
Inexact Augmented Lagrangian Framework for Non-Convex Optimization

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We propose a practical inexact augmented Lagrangian method (iALM) for noncon-
2 vex problems with nonlinear constraints. We characterize the total computational
3 complexity of our method subject to a verifiable geometric condition, which is
4 closely related to the Polyak-Lojswicz and Mangasarian-Fromowitz conditions.

5 In particular, when a first-order solver is used for the inner iterates, we prove that
6 iALM finds a first-order stationary point with $\tilde{O}(1/\epsilon^3)$ calls to the first-order oracle.
7 If, in addition, the problem is smooth and a second-order solver is used for the
8 inner iterates, iALM finds a second-order stationary point with $\tilde{O}(1/\epsilon^5)$ calls to
9 the second-order oracle. These complexity results match the known theoretical
10 results in the literature with a simple, implementable and versatile algorithm.

11 We provide numerical evidence on large-scale machine learning problems, in-
12 cluding the Burer-Monteiro factorization of semidefinite programs, and a novel
13 nonconvex relaxation of the standard basis pursuit template. We verify our geomet-
14 ric condition in all these examples.

15 1 Introduction

16 We study the nonconvex optimization problem

$$\begin{cases} \min_{x \in \mathbb{R}^d} f(x) + g(x) \\ A(x) = 0, \end{cases} \quad (1)$$

17 where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a continuously-differentiable nonconvex function and $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a
18 nonlinear operator. We assume that $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a possibly nonsmooth but proximal-friendly
19 convex function [48].

20 A host of problems in computer science [35, 38, 68], machine learning [41, 57], and signal pro-
21 cessing [55, 56] naturally fall under the template (1), including max-cut, clustering, generalized
22 eigenvalue decomposition, as well as the quadratic assignment problem (QAP) [68].

23 To solve (1), this paper proposes an intuitive and easy-to-implement augmented Lagrangian algorithm,
24 and provides its total iteration complexity under an interpretable geometric condition. Before we
25 elaborate on the results, let us first motivate (1) with an application to semidefinite programming
26 (SDP):

27 **Vignette: Burer-Monteiro splitting.** A powerful convex relaxation for max-cut, clustering, and
28 many others is provided by the SDP

$$\begin{cases} \min_{X \in \mathbb{S}^{d \times d}} \langle C, X \rangle \\ B(X) = b, X \succeq 0, \end{cases} \quad (2)$$

29 where $C \in \mathbb{R}^{d \times d}$, X is a positive semidefinite $d \times d$ matrix, and $B : \mathbb{S}^{d \times d} \rightarrow \mathbb{R}^m$ is a linear operator.
 30 If the unique-games conjecture is true, SDPs achieve the best approximation for the underlying
 31 discrete problem [53].

32 Since d is often large, many first- and second-order methods for solving such SDPs are immedi-
 33 ately ruled out, not only due to their high computational complexity, but also due to their storage
 34 requirements, which are $\mathcal{O}(d^2)$.

35 A contemporary challenge in optimization is therefore to solve SDPs using little space and in a
 36 scalable fashion. The recent homotopy conditional gradient method, which is based on linear
 37 minimization oracles (LMOs), can solve (2) in a small space via sketching [67]. However, such
 38 LMO-based methods are extremely slow in obtaining accurate solutions.

39 A different approach for solving (1), dating back to [15, 16], is the so-called Burer-Monteiro (BM)
 40 factorization $X = UU^\top$, where $U \in \mathbb{R}^{d \times r}$ and r is selected according to the guidelines in [50, 2],
 41 which are shown to be optimal [61]. This factorization does not introduce any extraneous local
 42 minima [16] and, moreover, [14] established the connection between the local minimizers of the
 43 factorized problem (3) and the global minimizers for (2).

