1}, \varepsilon_2^{-3}\})$ iterations [14], and the adaptive regularization with cubics (ARC) method takes $O(\max\{\varepsilon_1^{-2}, \varepsilon_2^{-3}\})$ iterations [13]. We formally state the result for the ARC method.

Theorem 2.2 ([13])**.** *Assume that there exists $\alpha > 0$ such that:*

- *A point y_0 with $f(y_0) \leq \alpha$ is known.*
- *$f, \nabla f, \nabla^2 f$ are uniformly bounded and Lipschitz continuous on the lower set $\{y : f(y) \leq \alpha\}$.*

The ARC method initialized at y_0 requires $O(\max\{\varepsilon_1^{-2}, \varepsilon_2^{-3}\})$ iterations to produce an $(\varepsilon_1, \varepsilon_2)$-AC point y. Furthermore, each iteration requires $O(1)$ evaluations of f and its derivatives.

2.2. **Constrained case.** Consider the nonlinear program

$$(P_{\text{con}}) \qquad \min_{y \in \mathfrak{M}} \; f(y), \quad \text{where} \quad \mathfrak{M} := \{y \in \mathbb{R}^m : h(y) = 0\},$$

with $f : \mathbb{R}^n \to \mathbb{R}$, $h : \mathbb{R}^n \to \mathbb{R}^m$ twice continuously differentiable. The Lagrangian function is $L(y, \lambda) = f(y) + \lambda \cdot h(y)$. A vector $y \in \mathbb{R}^n$ is a *2-critical point* if there are multipliers $\lambda \in \mathbb{R}^m$ such that

$$h(y) = 0, \qquad \nabla_y L(y, \lambda) = 0,$$
$$u^T \nabla^2_{yy} L(y, \lambda) u \geq 0, \quad \forall\, u \text{ such that } \nabla h(y) u = 0.$$

In the constrained case, for a local minima to be a critical point we need some regularity conditions. One such condition is the *linear independence constraint qualification* (LICQ), that states that $\nabla h(y)$ is full rank. This is equivalent to $\mathfrak{M}$ being smooth at y for the case of complete intersections (i.e., $\operatorname{codim} \mathfrak{M} = m$).

We now consider a relaxation of the criticality conditions.

Definition 2.3. Given $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$, $\gamma \in \mathbb{R}_+$, we say that y is $(\boldsymbol{\varepsilon}, \gamma)$-*approximately feasible approximately 2-critical* (AFAC) for (P_{con}) if there exists $\lambda \in \mathbb{R}^m$ such that:

$$\|h(y)\| \le \varepsilon_0, \qquad \|\nabla_y L(y,\lambda)\| \le \varepsilon_1, \tag{2a}$$

$$u^T \nabla^2_{yy} L(y,\lambda) u \ge -\varepsilon_2, \quad \forall\, u \text{ of unit norm such that } \|\nabla h(y) u\| \le \gamma. \tag{2b}$$

Our subsequent analysis requires a bound on the Lagrange multipliers λ. We can obtain such a bound through a quantitative version of LICQ.

Definition 2.4. For $\varrho > 0$, we say that ϱ-*LICQ* holds at y if the smallest singular value of $\nabla h(y)$ is at least ϱ.

Lemma 2.5. *Let (y,λ) be such that $\|\nabla_y L(y,\lambda)\| \le \varepsilon_1$. If ϱ-LICQ holds at y, then $\|\lambda\| \le \varrho^{-1}(\varepsilon_1 + \|\nabla f(y)\|)$.*

Proof. Let $J := \nabla h(y)$. Since $\nabla L(y,\lambda) = \nabla f(y) + J^T\lambda$, and J is full rank, then $\lambda = (J^\dagger)^T(\nabla L(y,\lambda) - \nabla f(y))$, where $J^\dagger$ is the pseudo-inverse of J. Hence $\|\lambda\| \le \varrho^{-1}(\varepsilon_1 + \|\nabla f(y)\|)$. □

Several local optimization methods for (P_{con}) with provable convergence guarantees are known. In particular, augmented Lagrangians [1] and trust-region methods [19, §15.4] converge to 2-critical points. Recent work has focused on finding algorithms with finite-time guarantees. The complexity of computing approximately 1-critical points was studied in, e.g., [6, 15, 16, 20].

As for approximately 2-critical points, we are only aware of [17, 31]. But both papers use a different 2nd-order condition, which is not easy to translate into our setting. Nonetheless, in Theorem 2.6 below we show that AFAC points can be computed in polynomial time. The proof of this theorem is in Appendix A, and relies on a variant of the method from [17]. To the best of our knowledge, this is the first polynomial time bound for computing 2-critical points that works with the standard notions of criticality.

Theorem 2.6. *Assume that there exist $\beta, \varrho \in \mathbb{R}_+$ such that:*

- *A point y_0 in the set $\mathfrak{M}_\beta := \{y : \|h(y)\| \le \beta\}$ is known.*
- *$f, \nabla f, \nabla^2 f$ and $\nabla h_i, \nabla^2 h_i$ for $i \in [m]$ are uniformly bounded and Lipschitz continuous on $\mathfrak{M}_\beta$.*
- *ϱ-LICQ holds at all $y \in \mathfrak{M}_\beta$.*

Let $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$, $\gamma \in \mathbb{R}_+$ be such that

$$\varepsilon_0 \le \beta, \quad \varepsilon_1 \le 1, \quad \varepsilon_1^2 \le \tfrac{1}{16} R_\lambda^{-1}\, \varepsilon_0\, \varepsilon_2, \quad \gamma^2 \le \tfrac{1}{16} R_\lambda^{-3}\, \varepsilon_0\, \varepsilon_2, \tag{3}$$

where $R_\lambda := 2 + 2\,\varrho^{-1} L_f$, $L_f := \max_y \|\nabla f(y)\|$. There is an algorithm that, when initialized at y_0, requires $O(\max\{\varepsilon_0^{-2}\,\varepsilon_1^{-2},\, \varepsilon_0^{-3}\,\varepsilon_2^{-3}\})$ evaluations of f, h and their derivatives to produce an $(\boldsymbol{\varepsilon}, \gamma)$-AFAC pair (y,λ), with $\|\lambda\| \le R_\lambda$.

Remark. Theorem 2.6 assumes that ϱ-LICQ holds everywhere in $\mathfrak{M}_\beta$. This is a strong assumption, that can only be guaranteed when β is very small. Hence the initial point y_0 must be approximately feasible.

3. Optimality of critical points

In this section we consider a smoothed analysis setting in which the constraint variables $\mathcal{A}, b$ are fixed, and the cost matrix C is subject to a small random perturbation. We will show that problem (BM) has no *spurious* approximately critical (AFAC) points with high probability. This means that any AFAC point of (BM) is approximately optimal for (SDP). We will restrict our attention to AFAC pairs (Y, λ) of bounded norm. Hence, we assume that $\|Y\| \leq R_Y$ and $\|\lambda\| \leq R_\lambda$ for some fixed constants R_Y, R_λ.

A crucial step toward our theorem is a geometric characterization of the spurious AFAC points in terms of tubes around algebraic varieties. We then take advantage of known effective bounds for the volume of such tubes [27].

From now on, we use the Frobenius norm for all matrices.

3.1. Spurious approximately critical points.

We proceed to describe the optimality conditions of problems (SDP) and (BM). Consider first the convex problem (SDP). Its conic dual is

$$\max_{\lambda \in \mathbb{R}^m} \; b^T \lambda \quad \text{such that} \quad S(\lambda) \succeq 0,$$

where

$$\begin{aligned} S(\lambda) \;&:=\; C - \mathcal{A}^*(\lambda) \;\in\; \mathbb{S}^n \qquad \text{is the } \textit{slack} \text{ matrix,} \\ \mathcal{A}^* : \mathbb{R}^m \to \mathbb{S}^n, \quad \lambda \mapsto &\textstyle\sum_i \lambda_i A_i \quad \text{is the adjoint of } \mathcal{A}. \end{aligned}$$

The optimality conditions for the dual pair of SDPs are:

$$\mathcal{A}(X) = b, \quad X \succeq 0, \quad S(\lambda) \succeq 0, \quad S(\lambda) X = 0.$$

The first three conditions correspond to primal and dual feasibility, and the last one is complementary slackness. We define an approximately optimal solution as a relaxation of the above conditions.

Definition 3.1. Let $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$. We say that $X \in \mathbb{S}^n$ is $\boldsymbol{\varepsilon}$-*approximately optimal* for (SDP) if there is $\lambda \in \mathbb{R}^m$ such that:

$$\|\mathcal{A}(X) - b\| \leq \varepsilon_0, \quad \|S(\lambda) X\| \leq \varepsilon_1, \quad X \succeq 0, \quad S(\lambda) \succeq -\varepsilon_2 \, I_n. \tag{4}$$

It is known that an $\boldsymbol{\varepsilon}$-approximately optimal solution is at distance $O(\|\boldsymbol{\varepsilon}\|)$ from an optimal solution under nondegeneracy assumptions [33]. We can also give a simple bound on the optimality gap.

Lemma 3.2. *If $(\bar{X}, \bar{\lambda})$ is $\boldsymbol{\varepsilon}$-approximately optimal for (SDP) then*

$$C \bullet \bar{X} \;\leq\; C \bullet X + \varepsilon_0 \|\bar{\lambda}\| + \varepsilon_1 \sqrt{n} + \varepsilon_2 \|X\| \sqrt{n} \qquad \forall \textit{ feasible } X.$$

Proof. The lemma follows from the following equations:

$$\begin{aligned} C \bullet X \;&=\; \bar{\lambda} \cdot \mathcal{A}(X) + S(\bar{\lambda}) \bullet X \;\geq\; \bar{\lambda} \cdot b - (\varepsilon_2 I_n) \bullet X \;=\; \bar{\lambda} \cdot b - \varepsilon_2 \|X\| \sqrt{n}, \\ \bar{\lambda} \cdot b \;&\geq\; \bar{\lambda} \cdot \mathcal{A}(\bar{X}) - \|\bar{\lambda}\| \, \|b - \mathcal{A}(\bar{X})\| \;\geq\; \bar{\lambda} \cdot \mathcal{A}(\bar{X}) - \varepsilon_0 \|\bar{\lambda}\|, \\ \bar{\lambda} \cdot \mathcal{A}(\bar{X}) \;&\geq\; \bar{\lambda} \cdot \mathcal{A}(\bar{X}) + S(\bar{\lambda}) \bullet X - \|S(\bar{\lambda}) X\|_* \;\geq\; C \bullet \bar{X} - \varepsilon_1 \sqrt{n}. \end{aligned}$$

□

We proceed to problem (BM). This is a special instance of the nonlinear program (P_{con}) with $f(Y) = C \bullet YY^T$ and $h(Y) = \mathcal{A}(YY^T) - b$. The Lagrangian function is $L(y, \lambda) = S(\lambda) \bullet YY^T + b^T\lambda$. The criticality conditions for (BM) are obtained by specializing (2).

Definition 3.3. Let $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$, $\gamma \in \mathbb{R}_+$. We say that $Y \in \mathbb{R}^{n\times p}$ is $(\boldsymbol{\varepsilon}, \gamma)$-*approximately feasible approximately 2-critical* (AFAC) for (BM) if there is $\lambda \in \mathbb{R}^m$ such that:

$$\|\mathcal{A}(YY^T) - b\| \leq \varepsilon_0, \qquad \|S(\lambda)Y\| \leq \varepsilon_1, \tag{5a}$$

$$S(\lambda) \bullet UU^T \geq -\varepsilon_2, \quad \forall\, U \in \mathbb{R}^{n\times p} \text{ with } \|U\| = 1,\ \|\mathcal{A}(UY^T)\| \leq \gamma. \tag{5b}$$

Remark. By fixing $\varepsilon_0 = 0$ we get a notion of exactly feasible approximately 2-critical (EFAC) point, which was used in [4, 30].

We are ready to formalize the concept of spurious critical points.

Definition 3.4. Let $R_Y, R_\lambda \in \mathbb{R}_+$ be fixed constants. For arbitrary $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$, $\gamma \in \mathbb{R}_+$, we say that a pair (Y, λ) is *spurious* $(\boldsymbol{\varepsilon}, \gamma)$-AFAC if:

- (Y, λ) is $(\boldsymbol{\varepsilon}, \gamma)$-AFAC for (BM).
- (YY^T, λ) is not $\boldsymbol{\varepsilon}'$-approx. optimal for (SDP), with $\boldsymbol{\varepsilon}' := (\varepsilon_0, R_Y\varepsilon_1, \varepsilon_2)$.
- $\|Y\| \leq R_Y$ and $\|\lambda\| \leq R_\lambda$.

A pair (Y, λ) is spurious *exactly* critical if the above holds for $\boldsymbol{\varepsilon} = 0, \gamma = 0$.

3.2. **Statement of the theorem.** We present the main result of this section. Let $\mathcal{A}, b$ be fixed, and let C be obtained from a random perturbation of magnitude σ around some fixed $\bar{C}$. Concretely, C is *uniformly* distributed on the Frobenius ball $\mathbf{B}_\sigma(\bar{C}) \subset \mathbb{S}^n$ of radius σ centered at $\bar{C}$. Consider the set $\mathscr{C}_{\boldsymbol{\varepsilon},\gamma} \subset \mathbb{S}^n$, consisting of all cost matrices for which there is a spurious AFAC point:

$$\mathscr{C}_{\boldsymbol{\varepsilon},\gamma} \;:=\; \{C \in \mathbb{S}^n \,:\, \exists\, (Y, \lambda) \text{ an spurious } (\boldsymbol{\varepsilon}, \gamma)\text{-AFAC pair}\}\,.$$

We show that if $\tau(p) > m$, then the matrix C avoids the "bad set" $\mathscr{C}_{\boldsymbol{\varepsilon},\gamma}$ with high probability. More precisely, the probability $\Pr[C \in \mathscr{C}_{\boldsymbol{\varepsilon},\gamma}]$ is vanishingly small as the ratio ε_1/γ goes to zero.

Theorem 3.5 (critical $\Rightarrow$ optimal)**.** *Let p such that $\tau(p) > m$. Let $\boldsymbol{\varepsilon} \in \mathbb{R}^3_+$, $\gamma \in \mathbb{R}_+$. Let C be uniformly distributed on the Frobenius ball $\mathbf{B}_\sigma(\bar{C})$. Then*

$$\Pr[C \in \mathscr{C}_{\boldsymbol{\varepsilon},\gamma}] \;\leq\; 8e\,\delta^{\tau(p)-m}\,(3\kappa)^m\,(2en^3/\sigma)^{\tau(p)},$$

where $\delta := \varepsilon_1\|\mathcal{A}\|/\gamma$ and $\kappa := R_\lambda\|\mathcal{A}\|$, provided that $\delta < \sigma/4en^3$.

The following corollary shows that when the stronger condition $\tau(p) > (1+\eta)m$ holds, where η is a fixed constant, then we can derive a high probability bound while maintaining δ polynomially bounded.

Corollary 3.6. *Consider the setup from Theorem 3.5. Assume that*

$$\tau(p) \geq (1+\eta)m + \eta\, t, \qquad \delta \leq (1/3\kappa)^{1/\eta}(\sigma/2en^3)^{1+1/\eta},$$

for some arbitrary constants $\eta, t > 0$. Then

$$\Pr[C \in \mathscr{C}_{\boldsymbol{\varepsilon},\gamma}] \;\leq\; 8e\,(\sigma/6\kappa en^3)^t.$$

Proof. It is a straightforward manipulation. □

3.3. **Tubes around varieties.** Our proof of Theorem 3.5 relies on a geometric characterization of the set $\mathscr{C}_{\boldsymbol{\varepsilon},\gamma}$. Such characterization is known for the case $\boldsymbol{\varepsilon} = 0, \gamma = 0$, corresponding to *exactly* critical points. It was shown in [9], see also [18], that the existence of a spurious exactly critical point implies that C lies in a certain algebraic variety of $\mathbb{S}^n$, as follows:

$$\exists \text{ (spurious exactly critical point)} \qquad \Longrightarrow \qquad C \,\in\, \mathbb{S}^n_{n-p} + \operatorname{Im}\mathcal{A}^*,$$

where $\mathbb{S}^n_{n-p} := \{X : \operatorname{rank} X \leq n-p\}$ is a variety of low rank matrices, and $\operatorname{Im}\mathcal{A}^*$ is the linear space spanned by $A_1, \ldots, A_m$. Hence, we have that

$$\mathscr{C}_{0,0} \;\subset\; \mathbb{S}^n_{n-p} + \operatorname{Im}\mathcal{A}^*.$$

When $\tau(p) > m$, the variety $\mathbb{S}^n_{n-p} + \operatorname{Im}\mathcal{A}^*$ is properly contained in $\mathbb{S}^n$. It follows that $\mathscr{C}_{0,0}$ has measure zero and hence, for generic C, there are no spurious exactly critical points.

We show below that for approximately critical points the situation is analogous, except that we need to consider a *tubular* neighborhood around the variety $\mathbb{S}^n_{n-p} + \operatorname{Im}\mathcal{A}^*$.

Definition 3.7. For a set $\mathcal{W} \subset \mathbb{S}^n$ and a positive number $\delta \in \mathbb{R}_+$, we define

$$\operatorname{tube}_\delta \mathcal{W} \;:=\; \{X \in \mathbb{S}^n : \|X - W\| \leq \delta \text{ for some } W \in \mathcal{W}\}.$$

Proposition 3.8. *Let $\delta := \varepsilon_1\|\mathcal{A}\|/\gamma$ and $B_\lambda := \{\lambda \in \mathbb{R}^m : \|\lambda\| \leq R_\lambda\}$. Then*

$$\mathscr{C}_{\boldsymbol{\varepsilon},\gamma} \;\subset\; \operatorname{tube}_\delta(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(B_\lambda) \;=\; \operatorname{tube}_\delta(\mathbb{S}^n_{n-p} + \mathcal{A}^*(B_\lambda)).$$

The next lemma is the key ingredient for Proposition 3.8. It shows that if Y is a spurious AFAC point, then its smallest singular value $\sigma_p(Y)$ is bounded from below. This is a generalization of a previous result by Burer and Monteiro [11] about exactly critical points. An analogue of this lemma is found in [30, Lem.3.2] for the case of EFAC points.

Lemma 3.9. *Let Y be an $(\boldsymbol{\varepsilon}, \gamma)$-AFAC point of (BM). If $\sigma_p(Y) \leq \gamma/\|\mathcal{A}\|$, then YY^T is $\boldsymbol{\varepsilon}'$-approximately optimal for (SDP), with $\boldsymbol{\varepsilon}' := (\varepsilon_0, R_Y\varepsilon_1, \varepsilon_2)$.*

Proof. Let (Y, λ) satisfy (5), and let us show that (YY^T, λ) satisfies (4). The first three conditions in (4) are easy to check. We proceed to show the last one: $S(\lambda) \succeq -\varepsilon_2 I_n$. Given a unit vector $u \in \mathbb{R}^n$, we need to show that $u^T S(\lambda) u \geq -\varepsilon_2$. There is a unit vector $z \in \mathbb{R}^p$ such that $\|Yz\| = \sigma_p(Y)$. The matrix $U := uz^T$ satisfies $\|U\| = 1$ and

$$\|UY^T\| \leq \|u\|\|Yz\| = \sigma_p(Y) \leq \gamma/\|\mathcal{A}\|.$$

Then $\|\mathcal{A}(UY^T)\| \leq \gamma$, so by (5b) we have

$$-\varepsilon_2 \;\leq\; S(\lambda) \bullet UU^T \;=\; \|z\|^2 (u^T S(\lambda) u) \;=\; u^T S(\lambda) u. \qquad \square$$

Proof of Proposition 3.8. Let $C \in \mathscr{C}_{\boldsymbol{\varepsilon},\gamma}$, so there is a spurious $(\boldsymbol{\varepsilon},\gamma)$-AFAC pair (Y,λ). By Lemma 3.9, we must have $\sigma_p(Y) > \gamma/\|\mathcal{A}\|$. Let $S := S(\lambda)$, and note that $\|SY\| \leq \varepsilon_1 = \gamma\,\delta/\|\mathcal{A}\|$ by (5a). Then

$$\operatorname{dist}(S,\, \mathbb{S}^n_{n-p}) \;=\; \sqrt{\textstyle\sum_{i=1}^{p} \sigma^2_{n-i+1}(S)} \;\leq\; \|SY\|/\sigma_p(Y) \;<\; \delta, \tag{6}$$

and hence $C \;=\; S(\lambda) + \mathcal{A}^*(\lambda) \;\in\; \operatorname{tube}_\delta(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(B_\lambda)$. $\square$

By Proposition 3.8, the set of "bad" cost matrices $\mathscr{C}_{\boldsymbol{\varepsilon},\gamma}$ is contained in a tube around a variety. To prove Theorem 3.5, we need an upper bound on the probability mass of this tube. The computation of integrals over tubes has a long history in differential geometry [22]. Effective bounds for these integrals were shown in [12, 27]; they were used for smooth analysis in the first reference. The following bound follows from [27].