44 This factorization leads to the nonconvex problem

$$\begin{cases} \min_{U \in \mathbb{R}^{d \times r}} \langle C, UU^\top \rangle \\ B(UU^\top) = b, \end{cases} \quad (3)$$

45 which can be easily written in the form of (1). To solve (3), the inexact Augmented Lagrangian
 46 method (iALM) is widely used [15, 16, 36], due to its cheap per iteration cost and its empirical
 47 success. Every (outer) iteration of iALM calls a solver to solve an intermediate augmented Lagrangian
 48 subproblem to near stationarity. The user is free in the choice of this solver, which could use first-
 49 order, such as the proximal gradient descent [48], or second-order information, such as the trust
 50 region method and BFGS [45].¹

51 Unlike its convex counterpart [42, 37, 63], the convergence rate and the complexity of iALM for (3)
 52 are not well-understood, see Section 5 for a review of the related literature. Indeed, addressing this
 53 important theoretical gap is one of the contributions of our work.

54 **Summary of contributions:**

- 55 ◦ We derive the convergence rate of iALM for solving (1) to first- or second-order optimality, and
 56 find the total iteration complexity of iALM using different solvers for the augmented Lagrangian
 57 subproblems. Our complexity bounds match the best theoretical results in optimization, see Section 5.
- 58 ◦ Our results are future-proof in the sense that they are independent of the choice of solver called by
 59 iALM.
- 60 ◦ We propose a geometric condition that simplifies the algorithmic analysis for iALM, and clarify its
 61 connection to well-known Polyak-Lojasiewicz [34] and Mangasarian-Fromovitz [4] conditions. We
 62 also verify this condition for key problems in Section 6.

63 **Roadmap.** Section 2 collects the main tools and our notation. We present the iALM in Section 3
 64 and obtain its convergence rate to first- and second-order stationary points in Section 4, alongside
 65 their iteration complexities. We provide a comprehensive review of the literature and highlight our
 66 key differences in Section 5. Section 6 presents the numerical evidence and comparisons with the
 67 state-of-the-art techniques.

68 **2 Preliminaries**

69 **Notation.** We use the notation $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ for the standard inner product and the norm on \mathbb{R}^d . For
 70 matrices, $\|\cdot\|$ and $\|\cdot\|_F$ denote the spectral and the Frobenius norms, respectively. For the convex
 71 function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, the subdifferential set at $x \in \mathbb{R}^d$ is denoted by $\partial g(x)$ and we will occasionally
 72 use the notation $\partial g(x)/\beta = \{z/\beta : z \in \partial g(x)\}$. When presenting iteration complexity results, we
 73 often use $\tilde{O}(\cdot)$ which suppresses the logarithmic dependencies.

¹Strictly speaking, BFGS is in fact a quasi-Newton method that emulates second-order information.

74 We use the indicator function $\delta_{\mathcal{X}} : \mathbb{R}^d \rightarrow \mathbb{R}$ of a set $\mathcal{X} \subset \mathbb{R}^d$, which takes x to

$$\delta_{\mathcal{X}}(x) = \begin{cases} 0 & x \in \mathcal{X} \\ \infty & x \notin \mathcal{X}. \end{cases} \quad (4)$$

75 The distance function from a point x to \mathcal{X} is denoted by $\text{dist}(x, \mathcal{X}) = \min_{z \in \mathcal{X}} \|x - z\|$. For integers
76 $k_0 \leq k_1$, we denote $[k_0 : k_1] = \{k_0, \dots, k_1\}$.

77 For an operator $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with components $\{A_i\}_{i=1}^m$, we let $DA(x) \in \mathbb{R}^{m \times d}$ denote the
78 Jacobian of A , where the i th row of $DA(x)$ is the gradient vector $\nabla A_i(x) \in \mathbb{R}^d$.

79 **Smoothness.** We require $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $A : \mathbb{R}^d \rightarrow \mathbb{R}^m$ to be smooth, namely, there exist
80 $\lambda_f, \lambda_A \geq 0$ such that

$$\begin{aligned} \|\nabla f(x) - \nabla f(x')\| &\leq \lambda_f \|x - x'\|, \\ \|DA(x) - DA(x')\| &\leq \lambda_A \|x - x'\|, \end{aligned} \quad (5)$$

81 for every $x, x' \in \mathbb{R}^d$.