Theorem 3.10. *Let $V \subset \mathbb{R}^k$ be a real variety of codimension c defined by polynomials of degree at most D. Let x be uniformly distributed on the Euclidean ball $\mathbf{B}_\sigma(\bar{x}) \subset \mathbb{R}^k$. Then*

$$\Pr[x \in \operatorname{tube}_\delta(V)] \;\leq\; 4 \sum_{i=c}^{k} \binom{k}{i} \left(\frac{2D\delta}{\sigma}\right)^i \left(1 + \frac{\delta}{\sigma}\right)^{k-i} \;\leq\; 8e \left(\frac{2ekD\delta}{c\sigma}\right)^c, \tag{7}$$

where in the second inequality we require that $\max\{4ekD\delta/c\sigma,\, k\delta/\sigma\} \leq 1$.

Proof. We start with the left inequality. This was shown in [27] in the case that V is a complete intersection. Assume now that V is not a complete intersection. Then $V = V(f)$ with $f \in \mathbb{R}[x]^m$ and $m > c$. For a generic matrix $A \in \mathbb{R}^{c \times m}$ we have that $f' = Af \in \mathbb{R}[x]^c$ is a complete intersection by Bertini's theorem [25]. Since $V \subset V(f')$, the result follows from the previous case.

We proceed to the right inequality. Since $k \leq \sigma/\delta$ then $(1+\delta/\sigma)^k \leq e$. Denoting $y := 2ekD\delta/c\sigma$, the summation in (7) is less than

$$4e \sum_{i=c}^{k} \binom{k}{i} \left(\frac{2D\delta}{\sigma}\right)^i \;\leq\; 4e \sum_{i=c}^{k} \left(\frac{ke}{i} \cdot \frac{2D\delta}{\sigma}\right)^i \;\leq\; 4e \sum_{i=c}^{\infty} y^i \;=\; \frac{4ey^c}{1-y} \;\leq\; 8ey^c,$$

where we used that $y \leq 1/2$ by assumption. $\square$

Proof of Theorem 3.5. Let $B_\lambda \subset \mathbb{R}^m$ be the of radius R_λ centered at zero. Consider an ϵ-net $\mathcal{N}$ of B_λ, where $\epsilon := \delta/\|\mathcal{A}\|$. It is known that $(3R_\lambda/\epsilon)^m = (3\kappa/\delta)^m$ points suffices for the ϵ-net. Observe that

$$\mathcal{A}^*(B_\lambda) \;\subset\; \mathcal{A}^*(\operatorname{tube}_\epsilon(\mathcal{N})) \;\subset\; \operatorname{tube}_\delta(\mathcal{A}^*(\mathcal{N})), \tag{8a}$$

$$\mathscr{C}_{\boldsymbol{\varepsilon},\gamma} \;\subset\; \operatorname{tube}_\delta(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(B_\lambda) \;\subset\; \operatorname{tube}_{2\delta}(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(\mathcal{N}). \tag{8b}$$

Recall that $\mathbb{S}^n_{n-p}$ is a variety of codimension $\tau(p)$ defined by equations of degree $n-p+1$. For any $\ell \in \mathcal{N}$, Theorem 3.10 gives

$$\Pr[C-\mathcal{A}^*(\ell) \in \mathrm{tube}_{2\delta}(\mathbb{S}^n_{n-p})] \;<\; 8e\left(\frac{4e\tau(n)(n-p+1)\delta}{\tau(p)\sigma}\right)^{\tau(p)} \;<\; 8e\left(\frac{2en^3\delta}{\sigma}\right)^{\tau(p)}.$$

Finally, the union bound gives

$$\begin{aligned}\Pr[C \in \mathscr{C}_{\boldsymbol{\varepsilon},\gamma}] \;&\leq\; \#\mathcal{N} \cdot \Pr[C \in \mathrm{tube}_{2\delta}(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(\ell)] \\ &<\; (3\kappa/\delta)^m \cdot 8e\left(2en^3\delta/\sigma\right)^{\tau(p)}. \qquad \square\end{aligned}$$

4. Feasibility of critical points

In this section we address the computation of a feasible solution for (*SDP*). To do so, we consider the least squares problem (SDP_{ls}) and its Burer-Monteiro relaxation (BM_{ls}). Note that if (*SDP*) is feasible, then the optimal value of (SDP_{ls}) is zero. Nevertheless, the results of this section apply to an arbitrary instance of (SDP_{ls}), even if the optimal value is nonzero.

We consider a smoothed analysis setting in which b is fixed, and $\mathcal{A}$ is subject to a small random perturbation. We will show that problem (BM_{ls}) has no *spurious* approximately critical (AC) points with high probability. This means that any AC point of (BM_{ls}) is approximately optimal for (SDP_{ls}). We will restrict our attention to AC points Y of bounded norm. Hence, we assume that $\|Y\| \leq R_Y$ for some fixed constant R_Y.

4.1. Spurious approximately critical points.

We proceed to describe the optimality conditions for (SDP_{ls}) and (BM_{ls}). Let $f(X) := \|\mathcal{A}(X)-b\|^2$ be the least squares objective function. The optimality conditions for the convex problem (SDP_{ls}) are:

$$S(X)\,X = 0, \qquad X \succeq 0, \qquad S(X) \succeq 0,$$

where $S(X)$ is the gradient of the objective function

$$S(X) \;:=\; \nabla f(X) \;=\; 2\mathcal{A}^*(\mathcal{A}(X)-b) \;\in\; \mathbb{S}^n.$$

We call X approximately optimal if either $f(X)$ is very close to zero, or the above conditions are approximately satisfied.

Definition 4.1. Let $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$. We say that $X \in \mathbb{S}^n$ is $\boldsymbol{\varepsilon}$-*approximately optimal* for (SDP_{ls}) if $X \succeq 0$ and either

$$\|\mathcal{A}(X)-b\| \leq \varepsilon_0 \quad \text{or} \quad \left(\,\|S(X)X\| \leq \varepsilon_1 \;\text{ and }\; S(X) \succeq -\varepsilon_2\, I_n\,\right). \tag{9}$$

The following lemma bounds the optimality gap for the second case in (9).

Lemma 4.2. *Let $\bar{X} \in \mathbb{S}^n_+$ such that $\|S(\bar{X})\bar{X}\| \leq \varepsilon_1$, $S(\bar{X}) \succeq -\varepsilon_2 I_n$. Then*

$$f(\bar{X}) \;\leq\; f(X) + \varepsilon_1\sqrt{n} + \varepsilon_2\|X\|\sqrt{n} \qquad \forall\, X \in \mathbb{S}^n_+.$$

Proof. Let $L(X) := f(X) - \bar{S} \bullet X$, with $\bar{S} := S(\bar{X})$. This is a convex function with $\nabla L(\bar{X}) = 0$, so $\bar{X}$ is its global minimum. Note that

$$f(X) \;=\; L(X) + \bar{S} \bullet X \;\geq\; L(X) - (\varepsilon_2 I_n) \bullet X \;=\; L(X) - \varepsilon_2 \|X\| \sqrt{n},$$
$$L(\bar{X}) \;=\; f(\bar{X}) - \bar{S} \bullet \bar{X} \;\geq\; f(\bar{X}) - \|\bar{S}\bar{X}\|_* \;\geq\; f(\bar{X}) - \varepsilon_1 \sqrt{n}.$$

Since $L(X) \geq L(\bar{X})$, the result follows from the above equations. □

We proceed now to the Burer-Monteiro problem (BM_{ls}). This is a special instance of the unconstrained optimization problem (P_{un}). The criticality conditions for (BM_{ls}) are obtained by specializing (1).

Definition 4.3. Let $(\varepsilon_1, \varepsilon_2) \in \mathbb{R}^2_+$. We say that $Y \in \mathbb{R}^{n \times p}$ is $(\varepsilon_1, \varepsilon_2)$-*approximately 2-critical* (AC) for (BM_{ls}) if

$$\|S(YY^T)Y\| \leq \varepsilon_1, \tag{10a}$$
$$S(YY^T) \bullet UU^T + 4\|\mathcal{A}(UY^T)\|^2 \geq -\varepsilon_2, \quad \forall\, U \in \mathbb{R}^{n \times p},\ \|U\| = 1 \tag{10b}$$

We are ready to formalize the concept of spurious critical points.

Definition 4.4. Let R_Y be fixed. For an arbitrary $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \varepsilon_2) \in \mathbb{R}^3_+$, we say that a point Y is *spurious* $\boldsymbol{\varepsilon}$-AC if:

- Y is $(\varepsilon_1, \varepsilon_2)$-AC for ($BM_{ls}$).
- YY^T is not $\boldsymbol{\varepsilon}'$-approx. optimal for ($SDP_{ls}$), with $\boldsymbol{\varepsilon}' := (\varepsilon_0, R_Y \varepsilon_1, 5\varepsilon_2)$.
- $\|Y\| \leq R_Y$.

A pair (Y, λ) is spurious *exactly* critical if the above holds for $\boldsymbol{\varepsilon} = 0$.

4.2. **Statement of the theorem.** We assume here that $b \in \mathbb{R}^m$ is fixed, and we vary the linear map $\mathcal{A}$. We identify $\mathcal{A}$ with the tuple of matrices $(A_1, \ldots, A_m)$, and hence view it as an element of the Euclidean space $(\mathbb{S}^n)^m$. We assume that $\mathcal{A}$ is *uniformly* distributed on the Euclidean ball $\mathbf{B}_\sigma(\bar{\mathcal{A}}) \subset (\mathbb{S}^n)^m$ of radius σ centered at $\bar{\mathcal{A}}$. Consider the set $\mathscr{A}_{\boldsymbol{\varepsilon}} \subset (\mathbb{S}^n)^m$, consisting of all $\mathcal{A}$ for which there is a spurious AC point:

$$\mathscr{A}_{\boldsymbol{\varepsilon}} \;:=\; \{\mathcal{A} \in (\mathbb{S}^n)^m \,:\, \exists\, Y \text{ a spurious } \boldsymbol{\varepsilon}\text{-AC point}\}\,.$$

We show that if $\tau(p) > m$, then the probability $\Pr[\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}]$ is vanishingly small as the ratio $\varepsilon_1/\sqrt{\varepsilon_2}$ goes to zero.

Theorem 4.5 (critical ⇒ feasible)**.** *Let p such that $\tau(p) > m$, and $\boldsymbol{\varepsilon} \in \mathbb{R}^3_+$. Let $\mathcal{A}$ be uniformly distributed on the Euclidean ball $\mathbf{B}_\sigma(\bar{\mathcal{A}})$. Then*

$$\Pr[\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}] \;\leq\; 8e\, \delta^{\tau(p)-m}\, (3\kappa)^m\, (en^3/\sigma\varepsilon_0)^{\tau(p)},$$

where $\delta := \varepsilon_1 R_A/\sqrt{\varepsilon_2}$, $\kappa := 2(R_A R_Y^2 + \|b\|)R_A$, and $R_A := \|\bar{\mathcal{A}}\| + \sigma$, provided that $\delta < \sigma\varepsilon_0/4en^3$.

As before, we can derive a high probability bound with δ polynomially bounded when $\tau(p) > (1+\eta)m$.

Corollary 4.6. *Consider the setup from Theorem 4.5. Assume that*

$$\tau(p) \geq (1+\eta)m + \eta\, t, \quad \delta \leq (1/3\kappa)^{1/\eta}(\rho\sigma^2/e^2n^6)^{1+1/\eta}, \quad \varepsilon_0 \geq \rho\,\sigma,$$

for some arbitrary constants $\eta, t, \rho > 0$. Then

$$\Pr[\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}] \;\leq\; 8e\,(\rho\sigma^2/3\kappa e^2n^6)^t.$$

4.3. **Tubes around varieties.** As in Section 3, our proof of Theorem 3.5 relies on a geometric characterization of the spurious AC points. First consider the simpler case of spurious *exactly* critical points ($\boldsymbol{\varepsilon} = 0$). The following equation is a consequence of our analysis:

$$\exists \text{ (spurious exactly critical point)} \quad \Longrightarrow \quad \mathbb{S}^n_{n-p} \cap \mathrm{Im}\mathcal{A}^* \text{ is nontrivial.}$$

This implies that the set $\mathscr{A}_0$ has measure zero when $\tau(p) > m$. Hence, for generic $\mathcal{A}$, there are no spurious exactly critical points.

The case $\boldsymbol{\varepsilon} > 0$ is similar, but we need consider a tube around $\mathbb{S}^n_{n-p}$.

Proposition 4.7. *Let $\delta := \varepsilon_1 R_A/\sqrt{\varepsilon_2}$, $D_\lambda := \{\lambda \in \mathbb{R}^m : 2\varepsilon_0 \leq \|\lambda\| \leq R_\lambda\}$, $R_\lambda := 2(R_A R_Y^2 + \|b\|)$. Then for $\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}$ we have that*

$$\mathrm{tube}_\delta(\mathbb{S}^n_{n-p}) \cap \mathcal{A}^*(D_\lambda) \neq \emptyset.$$

The above proposition relies on an analogue of Lemma 3.9.

Lemma 4.8. *Let Y be an $(\varepsilon_1, \varepsilon_2)$-AC point of (BM_{ls}). If $\sigma_p(Y) \leq \sqrt{\varepsilon_2}/R_A$, then YY^T is $\boldsymbol{\varepsilon}'$-approximately optimal for (SDP_{ls}), with $\boldsymbol{\varepsilon}' := (0, R_Y\varepsilon_1, 5\varepsilon_2)$.*

Proof. Let Y satisfy (10), and let us show that YY^T satisfies (9). The first-order condition is easy to check. We proceed to show that $u^T S(X) u \geq -\varepsilon_2'$ for any unit vector $u \in \mathbb{R}^n$. Let $z \in \mathbb{R}^p$ be a unit vector such that $\|Yz\| = \sigma_p(Y)$. The matrix $U := uz^T$ satisfies $\|U\| = 1$ and $\|UY^T\| \leq \|u\|\|Yz\| = \sigma_p(Y)$. Then $\|\mathcal{A}(UY^T)\| \leq \sqrt{\varepsilon_2}$ and by (10b) we have

$$u^T S(X) u = S(YY^T) \bullet UU^T \geq -\varepsilon_2 - 4\|\mathcal{A}(UY^T)\|^2 \geq -5\varepsilon_2. \qquad \square$$

Proof of Proposition 4.7. As $\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}$, there is a spurious $\boldsymbol{\varepsilon}$-AC point Y. By Lemma 3.9, we must have $\sigma_p(Y) > \sqrt{\varepsilon_2}/R_A$. Note that $\|S(YY^T)Y\| \leq \varepsilon_1$. Together with (6), we conclude that $S(YY^T) \in \mathrm{tube}_\delta(\mathbb{S}^n_{n-p})$. Let $\lambda := 2(\mathcal{A}(YY^T) - b)$. Note that $\|\lambda\| > 2\varepsilon_0$ and

$$\|\lambda\| = 2\,\|\mathcal{A}(YY^T) - b\| \leq 2(\|\mathcal{A}\|\|Y\|^2 + \|b\|) \leq R_\lambda.$$

Then $\lambda \in D_\lambda$ and $\mathcal{A}^*(\lambda) = S(YY^T) \in \mathrm{tube}_\delta(\mathbb{S}^n_{n-p})$. $\square$

Proof of Theorem 4.5. The result in Proposition 4.7 can be expressed as:

$$\mathscr{A}_{\boldsymbol{\varepsilon}} \;\subset\; \{\mathcal{A} \in (\mathbb{S}^n)^m : 0 \in \mathrm{tube}_\delta(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(D_\lambda)\},$$

which is closer to the formula in Proposition 3.8. Consider an ϵ-net $\mathcal{N}$ of D_λ, where $\epsilon := \delta/R_A$. It suffices to take $(3R_\lambda/\epsilon)^m = (3\kappa/\delta)^m$ points for

the ϵ-net. A reasoning similar to (8) gives

$$\begin{aligned}\mathscr{A}_{\boldsymbol{\varepsilon}} \;&\subset\; \{\mathcal{A} \in (\mathbb{S}^n)^m : 0 \in \operatorname{tube}_{2\delta}(\mathbb{S}^n_{n-p}) + \mathcal{A}^*(\mathcal{N}))\} \\ &=\; \bigcup_{\ell \in \mathcal{N}} \{\mathcal{A} \in (\mathbb{S}^n)^m : \mathcal{A}^*(\ell) \in \operatorname{tube}_{2\delta}(\mathbb{S}^n_{n-p}))\}.\end{aligned}$$

Let $\ell \in \mathcal{N}$, and consider the linear map

$$\phi_\ell : (\mathbb{S}^n)^m \to \mathbb{S}^n, \quad \mathcal{A} \mapsto \mathcal{A}^*(\ell).$$

This is a surjective map. Moreover, the scaled map $\frac{1}{\|\ell\|}\,\phi_\ell$ gives an isometry $(\ker \phi_\ell)^\perp \cong \mathbb{S}^n$. It follows that

$$\phi_\ell(\mathcal{A}) \;\in\; \operatorname{tube}_{2\delta}(\mathbb{S}^n_{n-p}) \quad\iff\quad \mathcal{A} \;\in\; \operatorname{tube}_{2\delta/\|\ell\|}\left(\phi_\ell^{-1}(\mathbb{S}^n_{n-p})\right).$$

Since $\|\ell\| \geq 2\varepsilon_0$, we conclude that

$$\mathscr{A}_{\boldsymbol{\varepsilon}} \;\subset\; \bigcup_{\ell \in \mathcal{N}} \operatorname{tube}_{\delta/\varepsilon_0}(V_\ell), \quad \text{with } \; V_\ell := \phi_\ell^{-1}(\mathbb{S}^n_{n-p}).$$

The final part of the proof is similar to the one in Theorem 3.5. The variety V_ℓ is a cylinder over $\mathbb{S}^n_{n-p}$, so it has the same codimension $\tau(p)$ and degree $n-p+1$ as $\mathbb{S}^n_{n-p}$. Using the union bound and Theorem 3.10, we get

$$\begin{aligned}\Pr[\mathcal{A} \in \mathscr{A}_{\boldsymbol{\varepsilon}}] \;&<\; \#\mathcal{N} \cdot \Pr[\mathcal{A} \in \operatorname{tube}_{\delta/\varepsilon_0}(V_\ell)] \\ &<\; (3\kappa/\delta)^m \cdot 8e\left(en^3\delta/\sigma\varepsilon_0\right)^{\tau(p)}.\end{aligned}$$

□

5. Overall complexity estimates

In this section we will derive polynomial time guarantees for the Burer-Monteiro method. In particular, we will formalize and prove Theorems 1.2 and 1.3 from the introduction.

We introduce some notation that will be used throughout this section. We consider two constants $\alpha \geq \beta > 0$ and associated sets $\mathfrak{M}_\alpha \supset \mathfrak{M}_\beta$, where $\mathfrak{M}_t := \{Y : \|\mathcal{A}(YY^T) - b\| \leq t\}$. We assume that β is small enough so that ρ-LICQ holds globally on $\mathfrak{M}_\beta$. On the other hand, $\alpha > 0$ is sufficiently large so that a point $Y_0 \in \mathfrak{M}_\alpha$ is always known. We further assume that $\mathfrak{M}_\alpha$ is compact. This compactness assumption is satisfied, for instance, if one of the constraints matrices A_i is positive semidefinite.

5.1. **Optimality.** We assume first that an approximately feasible solution Y_0 is known. Consider the following setting:

- p satisfies $\tau(p) \geq (1+\eta)m + \eta\, t$ for some given constants $\eta, t \in \mathbb{R}_+$.
- $\mathcal{A}, b$ are fixed and C is uniformly distributed on a ball $\mathbf{B}_\sigma(\bar{C})$.
- $\exists \beta \in \mathbb{R}_+$ such that: $\mathfrak{M}_\beta$ is compact, a point $Y_0 \in \mathfrak{M}_\beta$ is known, and ϱ-LICQ holds on $\mathfrak{M}_\beta$.
- $R_Y, L_f \in \mathbb{R}_+$ are constants that bound $\|Y\|$ and $\|CY\|$, for $Y \in \mathfrak{M}_\beta$.
- (BM) is solved with the method from Theorem 2.6 initialized at Y_0.

The next theorem shows that the Burer-Monteiro method solves (SDP) in polynomial time with high probability.

Theorem 5.1 (Polytime optimality)**.** *Let $\rho \in (0,1]$ arbitrary, and let*

$$\varepsilon_0 := \gamma := \epsilon, \quad \varepsilon_1 := \epsilon^2, \quad \varepsilon_2 := 16\, R_\lambda^3\, \epsilon, \quad \textit{with} \quad \epsilon := K^{-1}\, \rho\, (\sigma/2en^3)^{1+1/\eta},$$

where R_λ and K are the problem dependent constants

$$R_\lambda := 2 + 2\varrho^{-1} L_f, \quad K := \|\mathcal{A}\|\, (3\kappa)^{1/\eta}, \qquad \textit{with} \qquad \kappa := R_\lambda \|\mathcal{A}\|.$$

The algorithm from Theorem 2.6 returns a pair (Y, λ) after $O(\epsilon^{-6})$ function evaluations. With probability at least $1 - O(\sigma/n^3)^t$, the pair (YY^T, λ) is $(\epsilon, \epsilon^2 R_Y, 16 R_\lambda^3 \epsilon)$-approximately optimal for (SDP).