82 **Augmented Lagrangian method (ALM).** ALM is a classical algorithm, which first appeared in [31,
83 52] and extensively studied afterwards in [4, 9]. For solving (1), ALM suggests solving the problem

$$\min_x \max_y \mathcal{L}_\beta(x, y) + g(x), \quad (6)$$

84 where, for penalty weight $\beta > 0$, \mathcal{L}_β is the corresponding augmented Lagrangian, defined as

$$\mathcal{L}_\beta(x, y) := f(x) + \langle A(x), y \rangle + \frac{\beta}{2} \|A(x)\|^2. \quad (7)$$

85 The minimax formulation in (6) naturally suggests the following algorithm for solving (1). For dual
86 step sizes $\{\sigma_k\}_k$, consider the iterations

$$x_{k+1} \in \underset{x}{\text{argmin}} \mathcal{L}_\beta(x, y_k) + g(x), \quad (8)$$

87

$$y_{k+1} = y_k + \sigma_k A(x_{k+1}).$$

88 However, computing x_{k+1} above requires solving the nonconvex problem (8) to optimality, which is
89 typically intractable. Instead, it is often easier to find an approximate first- or second-order stationary
90 point of (8).

91 Hence, we argue that by gradually improving the stationarity precision and increasing the penalty
92 weight β above, we can reach a stationary point of the main problem in (1), as detailed in Section 3.

93 **Optimality conditions.** First-order necessary optimality conditions for (1) are well-studied. Indeed,
94 $x \in \mathbb{R}^d$ is a first-order stationary point of (1) if there exists $y \in \mathbb{R}^m$ such that

$$\begin{cases} -\nabla f(x) - DA(x)^\top y \in \partial g(x) \\ A(x) = 0, \end{cases} \quad (9)$$

95 where $DA(x)$ is the Jacobian of A at x . Recalling (7), we observe that (9) is equivalent to

$$\begin{cases} -\nabla_x \mathcal{L}_\beta(x, y) \in \partial g(x) \\ A(x) = 0, \end{cases} \quad (10)$$

96 which is in turn the necessary optimality condition for (6). Inspired by this, we say that x is an (ϵ_f, β)
97 first-order stationary point of (6) if there exists a $y \in \mathbb{R}^m$ such that

$$\begin{cases} \text{dist}(-\nabla_x \mathcal{L}_\beta(x, y), \partial g(x)) \leq \epsilon_f \\ \|A(x)\| \leq \epsilon_f, \end{cases} \quad (11)$$

98 for $\epsilon_f \geq 0$. In light of (11), a suitable metric for evaluating the stationarity of a pair $(x, y) \in \mathbb{R}^d \times \mathbb{R}^m$
99 is

$$\text{dist}(-\nabla_x \mathcal{L}_\beta(x, y), \partial g(x)) + \|A(x)\|, \quad (12)$$

100 which we use as the first-order stopping criterion. As an example, for a convex set $\mathcal{X} \subset \mathbb{R}^d$, suppose
 101 that $g = \delta_{\mathcal{X}}$ is the indicator function on \mathcal{X} . Let also $T_{\mathcal{X}}(x) \subseteq \mathbb{R}^d$ denote the tangent cone to \mathcal{X} at x ,
 102 and with $P_{T_{\mathcal{X}}(x)} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ we denote the orthogonal projection onto this tangent cone. Then, for
 103 $u \in \mathbb{R}^d$, it is not difficult to verify that

$$\text{dist}(u, \partial g(x)) = \|P_{T_{\mathcal{X}}(x)}(u)\|. \quad (13)$$

104 When $g = 0$, a first-order stationary point $x \in \mathbb{R}^d$ of (1) is also second-order stationary if

$$\lambda_{\min}(\nabla_{xx}\mathcal{L}_{\beta}(x, y)) \geq 0, \quad (14)$$

105 where $\nabla_{xx}\mathcal{L}_{\beta}$ is the Hessian of \mathcal{L}_{β} with respect to x , and $\lambda_{\min}(\cdot)$ returns the smallest eigenvalue of
 106 its argument. Analogously, x is an $(\epsilon_f, \epsilon_s, \beta)$ second-order stationary point if, in addition to (11), it
 107 holds that

$$\lambda_{\min}(\nabla_{xx}\mathcal{L}_{\beta}(x, y)) \geq -\epsilon_s, \quad (15)$$