Proof. The smoothness assumptions in Theorem 2.6 are satisfied since $\mathfrak{M}_\beta$ is compact. Then (Y, λ) is an $(\boldsymbol{\varepsilon}, \gamma)$-AFAC pair with $\|\lambda\| \leq R_\lambda$. Note that

$$\delta \;:=\; \varepsilon_1 \|\mathcal{A}\|/\gamma \;=\; \epsilon\, \|\mathcal{A}\| \;\leq\; (1/3\kappa)^{1/\eta}\, (\sigma/2en^3)^{1+1/\eta}$$

is as in Corollary 3.6. Hence (YY^T, λ) is $(\varepsilon_0, \varepsilon_1 R_Y, \varepsilon_2)$-approximately optimal for (*SDP*) with probability $1 - O(\sigma/n^3)^t$. □

The above theorem shows that YY^T obtained is approximately optimal for the perturbed problem (*SDP*) with high probability. Let ($\overline{SDP}$) denote the SDP problem in which we use the unperturbed cost matrix $\bar{C}$. We can also show that YY^T is also approximately optimal for ($\overline{SDP}$).

Corollary 5.2. *Consider the setup of Theorem 5.1. With probability at least $1 - O(\sigma/n^3)^t$, the pair (YY^T, λ) is $(\varepsilon_0'', \varepsilon_1'', \varepsilon_2'')$-approximately optimal for ($\overline{SDP}$), where $\varepsilon_0'', \varepsilon_1'', \varepsilon_2'' = O(\sigma)$.*

Proof. Let $X := YY^T$. We know that (X, λ) is $(\varepsilon_0', \varepsilon_1', \varepsilon_2')$-approximately optimal for (*SDP*) with high probability. Let $S := C - \mathcal{A}^*(\lambda)$, $\bar{S} := \bar{C} - \mathcal{A}^*(\lambda)$ be the slack matrices for (*SDP*) and ($\overline{SDP}$). Observe that

$$\|\mathcal{A}(X) - b\| \;\leq\; \varepsilon_0' \;\leq\; O(\sigma), \tag{11}$$

$$\|\bar{S}X\| \;\leq\; \|SX\| + \|(\bar{S} - S)X\| \;\leq\; \varepsilon_1' + \sigma\|X\| \;\leq\; O(\sigma), \tag{12}$$

$$\bar{S} \;\succeq\; S - \|\bar{S} - S\| I_n \succeq -(\varepsilon_2' + \sigma) I_n \succeq -O(\sigma) I_n. \tag{13}$$

So the optimality conditions of ($\overline{SDP}$) hold with $\varepsilon_0'', \varepsilon_1'', \varepsilon_2'' = O(\sigma)$. □

5.2. **Feasibility.** We now consider the computation of a feasible point. Consider the following setting:

- p satisfies $\tau(p) \geq (1+\eta)m + \eta\, t$ for some given constants $\eta, t \in \mathbb{R}_+$.
- b is fixed and $\mathcal{A}$ is uniformly distributed on a ball $\mathbf{B}_\sigma(\bar{\mathcal{A}})$.
- $\exists \alpha \in \mathbb{R}_+$ and a matrix Y_0 such that $\mathfrak{M}_\alpha$ is compact and $Y_0 \in \mathfrak{M}_\alpha$
- $R_Y \in \mathbb{R}_+$ is a constant that bounds $\|Y\|$, for $Y \in \mathfrak{M}_\alpha$.
- (BM_{ls}) is solved with the method from Theorem 2.2 initialized at Y_0.

Theorem 5.3 (Polytime feasibility)**.** *Let $\rho \in (0,1]$ arbitrary, and let*

$$\varepsilon_1 \;:=\; \epsilon^{3/2}, \quad \varepsilon_2 \;:=\; \epsilon, \qquad \textit{with} \qquad \epsilon \;:=\; K^{-1}\, \left(\rho\, \sigma^2/e^2 n^6\right)^{1+1/\eta},$$

where K is the problem dependent constant

$$K := R_A\,(3\kappa)^{1/\eta}, \quad \textit{with} \quad \kappa := 2(R_A R_Y^2 + \|b\|)R_A, \quad R_A := \|\bar{\mathcal{A}}\| + \sigma.$$

The algorithm from Theorem 2.2 returns a point Y after $O(\epsilon^{-3})$ function evaluations. With probability at least $1-O(\sigma/n^3)^{2t}$, we have that YY^T is $(\rho\,\sigma, \epsilon^{3/2}R_Y, 5\epsilon)$-approximately optimal for (SDP_{ls}).

Proof. The smoothness assumptions in Theorem 2.2 are satisfied since $\mathfrak{M}_\alpha$ is compact. Therefore Y is an $(\varepsilon_1, \varepsilon_2)$-AC point. Note that

$$\delta \;:=\; \varepsilon_1 R_A/\sqrt{\varepsilon_2} \;=\; \epsilon\, R_A \;=\; (1/3\kappa)^{1/\eta}(\rho\,\sigma^2/e^2n^6)^{1+1/\eta}$$

is as in Corollary 4.6. Hence YY^T is $(\rho\,\sigma, \varepsilon_1 R_Y, 5\varepsilon_2)$-approximately optimal for (SDP_{ls}) with probability $1-O(\sigma/n^3)^{2t}$. □

Remark. The above theorem holds even if the optimal value of (SDP_{ls}) is nonzero. In the special case that the optimal value is zero, then by Lemma 4.2 we have that

$$\|\mathcal{A}(YY^T)-b\| \le \max\{\varepsilon_0',\; n^{1/4}(\varepsilon_1' + \varepsilon_2' R_Y)^{1/2}\}.$$

Let $(\overline{SDP}_{ls})$ denote the instance of problem (SDP_{ls}) in which we use the unperturbed constraints $\bar{\mathcal{A}}$. We next show that YY^T is also approximately optimal for $(\overline{SDP}_{ls})$.

Corollary 5.4. *Consider the setup of Theorem 5.3. With probability at least $1-O(\sigma/n^3)^{2t}$, the matrix YY^T is $(\varepsilon_0'', \varepsilon_1'', \varepsilon_2'')$-approximately optimal for $(\overline{SDP}_{ls})$, where $\varepsilon_0'', \varepsilon_1'', \varepsilon_2'' = O(\sigma)$.*

Proof. We know that the matrix $X := YY^T$ is $(\varepsilon_0', \varepsilon_1', \varepsilon_2')$-approximately optimal for (SDP) with high probability. There are two cases. The first case is that $\|\mathcal{A}(X)-b\| \le \varepsilon_0'$, which implies that

$$\|\bar{\mathcal{A}}(X)-b\| \;\le\; \|\mathcal{A}(X)-b\| + \|(\bar{\mathcal{A}}-\mathcal{A})X\| \;\le\; \varepsilon_0' + \sigma\|X\| \;\le\; O(\sigma).$$

This means that $\varepsilon_0'' = O(\sigma)$. Consider the following variables:

$$\lambda := 2(\mathcal{A}(X)-b), \quad S := \mathcal{A}^*(\lambda), \quad \bar{\lambda} := 2(\bar{\mathcal{A}}(X)-b), \quad \bar{S} := \bar{\mathcal{A}}^*(\bar{\lambda}).$$

The second case is that $\|SX\| \le \varepsilon_1$ and $S \succeq -\varepsilon_2 I_n$. Observe that

$$\begin{aligned} \|\bar{\lambda}-\lambda\| \;&\le\; 2\,\|\bar{\mathcal{A}}-\mathcal{A}\|\,\|X\| \;\le\; O(\sigma), \\ \|\bar{S}-S\| \;&\le\; \|\bar{\mathcal{A}}^*(\bar{\lambda}-\lambda)\| + \|(\bar{\mathcal{A}}^*-\mathcal{A}^*)\lambda\| \;\le\; O(\sigma). \end{aligned}$$

From (12) and (13) we get that $\|\bar{S}X\| \le O(\sigma)$ and $\bar{S} \succeq -O(\sigma)I_n$. So the optimality conditions of $(\overline{SDP}_{ls})$ hold with $\varepsilon_0'', \varepsilon_1'', \varepsilon_2'' = O(\sigma)$. □

6. Experiments

We present some experimental results to complement our theorems. We rely on the open source library `NLopt` [23] to solve the nonlinear problem (BM). Concretely, we use the *augmented Lagrangian method* (ALM) implemented in `NLopt` (which is based on [7]), and we use the preconditioned truncated Newton method as the subroutine. We also rely on `Mosek` to solve semidefinite programs.

For our first experiment we consider a random SDP with a planted solution. More precisely, we consider a matrix $X_0 \in \mathbb{S}^n$, $X_0 \succeq 0$ of rank r, where $n := 50$ and $r \in \{4, 7, 12\}$. We then generate a random SDP for which X_0 is an optimal solution. To do so, we first generate $m := \tau(r)$ random constraints that are satisfied at X_0. Afterwards, we find a cost matrix C in the normal cone of X_0 (this requires solving an auxiliary SDP).

For each $r \in \{4, 7, 12\}$ we generate 100 random SDPs as above. We solve these SDPs with the Burer-Monteiro method, using different values of p (the rank of Y) and random initializations. Figure 1 shows the percentage of experiments solved correctly for each value of r and p. We regard an experiment as "correct" if the criticality conditions from (4) are satisfied.

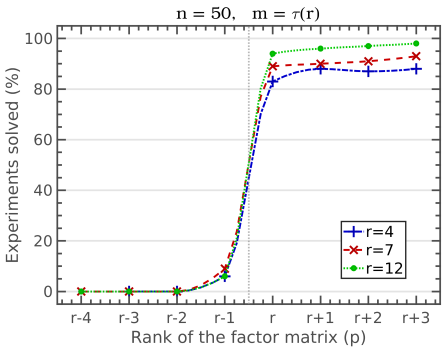

Figure 1. Performance of the Burer-Monteiro method for a class of random SDPs with a planted solution of rank r.

Figure 1 illustrates that there is a sharp phase transition at the Barvinok-Pataki bound $p = r$. Above the Barvinok-Pataki bound, the Burer-Monteiro method solves most instances. Beneath the Barvinok-Pataki bound, it is not just that our techniques stop working, but that the method itself usually fails. Note that previous work [4, 30] required p to be larger than $3r/\sqrt{2}$. However, we see from experiments that the phase transition is much sharper than this. It remains an interesting open question to investigate for what types of structured SDPs having a smaller p might suffice [2, 28].

Observe that, even for $p \geq r$, the number of experiments solved is always below 100%. Nonetheless, the number of bad instances seems to get smaller for larger values of p. This agrees with our result from Theorem 3.5.

For the second experiment we fix the parameters $n := 50$, $m := 28$, and $p := r := 7$. Among the 100 random SDPs considered in Figure 1, we take an instance for which the Burer-Monteiro problem performed badly. We then perturb this seemingly bad SDP by adding varying amounts of noise σ. For each noise level we solve 70 random experiments, in which both the perturbations and the initializations are random. Figure 2 summarizes the results obtained.

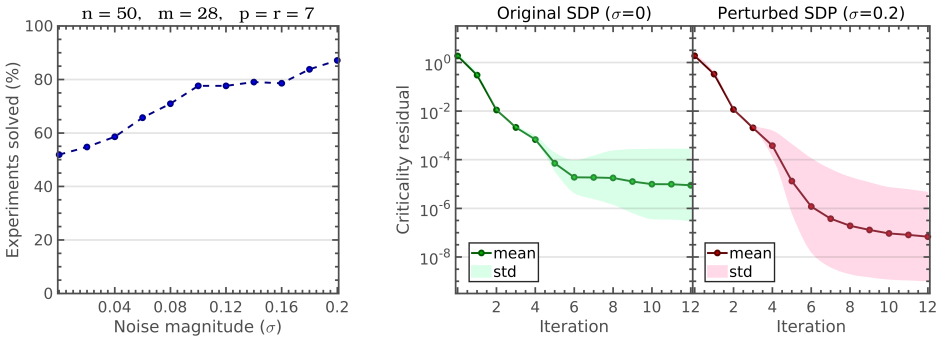

(a) Fraction of SDPs solved

(b) Iteration progress

FIGURE 2. Performance of the Burer-Monteiro method after perturbing a "bad" SDP with different levels of noise.

Figure 2a shows the percentage of instances for which the Burer-Monteiro method succeeded for each noise level. The percentage is with respect to the random perturbation and the random initialization. For the unperturbed problem ($\sigma = 0$) the method succeeds only for 52% of the random initializations. This percentage increases as we add noise. For $\sigma = 0.2$ the method succeeds 87% of the time. Figure 2b shows the progress of the algorithm for the cases $\sigma = 0$ and $\sigma = 0.2$. The progress is measured in terms of the residual of the criticality conditions. The figure shows the mean and standard deviation of the residual for each iteration of ALM.

Figures 2a and 2b illustrate the advantages of smoothing a badly behaved SDP. Theorem 5.1 predicts that the complexity of the algorithm is proportional to σ^{-d} for some exponent d. So for a fixed number of iterations N, we should set the noise level proportional to $N^{-1/d}$. However, our bounds were shown for the algorithm from Theorem 2.6. We do not know if they also apply to ALM.

7. Discussion

The Burer-Monteiro problem is often solved in practice using ALM. In particular, this was the method of choice in the original work by Burer and Monteiro [10, 11]. This motivates the following open problem.

Problem 7.1. *If we solve problem (BM) using ALM, can we provably find an approximately optimal solution for (SDP) in polynomial time?*

In order to answer the above question, one would need an analogue of Theorem 2.6. Namely, we need to show that ALM converges to an AFAC point in polynomial time. A step in this direction was recently given by Sahin et al. [31]. They proved that ALM computes a point satisfying the 2nd-order condition for the augmented Lagrangian function. However, this does not imply that the 2nd-order condition holds for the original problem.

Appendix A. Computing AFAC points

In this appendix we prove Theorem 2.6, which states that AFAC points can be computed in polynomial time.

A.1. **The algorithm.** Cartis et al. [17] proposed a method for computing q-th order critical points for $q \in \{1,2,3\}$. However, they use a nonstandard notion of criticality which is not easy to translate into our setting. We present here a slight modification of this algorithm that accommodates more general criticality conditions.

We focus on the constrained optimization problem (P_{con}). Consider the least squares functions

$$\nu(y) := \|h(y)\|^2, \qquad \mu(t,y) := (f(y)-t)^2 + \|h(y)\|^2.$$

We denote $\mu_t = \mu(t,\cdot)$ the function obtained by fixing the value of t. Algorithm 1 below is a variant of the method from [17]. It consists of two phases. The first phase attempts to find an approximately feasible solution through the unconstrained problem $\min_y \nu(y)$. If successful, the second phase minimizes f while preserving feasibility. To do so, it solves a sequence of problems $\min_y \mu(t_k,y)$, where the values $\{t_k\}_{k\geq 0}$ are decreasing.

Algorithm 1 relies on an *inner method* for solving the unconstrained problem $\min_y \psi(y)$, where ψ is either ν or $\mu_t = \mu(t,\cdot)$. Given $\boldsymbol{\epsilon} = (\epsilon_1,\ldots,\epsilon_q) \in \mathbb{R}^q_+$, the inner method looks for a point y such that $\boldsymbol{\chi}(\psi,y) \leq \boldsymbol{\epsilon}$, for some *criticality measure* $\boldsymbol{\chi} = (\chi_1,\ldots,\chi_q)$. We assume that the j-th component $\chi_j(\psi,y)$ only involves derivatives $\{\nabla^d \psi(y)\}_{d\leq j}$ up to order j. For instance, the AC-criticality condition from (1) corresponds to the case

$$\boldsymbol{\chi}^{\text{AC}}(\psi,y) := \big(\, \|\nabla\psi(y)\|, -\min \operatorname{eig}(\nabla^2\psi(y)) \,\big). \tag{14}$$

Given an initial point y^0 and tolerances $\boldsymbol{\epsilon} \in \mathbb{R}^q_+$, the inner method produces iterates $\{y^i\}_{i=1}^N$. We assume that the final point y^N achieves these tolerances and that the objective function decreases proportionately to N:

$$\boldsymbol{\chi}(\psi,y^N) \leq \boldsymbol{\epsilon} \qquad \text{and} \qquad \psi(y^0) - \psi(y^N) \geq N\,\kappa_\psi\, p(\boldsymbol{\epsilon}), \tag{15}$$

for some constant $\kappa_\psi > 0$ and some function p. Hence, the number of iterations N is proportional to $p(\boldsymbol{\epsilon})^{-1}$.

The next theorem provides guarantees for Algorithm 1. Our proof closely follows that of [17, Thm.4.5] but has the advantage that it applies to a general class of criticality measures, as opposed to [17], which relies on a

Algorithm 1 Constrained optimization algorithm based on [17]

Input: Initial point $y_0 \in \mathbb{R}^n$, tolerances $\epsilon_0 \in \mathbb{R}_+$, $\boldsymbol{\epsilon} \in \mathbb{R}^q_+$, constant $\delta \in (0,1)$.
Output: A point $y \in \mathbb{R}^n$ and a number $t \leq f(y)$.

PHASE I

$y_1 := \operatorname{local\,min}_y \nu(y)$ starting with y_0
$t_0 := f(y_1)$
if $\nu(y_1) > (\delta\epsilon_0)^2$ **then return** (y_1, t_0)

PHASE II

$t_1 := f(y_1) - (\epsilon_0^2 - \nu(y_1))^{1/2}$
for $k = 2, 3, 4, \ldots$ **do**
  $y_k := \operatorname{local\,min}_y \mu(t_{k-1}, y)$ starting with y_{k-1}
  if $\mu(t_{k-1}, y_k) < (\delta\epsilon_0)^2$ **then** ▷ case (a)
    $t_k := f(y_k) - (\epsilon_0^2 - \nu(y_k))^{1/2}$
    if $\boldsymbol{\chi}(\mu_{t_k}, y_k) \leq \boldsymbol{\epsilon}$ **then return** (y_k, t_k)
  if $\mu(t_{k-1}, y_k) \geq (\delta\epsilon_0)^2$ & $f(y_k) < t_{k-1}$ **then** ▷ case (b)
    $t_k := 2f(y_k) - t_{k-1}$
    if $\boldsymbol{\chi}(\mu_{t_k}, y_k) \leq \boldsymbol{\epsilon}$ **then return** (y_k, t_k)
  if $\mu(t_{k-1}, y_k) \geq (\delta\epsilon_0)^2$ & $f(y_k) \geq t_{k-1}$ **then** ▷ case (c)
    return (y_k, t_k), with $t_k := t_{k-1}$

particular nonstandard measure of criticality. However our complexity is larger than in [17] by a factor of ϵ_0^{-1}.

Theorem A.1. *Assume that:*

- *The inner method satisfies* (15) *for the function* ν *with constant* κ_ν.
- *The inner method satisfies* (15) *for the function* μ_t, *and the constant* κ_μ *is independent of* t.
- *There exists* $\beta > \epsilon_0$ *and* $f_{\text{low}} \in \mathbb{R}$ *such that* $f(y) \geq f_{\text{low}}$ *for all* $y \in \mathfrak{M}_\beta$, *where* $\mathfrak{M}_\beta := \{y : \|h(y)\| \leq \beta\}$.

Then the total number of inner iterations made in Algorithm 1 is at most

$$p(\boldsymbol{\epsilon})^{-1} \left(\kappa_\nu^{-1} \, \nu(y_0) \, + \, \epsilon_0 \, \kappa_\mu^{-1} \, (1-\delta)^{-1} \, (f(y_1) - f_{\text{low}} + \beta) \right), \tag{16}$$

and the algorithm returns a pair (y, t) *such that:*

$$\text{either} \qquad t < f(y), \qquad \|h(y)\| \leq \epsilon_0, \qquad \boldsymbol{\chi}(\mu_t, y) \leq \boldsymbol{\epsilon}, \tag{17a}$$

$$\text{or} \qquad t = f(y), \qquad \|h(y)\| > \delta\epsilon_0, \qquad \chi_1(\nu, y) \leq \epsilon_1. \tag{17b}$$

A.2. **Proof of Theorem A.1.** Let K be the number of outer iterations of Algorithm 1. Consider the sets of indices:

$$A := \{1\} \cup \{k : 2 \leq k \leq K \text{ and case (a) is applied}\},$$
$$B := \{k : 2 \leq k \leq K \text{ and case (b) is applied}\}.$$

The following lemma gives a few properties of Algorithm 1. Its proof is identical to [17, Lem.3.1].