108 for $\epsilon_s \geq 0$. Naturally, for second-order stationarity, we use $\lambda_{\min}(\nabla_{xx}\mathcal{L}_{\beta}(x, y))$ as the stopping
 109 criterion.

110 **Smoothness lemma.** This next result controls the smoothness of $\mathcal{L}_{\beta}(\cdot, y)$ for a fixed y . The proof
 111 is standard but nevertheless is included in Appendix C for completeness.

112 **Lemma 2.1 (smoothness)** *For fixed $y \in \mathbb{R}^m$ and $\rho, \rho' \geq 0$, it holds that*

$$\|\nabla_x \mathcal{L}_{\beta}(x, y) - \nabla_x \mathcal{L}_{\beta}(x', y)\| \leq \lambda_{\beta} \|x - x'\|, \quad (16)$$

113 *for every $x, x' \in \{x'' : \|x''\| \leq \rho, \|A(x'')\| \leq \rho'\}$, where*

$$\begin{aligned} \lambda_{\beta} &\leq \lambda_f + \sqrt{m}\lambda_A \|y\| + (\sqrt{m}\lambda_A \rho' + d\lambda_A^2)\beta \\ &=: \lambda_f + \sqrt{m}\lambda_A \|y\| + \lambda''(A, \rho, \rho')\beta. \end{aligned} \quad (17)$$

114 *Above, λ_f, λ_A were defined in (5) and*

$$\lambda'_A := \max_{\|x\| \leq \rho} \|DA(x)\|. \quad (18)$$

115 3 Algorithm

116 To solve the equivalent formulation of (1) presented in (6), we propose the inexact ALM (iALM),
 117 detailed in Algorithm 1.

118 At the k^{th} iteration, Step 2 of Algorithm 1 calls a solver that finds an approximate stationary point
 119 of the augmented Lagrangian $\mathcal{L}_{\beta_k}(\cdot, y_k)$ with the accuracy of ϵ_{k+1} , and this accuracy gradually
 120 increases in a controlled fashion.

121 The increasing sequence of penalty weights $\{\beta_k\}_k$ and the dual update (Steps 4 and 5) are responsible
 122 for continuously enforcing the constraints in (1). The appropriate choice for $\{\beta_k\}_k$ will be specified
 123 in Sections 4.1 and 4.2.

124 The particular choice of the dual step sizes $\{\sigma_k\}_k$ in Algorithm 1 ensures that the dual variable y_k
 125 remains bounded, see [3] for a precedent in the ALM literature where a similar dual step size is
 126 considered.

127 4 Convergence Rate

128 In this section, we derive the total iteration complexity of Algorithm 1 for finding first-order and
 129 second-order stationary points of problem (1). All the proofs are deferred to Appendix A. Theorem 4.1
 130 below characterizes the convergence rate of Algorithm 1 for finding stationary points in terms of the
 131 number of outer iterations.

132

Algorithm 1 Inexact ALM for solving (1)

Input: Non-decreasing, positive, unbounded sequence $\{\beta_k\}_{k \geq 1}$, stopping thresholds $\tau_f > 0$ and $\tau_s > 0$.