Lemma A.2. *If the algorithm reaches Phase II, then:*

$$\nu(y_k) \le \mu(t_k, y_k) \le \epsilon_0^2, \quad 0 \le f(y_k) - t_k \le \epsilon_0, \qquad \text{for } k \ge 1, \tag{18}$$

$$\mu(t_k, y_k) = \epsilon_0^2, \quad t_{k-1} - t_k \ge (1-\delta)\epsilon_0, \qquad \text{for } k \in A, \tag{19}$$

$$\mu(t_k, y_k) = \mu(t_{k-1}, y_k), \quad t_{k-1} > t_k, \qquad \text{for } k \in B, \tag{20}$$

$$\mu(t_k, y_k) \ge (\delta\epsilon_0)^2, \quad \boldsymbol{\chi}(\mu_{t_k}, y_k) \le \boldsymbol{\epsilon}, \qquad \text{for } k = K. \tag{21}$$

Let (y, t) be the output of Algorithm 1, and let us show (17). Assume first that the algorithm terminates in Phase I. Then y is a local minimum of ν, $\nu(y) > (\delta\epsilon_0)^2$, and $t = f(y)$. Hence (17b) holds. Assume now that the algorithm terminates in Phase II. By (18) and (21), we have

$$t \le f(y), \qquad (\delta\epsilon_0)^2 \le \mu_t(y) \le \epsilon_0^2, \qquad \boldsymbol{\chi}(\mu_t, y) \le \boldsymbol{\epsilon}.$$

If $f(y) < t$ then $\|h(y)\| \le \sqrt{\mu_t(y)} \le \epsilon_0$, so (17a) holds. Consider now the case that $f(y) = t$. Note that $\mu_t(y) = \nu(y)$, $\nabla\mu_t(y) = \nabla\nu(y)$. Then $\chi_1(\mu_t, y) = \chi_1(\nu, y)$, as they only involve derivatives up to order 1. Since $\|h(y)\| = \sqrt{\mu_t(y)} \ge \delta\epsilon_0$, then (17b) holds.

We proceed to show that the number of inner iterations is bounded by (16). Each outer iteration k of Algorithm 1 calls the inner method once. Let N_k be the number of inner iterations made in this call. The total number of inner iterations is $\sum_{k=1}^{K} N_k$.

We first analyze Phase I. The inner method is applied to the problem $\min_y \nu(y)$, starting with y_0 and terminating with y_1. By (15), we have

$$\nu(y_0) \;\ge\; \nu(y_0) - \nu(y_1) \;\ge\; N_1\, \kappa_\nu\, p(\boldsymbol{\epsilon}).$$

It follows that $N_1 \le \nu(y_0)/\kappa_\nu p(\epsilon)$.

We proceed to Phase II. For each $a \in A$, let $n(a)$ be the next integer that lies in A. For the largest $a \in A$ we define $n(a) := K$, where K is the final iteration. We can group the indices $k \ge 2$ as follows:

$$\{2, 3, \ldots, K\} = \bigcup_{a \in A} K_a, \qquad \text{where} \qquad K_a := \{a{+}1, a{+}2, \ldots, n(a)\}.$$

We will show that for any $a \in A$ we have that

$$N(K_a) \;:=\; \sum_{k \in K_a} N_k \;\le\; \epsilon_0^2/\kappa_\mu p(\epsilon). \tag{22}$$

Consider an iteration $k \in K_a$. The inner method is applied to $\min_y \mu(t_{k-1}, y)$, starting with y_{k-1} and terminating with y_k. By (15), we have

$$\mu(t_{k-1}, y_{k-1}) - \mu(t_{k-1}, y_k) \;\ge\; N_k\, \kappa_\mu\, p(\boldsymbol{\epsilon}).$$

Observe that $K_a \setminus \{n(a)\} \subset B$. By (20), we have

$$\mu(t_{k-1}, y_k) = \mu(t_k, y_k) \quad \text{for} \quad k \in K_a \setminus \{n(a)\}.$$

Also note that $\mu(t_a, y_a) = \epsilon_0^2$ by (19). Therefore,

$$\begin{aligned}\epsilon_0^2 \;&\geq\; \mu(t_a, y_a) - \mu(t_{n(a)}{-}1, y_{n(a)}) \\ &=\; \sum_{k\in K_a} \mu(t_{k-1}, y_{k-1}) - \mu(t_{k-1}, y_k) \;\geq\; \sum_{k\in K_a} N_k\, \kappa_\mu\, p(\boldsymbol{\epsilon}).\end{aligned}$$

By rearranging the above inequality we get (22).

Let us now upper bound the cardinality of A. By (19) and (20) we have that $t_{k-1}{-}t_k$ is at least $(1{-}\delta)\epsilon_0$ for $k \in A$, and is positive for $k \in B$. Also note that $t_0{=}f(y_1)$ and $t_K{\geq}f(y_K){-}\epsilon_0{\geq}f_{\text{low}}{-}\beta$ by (18). Then

$$f(y_1){-}f_{\text{low}}{+}\beta \;\geq\; t_0{-}t_K \;=\; \sum_{k=1}^{K}(t_{k-1}{-}t_k) \;\geq\; \sum_{k\in A}(t_{k-1}{-}t_k) \;\geq\; |A|\,(1{-}\delta)\epsilon_0,$$

and hence $|A| \leq (f(y_1){-}f_{\text{low}}{+}\beta)/(1{-}\delta)\epsilon_0$.

Combining everything, we derive

$$\sum_{k=1}^{K} N_k \;\leq\; N_1 + |A| \cdot \max_{a\in A} N(K_a) \;\leq\; \frac{\nu(y_0)}{\kappa_\nu p(\epsilon)} + \frac{f(y_1){-}f_{\text{low}}{+}\beta}{(1{-}\delta)\epsilon_0} \cdot \frac{\epsilon_0^2}{\kappa_\mu\, p(\epsilon)},$$

which is equal to (16).

A.3. **Proof of Theorem 2.6.** We finally show that AFAC points can be computed in polynomial time. Let $\varepsilon_0, \varepsilon_1, \varepsilon_2, \gamma, R_\lambda$ be as in the statement of Theorem 2.6. We consider Algorithm 1 with parameters

$$\begin{aligned}\delta \;&:=\; 1/2, \qquad q \;:=\; 2, \qquad \boldsymbol{\epsilon} \;:=\; (\epsilon_1, \epsilon_2), \\ \epsilon_0 \;&:=\; \varepsilon_0, \quad \epsilon_1 \;:=\; R_\lambda^{-1}\,\varepsilon_0\,\varepsilon_1, \quad \epsilon_2 \;:=\; \tfrac{1}{2} R_\lambda^{-1}\,\varepsilon_0\,\varepsilon_2.\end{aligned}$$

For the inner method we use the ARC algorithm from Theorem 2.2, using the criticality measure (14). Algorithm 1 returns a pair (y, t). The associated multiplier is $\lambda := (f(y){-}t)^{-1}\, h(y) \in \mathbb{R}^m$, which is defined only if $f(y) \neq t$.

In order to apply Theorem A.1, we have to check that the functions ν and $\mu_t{=}\mu(t,\cdot)$ are smooth enough so that the inner algorithm satisfies (15).

Lemma A.3 ([17, Lem.4.1])**.** *Assume that $\{\nabla^j f\}_{j=0}^q, \{\nabla^j h\}_{j=0}^q$ are uniformly bounded and Lipschitz continuous on a set $D \subset \mathbb{R}^n$. Then*

(i) $\{\nabla^j \nu\}_{j=0}^q$ are uniformly bounded and Lipschitz continuous on D.

(ii) $\{\nabla^j \mu_t\}_{j=0}^q$ are uniformly bounded and Lipschitz continuous on $D{\cap}B_t$, with $B_t := \{y : |f(y){-}t|{\leq}1\}$, and the constants are independent of t.

The above lemma shows that ν is smooth on $\mathfrak{M}_\beta$ and μ_t is smooth on $\mathfrak{M}_\beta \cap B_t$, with $B_t := \{y : |f(y){-}t|{\leq}1\}$. Note that all points y_k produced by Algorithm 1 lie in $\mathfrak{M}_\beta \cap B_t$ because of (18). Since ν, μ_t are sufficiently smooth, we can apply Theorem 2.2 (see also [13]). We conclude that the inner method satisfies (15) with

$$p(\boldsymbol{\epsilon}) \;=\; \min\{\epsilon_1^2,\, \epsilon_2^3\} \;=\; \Omega(\min\{\varepsilon_0^2\,\varepsilon_1^2,\, \varepsilon_0^3\,\varepsilon_2^3\}).$$

Hence, by Theorem A.1, the total number of inner iterations is $O(p(\boldsymbol{\epsilon})^{-1}) = O(\max\{\varepsilon_0^{-2}\,\varepsilon_1^{-2}, \varepsilon_0^{-3}\,\varepsilon_2^{-3}\})$. Since each inner iteration requires $O(1)$ function evaluations (see Theorem 2.2), then the total number of function evaluations has the same order of magnitude.

Let us see that the conditions (2) hold. Let (y,t) be the output of Algorithm 1. By Theorem A.1, this pair satisfies either (17a) or (17b). Let us see that (17b) cannot occur. Assume that

$$\|h(y)\| > \epsilon_0/2, \qquad \|\nabla\nu(y)\| \le \epsilon_1.$$

Observe that $\|h(y)\| \le \epsilon_0 \le \beta$ by (18), and hence ϱ-LICQ holds at y. Then

$$\varrho\,\epsilon_0 \;<\; 2\,\varrho\|h(y)\| \;\le\; 2\,\|h(y)^T\nabla h(y)\| \;=\; \|\nabla\nu(y)\| \;\le\; \epsilon_1.$$

Also note that that

$$R_\lambda^{-1} \;\le\; \tfrac{1}{2}\varrho(1+L_f)^{-1} \;\le\; \tfrac{1}{2}\varrho, \qquad \epsilon_1/\epsilon_0 \;=\; R_\lambda^{-1}\,\varepsilon_1 \;\le\; \tfrac{1}{2}\varrho\,\varepsilon_1 \;\le\; \tfrac{1}{2}\varrho.$$

The last two equations give a contradiction.