Initialization: Initial primal variable $x_1 \in \mathbb{R}^d$, initial dual variable $y_0 \in \mathbb{R}^m$, initial dual step size $\sigma_1 > 0$.

for $k = 1, 2, \dots$ **do**

1. **(Update tolerance)** $\epsilon_{k+1} = 1/\beta_k$.
2. **(Inexact primal solution)** Obtain $x_{k+1} \in \mathbb{R}^d$ such that

$$\text{dist}(-\nabla_x \mathcal{L}_{\beta_k}(x_{k+1}, y_k), \partial g(x_{k+1})) \leq \epsilon_{k+1}$$

for first-order stationarity

$$\lambda_{\min}(\nabla_{xx} \mathcal{L}_{\beta_k}(x_{k+1}, y_k)) \geq -\epsilon_{k+1}$$

for second-order-stationarity, if $g = 0$ in (1).

3. **(Update dual step size)**

$$\sigma_{k+1} = \sigma_1 \min \left(\frac{\|A(x_1)\| \log^2 2}{\|A(x_{k+1})\| (k+1) \log^2(k+2)}, 1 \right).$$

4. **(Dual ascent)** $y_{k+1} = y_k + \sigma_{k+1} A(x_{k+1})$.
5. **(Stopping criterion)** If

$$\text{dist}(-\nabla_x \mathcal{L}_{\beta_k}(x_{k+1}), \partial g(x_{k+1})) + \|A(x_{k+1})\| \leq \tau_f,$$

for first-order stationarity and if also $\lambda_{\min}(\nabla_{xx} \mathcal{L}_{\beta_k}(x_{k+1}, y_k)) \geq -\tau_s$ for second-order stationarity, then quit and return x_{k+1} as an (approximate) stationary point of (1).

end for

133 **Theorem 4.1 (convergence rate)** For integers $2 \leq k_0 \leq k_1$, consider the interval $K = [k_0 :$
 134 $k_1]$, and let $\{x_k\}_{k \in K}$ be the output sequence of Algorithm 1 on the interval K .² Let also $\rho :=$
 135 $\sup_{k \in [K]} \|x_k\|$.³ Suppose that f and A satisfy (5) and let

$$\lambda'_f = \max_{\|x\| \leq \rho} \|\nabla f(x)\|, \quad \lambda'_A = \max_{\|x\| \leq \rho} \|DA(x)\|, \quad (19)$$

136 be the (restricted) Lipschitz constants of f and A , respectively. With $\nu > 0$, assume that

$$\nu \|A(x_k)\| \leq \text{dist} \left(-DA(x_k)^\top A(x_k), \frac{\partial g(x_k)}{\beta_{k-1}} \right), \quad (20)$$

137 for every $k \in K$. We consider two cases:

- 138 • If a first-order solver is used in Step 2, then x_k is an $(\epsilon_{k,f}, \beta_k)$ first-order stationary point of (1)
 139 with

$$\begin{aligned} \epsilon_{k,f} &= \frac{1}{\beta_{k-1}} \left(\frac{2(\lambda'_f + \lambda'_A y_{\max})(1 + \lambda'_A \sigma_k)}{\nu} + 1 \right) \\ &=: \frac{Q(f, g, A, \sigma_1)}{\beta_{k-1}}, \end{aligned} \quad (21)$$

140 for every $k \in K$, where $y_{\max}(x_1, y_0, \sigma_1)$ is specified in (41) due to the limited space.

²The choice of $k_1 = \infty$ is valid here too.

³If necessary, to ensure that $\rho < \infty$, one can add a small factor of $\|x\|^2$ to \mathcal{L}_β in (7). Then it is easy to verify that the iterates of Algorithm 1 remain bounded, provided that the penalty weight β is large enough, $\sup_x \|\nabla f(x)\|/\|x\| < \infty$, $\sup_x \|A(x)\| < \infty$, and $\sup_x \|DA(x)\| < \infty$.

- 141 • If a second-order solver is used in Step 2, then x_k is an $(\epsilon_{k,f}, \epsilon_{k,s}, \beta_k)$ second-order stationary
 142 point of (1) with $\epsilon_{k,s}$ specified above and with

$$\begin{aligned} \epsilon_{k,s} &= \epsilon_{k-1} + \sigma_k \sqrt{m} \lambda_A \frac{2\lambda'_f + 2\lambda'_A y_{\max}}{\nu \beta_{k-1}} \\ &= \frac{\nu + \sigma_k \sqrt{m} \lambda_A 2\lambda'_f + 2\lambda'_A y_{\max}}{\nu \beta_{k-1}} =: \frac{Q'(f, g, A, \sigma_1)}{\beta_{k-1}}. \end{aligned} \quad (22)$$

143 Loosely speaking, Theorem 4.1 states that Algorithm 1 converges to a (first- or second-) order
 144 stationary point of (1) at the rate of $1/\beta_k$, further specified in Sections 4.1 and 4.2. A few remarks
 145 are in order about Theorem 4.1.

146 **Regularity.** The key geometric condition in Theorem 4.1 is (20) which, broadly speaking, ensures
 147 that the primal updates of Algorithm 1 reduce the feasibility gap as the penalty weight β_k grows. We
 148 will verify this condition for several examples in Section 6.

149 This condition in (20) is closely related to those in the existing literature. In the special case where
 150 $g = 0$ in (1), it is easy to verify that (20) reduces to the Polyak-Lojasiewicz (PL) condition for
 151 minimizing $\|A(x)\|^2$ [34]. PL condition itself is a special case of Kurdyka-Lojasiewicz with $\theta = 1/2$,
 152 see [64, Definition 1.1]. When $g = 0$, it is also easy to see that (20) is weaker than the Mangasarian-
 153 Fromovitz (MF) condition in nonlinear optimization [11, Assumption 1]. Moreover, **when g is**
 154 **the indicator on a convex set**, (20) is a consequence of the *basic constraint qualification* in [54],
 155 which itself generalizes the MF condition to the case when g is an indicator function of a convex
 156 set. **Note:AE: I'm not sure if the claim about basic constraint qualification is true because our**
 157 **condition should locally hold rather globally. Could you add the exact equation number in**
 158 **[55] and double check this? Also consider double checking other claims in this paragraph.**

159 Loosely speaking, we may think of (20) as a local condition, which should hold within a neighborhood
 160 of the constraint set $\{x : A(x) = 0\}$ rather than everywhere in \mathbb{R}^d . There is a constant complexity
 161 algorithm in [11] to reach this so-called “information zone”, which supplements Theorem 4.1. Lastly,
 162 in contrast to most conditions in the nonconvex optimization literature, such as [27], the condition
 163 in (20) appears to be easier to verify, as we see in Section 6.

164 **AE: the spars review talks about the "Pong-Li" work. Fatih, do you know what is that?**

165 (mfs:) I do not think this work is relevant to our condition but the template seems similar. We can
 166 mention in the related works instead. **Note:AE: Ok. Yes, consider adding it to the related work.**

167 **Penalty method.** A classical algorithm to solve (1) is the penalty method, which is characterized
 168 by the absence of the dual variable ($y = 0$) in (7). Indeed, ALM can be interpreted as an adaptive
 169 penalty or smoothing method with a variable center determined by the dual variable. It is worth noting
 170 that, with the same proof technique, one can establish the same convergence rate of Theorem 4.1 for
 171 the penalty method. However, while both methods have the same convergence rate in theory, iALM
 172 outperforms the penalty method in practice by virtue of its variable center and has been excluded
 173 from this presentation for this reason.

174 **Computational complexity.** Theorem 4.1 specifies the number of (outer) iterations that Algorithm 1
 175 requires to reach a near-stationary point of problem (1) with a prescribed precision and, in particular,
 176 specifies the number of calls made to the solver in Step 2. In this sense, Theorem 4.1 does not
 177 fully capture the computational complexity of Algorithm 1, as it does not take into account the
 178 computational cost of the solver in Step 2.

179 To better understand the total iteration complexity of Algorithm 1, we consider two scenarios in the
 180 following. In the first scenario, we take the solver in Step 2 to be the Accelerated Proximal Gradient
 181 Method (APGM), a well-known first-order algorithm [29]. In the second scenario, we will use the
 182 second-order trust region method developed in [19].