Then the output (y,t) satisfies (17a). Hence, $t < f(y)$ and

$$\|h(y)\| \;\le\; \epsilon_0, \qquad \|\nabla\mu_t(y)\| \;\le\; \epsilon_1, \qquad \nabla^2\mu_t(y) \;\succeq\; -\epsilon_2 I_n.$$

Let $\alpha := (f(y)-t)^{-1}$, so that $\lambda = \alpha\,h(y)$. It can be checked that $\alpha^2\mu_t(y) = \|(1,\lambda)\|^2$. Note that $\mu_t(y) \ge (\epsilon_0/2)^2$ by (21), and hence

$$\alpha \;=\; \mu_t(y)^{-1/2}\,\|(1,\lambda)\| \;\le\; 2\,\varepsilon_0^{-1}\|(1,\lambda)\|.$$

The Lagrangian function $L(y,\lambda) = f(y)+\lambda\cdot h(x)$ is closely related to $\mu_t(y)$. A simple calculation gives that

$$\nabla L(y,\lambda) \;=\; \alpha\cdot\tfrac{1}{2}\nabla\mu_t(y), \qquad \nabla^2 L(y,\lambda) \;=\; \alpha\,\big(\tfrac{1}{2}\nabla^2\mu_t(y) - \tilde{J}^T\tilde{J}\big), \tag{23}$$

where $\tilde{J} := \begin{pmatrix}\nabla f(y)\\ \nabla h(y)\end{pmatrix}$ is the augmented Jacobian.

We proceed to verify (2a). We already have that $\|h(y)\| \le \varepsilon_0$. Note that

$$\|\nabla L(y,\lambda)\| \;=\; \tfrac{1}{2}\,\alpha\,\|\nabla\mu_t(y)\| \;\le\; \varepsilon_0^{-1}\|(1,\lambda)\|\,\epsilon_1 \;=\; R_\lambda^{-1}\|(1,\lambda)\|\,\varepsilon_1. \tag{24}$$

We claim that $\|(1,\lambda)\| \le R_\lambda$. By Lemma 2.5 and $R_\lambda^{-1} \le \varrho/2$, $\varepsilon_1 \le 1$, we have

$$\|\lambda\| \;\le\; \varrho^{-1}\,(R_\lambda^{-1}\,\varepsilon_1\|(1,\lambda)\| + \|\nabla f(y)\|) \;\le\; \tfrac{1}{2}(1+\|\lambda\|) + \varrho^{-1}L_f.$$

It follows that $\|\lambda\| \le 1+2\varrho^{-1}L_f = R_\lambda - 1$ and hence $\|(1,\lambda)\| \le R_\lambda$, as we claimed. Then $\|\nabla L(y,\lambda)\| \le \varepsilon_1$ by (24).

We now verify (2b). Let $u \in \mathbb{R}^n$ of unit norm such that $\|Ju\| \le \gamma$, where $J := \nabla h(y)$. We need to show that $u^T\nabla^2 L(y,\lambda)u \ge -\varepsilon_2$. By (23), we have

$$u^T\nabla^2 L(y,\lambda)u \;=\; \alpha\,(\tfrac{1}{2}u^T\nabla^2\mu_t(y)u - \|\tilde{J}u\|^2). \tag{25}$$

Note that $u^T\nabla^2\mu_t(y)u \ge -\epsilon_2 = -\frac{1}{2}R_\lambda^{-1}\varepsilon_0\varepsilon_2$. We bound $\|\tilde{J}u\|$ next:

$$\tilde{J} \;=\; \begin{pmatrix}\nabla f(y)\\ \nabla h(y)\end{pmatrix} \;=\; \begin{pmatrix}\nabla L(y,\lambda)\\ 0\end{pmatrix} + \begin{pmatrix}-\lambda^T J\\ J\end{pmatrix},$$

$$\|\tilde{J}u\| \;\le\; \|\nabla L(y,\lambda)\| + \|(1,\lambda)\|\,\|Ju\| \;\le\; \varepsilon_1 + \gamma\,\|(1,\lambda)\| \;\le\; \tfrac{1}{2}(R_\lambda^{-1}\,\varepsilon_0\,\varepsilon_2)^{1/2},$$

where we used that ε_1 and γR_λ are at most $\frac{1}{4}(R_\lambda^{-1}\varepsilon_0\varepsilon_2)^{1/2}$ by (3). Hence

$$\alpha\,(\tfrac{1}{2}u^T\nabla^2\mu_t(y)u - \|\tilde{J}u\|^2) \;\geq\; -(2\varepsilon_0^{-1}\|(1,\lambda)\|)\cdot(\tfrac{1}{2}R_\lambda^{-1}\varepsilon_0\,\varepsilon_2) \;\geq -\varepsilon_2.$$

Together with (25), we get that $u^T\nabla^2 L(y,\lambda)u \geq -\varepsilon_2$.

References

[1] R. Andreani, E. G. Birgin, J. M. Martínez, and M. L. Schuverdt. Second-order negative-curvature methods for box-constrained and general constrained optimization. *Comput. Optim. Appl.*, 45(2):209–236, 2010.

[2] A. S. Bandeira, N. Boumal, and V. Voroninski. On the low-rank approach for semidefinite programs arising in synchronization and community detection. In *Proc. Conf. Learn. Theory*, pages 361–382, 2016.

[3] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete Comput. Geom.*, 13(2):189–202, 1995.

[4] S. Bhojanapalli, N. Boumal, P. Jain, and P. Netrapalli. Smoothed analysis for low-rank solutions to semidefinite programs in quadratic penalty form. In *Conf. Learn. Theory*, pages 3243–3270, 2018.

[5] S. Bhojanapalli, B. Neyshabur, and N. Srebro. Global optimality of local search for low rank matrix recovery. In *Adv. Neural Inf. Process. Syst.*, pages 3873–3881, 2016.

[6] E. G. Birgin, J. Gardenghi, J. M. Martinez, S. A. Santos, and P. L. Toint. Evaluation complexity for nonlinear constrained optimization using unscaled KKT conditions and high-order models. *SIAM J. Optim.*, 26(2):951–967, 2016.

[7] E. G. Birgin and J. M. Martínez. Improving ultimate convergence of an augmented Lagrangian method. *Optim. Methods Software*, 23(2):177–195, 2008.

[8] N. Boumal, V. Voroninski, and A. Bandeira. The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In *Adv. Neural Inf. Process. Syst.*, pages 2757–2765, 2016.

[9] N. Boumal, V. Voroninski, and A. Bandeira. Deterministic guarantees for Burer-Monteiro factorizations of smooth semidefinite programs. *Commun. Pure Appl. Math.*, 2019.

[10] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.*, 95(2):329–357, 2003.

[11] S. Burer and R. D. Monteiro. Local minima and convergence in low-rank semidefinite programming. *Math. Program.*, 103(3):427–444, 2005.

[12] P. Bürgisser, F. Cucker, and M. Lotz. The probability that a slightly perturbed numerical analysis problem is difficult. *Math. Comput.*, 77(263):1559–1583, 2008.

[13] C. Cartis, N. I. Gould, and P. L. Toint. Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case function-and derivative-evaluation complexity. *Math. Program.*, 130(2):295–319, 2011.

[14] C. Cartis, N. I. Gould, and P. L. Toint. Complexity bounds for second-order optimality in unconstrained optimization. *J. Complexity*, 28(1):93–108, 2012.

[15] C. Cartis, N. I. Gould, and P. L. Toint. On the evaluation complexity of cubic regularization methods for potentially rank-deficient nonlinear least-squares problems and its relevance to constrained nonlinear optimization. *SIAM J. Optim.*, 23(3):1553–1574, 2013.

[16] C. Cartis, N. I. Gould, and P. L. Toint. On the complexity of finding first-order critical points in constrained nonlinear optimization. *Math. Program.*, 144(1-2):93–106, 2014.

[17] C. Cartis, N. I. Gould, and P. L. Toint. Optimality of orders one to three and beyond: characterization and evaluation complexity in constrained nonconvex optimization. *J. Complexity*, 2018.

[18] D. Cifuentes. Burer-Monteiro guarantees for general semidefinite programs. *arXiv:1904.07147*, 2019.

[19] A. R. Conn, N. I. Gould, and P. L. Toint. *Trust region methods*, volume 1 of *MOS-SIAM Series Optim.* SIAM, 2000.
[20] F. E. Curtis, D. P. Robinson, and M. Samadi. Complexity analysis of a trust funnel algorithm for equality constrained optimization. *SIAM J. Optim.*, 28(2):1533–1563, 2018.
[21] R. Ge, C. Jin, and Y. Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proc. Int. Conf. Mach. Learn.*, pages 1233–1242, 2017.
[22] A. Gray. *Tubes*, volume 221. Birkhäuser, 2012.
[23] S. G. Johnson. The NLopt nonlinear-optimization package. Available at `http://github.com/stevengj/nlopt`.
[24] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM J. Optim.*, 20(5):2327–2351, 2010.
[25] S. L. Kleiman. Bertini and his two fundamental theorems. In *Studies in the History of Modern Mathematics III*, volume 55 of *Rend. Circ. Mat. Palermo (2) Suppl.*, pages 9–37, 1998.
[26] Q. Li and G. Tang. The nonconvex geometry of low-rank matrix optimizations with general objective functions. In *IEEE Global Conf. Signal Inf. Process.*, pages 1235–1239, 2017.
[27] M. Lotz. On the volume of tubular neighborhoods of real algebraic varieties. *Proc. Am. Math. Soc.*, 143(5):1875–1889, 2015.
[28] S. Mei, T. Misiakiewicz, A. Montanari, and R. I. Oliveira. Solving SDPs for synchronization and maxcut problems via the Grothendieck inequality. In *Proc. Conf. Learn. Theory*, pages 1476–1515, 2017.
[29] G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Math. Oper. Res.*, 23(2):339–358, 1998.
[30] T. Pumir, S. Jelassi, and N. Boumal. Smoothed analysis of the low-rank approach for smooth semidefinite programs. In *Adv. Neural Inf. Process. Syst.*, pages 2287–2296, 2018.
[31] M. F. Sahin, A. Eftekhari, A. Alacaoglu, F. Latorre, and V. Cevher. An inexact augmented Lagrangian framework for nonconvex optimization with nonlinear constraints. *Proc. Neural Inf. Process. Syst., arXiv:1906.11357*, 2019.
[32] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *JACM*, 51(3):385–463, 2004.
[33] J. F. Sturm. Error bounds for linear matrix inequalities. *SIAM J. Optim.*, 10(4):1228–1248, 2000.
[34] I. Waldspurger and A. Waters. Rank optimality for the Burer-Monteiro factorization. *arXiv:1812.03046*, 2018.
[35] R. Zhang, C. Josz, S. Sojoudi, and J. Lavaei. How much restricted isometry is needed in nonconvex matrix recovery? In *Adv. Neural Inf. Process. Syst.*, pages 5586–5597, 2018.

Massachusetts Institute of Technology, Cambridge, MA, USA
Email address: `diegcif@mit.edu`

Massachusetts Institute of Technology, Cambridge, MA, USA
Email address: `moitra@mit.edu`