183 4.1 First-Order Optimality

184 Let us first consider the case where the solver in Step 2 is the first-order algorithm APGM, described
 185 in detail in [29]. At a high level, APGM makes use of $\nabla_x \mathcal{L}_\beta(x, y)$ in (7), the proximal operator
 186 prox_g , and the classical Nesterov acceleration [44] to reach first-order stationarity for the subproblem

187 in (8). Suppose that $g = \delta_{\mathcal{X}}$ is the indicator function on a bounded convex set $\mathcal{X} \subset \mathbb{R}^d$ and let

$$\rho = \max_{x \in \mathcal{X}} \|x\|, \quad (23)$$

188 be the radius of a ball centered at the origin that includes \mathcal{X} . Then, adapting the results in [29] to our
189 setup, APMG reaches x_k in Step 2 of Algorithm 1 after

$$\mathcal{O} \left(\frac{\lambda_{\beta_k}^2 \rho^2}{\epsilon_{k+1}} \right) \quad (24)$$

190 (inner) iterations, where λ_{β_k} denotes the Lipschitz constant of $\nabla_x \mathcal{L}_{\beta_k}(x, y)$, bounded in (17). For
191 the clarity of the presentation, we have used a looser bound in (24) compared to [29]. Using (24), we
192 derive the following corollary, describing the total iteration complexity of Algorithm 1 in terms of the
193 number calls made to the first-order oracle in APMG.

194 **Corollary 4.2** *For $b > 1$, let $\beta_k = b^k$ for every k . If we use APMG from [29] for Step 2 of*
195 *Algorithm 1, the algorithm finds an (ϵ_f, β_k) first-order stationary point, after T calls to the first-order*
196 *oracle, where*

$$T = \mathcal{O} \left(\frac{Q^3 \rho^2}{\epsilon^3} \log_b \left(\frac{Q}{\epsilon} \right) \right) = \tilde{\mathcal{O}} \left(\frac{Q^3 \rho^2}{\epsilon^3} \right). \quad (25)$$

197 For Algorithm 1 to reach a near-stationary point with an accuracy of ϵ_f in the sense of (11) and
198 with the lowest computational cost, we therefore need to perform only one iteration of Algorithm 1,
199 with β_1 specified as a function of ϵ_f by (21) in Theorem 4.1. In general, however, the constants in
200 (21) are unknown and this approach is thus not feasible. Instead, the homotopy approach taken by
201 Algorithm 1 ensures achieving the desired accuracy by gradually increasing the penalty weight.⁴ This
202 homotopy approach increases the computational cost of Algorithm 1 only by a factor logarithmic in
203 the ϵ_f , as detailed in the proof of Corollary 4.2.

204 4.2 Second-Order Optimality

205 Let us now consider the second-order optimality case where the solver in Step 2 is the trust region
206 method developed in [19]. Trust region method minimizes a quadratic approximation of the function
207 within a dynamically updated trust-region radius. Second-order trust region method that we consider
208 in this section makes use of Hessian (or an approximation of Hessian) of the augmented Lagrangian
209 in addition to first order oracles.

210 As shown in [46], finding approximate second-order stationary points of convex-constrained problems
211 is in general NP-hard. For this reason, we focus in this section on the special case of (1) with $g = 0$.

212 Let us compute the total computational complexity of Algorithm 1 with the trust region method in
213 Step 2, in terms of the number of calls made to the second-order oracle. By adapting the result in [19]
214 to our setup, we find that the number of (inner) iterations required in Step 2 of Algorithm 1 to produce
215 x_{k+1} is

$$\mathcal{O} \left(\frac{\lambda_{\beta_k, H}^2 (\mathcal{L}_{\beta_k}(x_1, y) - \min_x \mathcal{L}_{\beta_k}(x, y))}{\epsilon_k^3} \right), \quad (26)$$

216 where $\lambda_{\beta, H}$ is the Lipschitz constant of the Hessian of the augmented Lagrangian, which is of the
217 order of β , as can be proven similar to Lemma 2.1 and x_1 is the initial iterate of the given outer loop.
218 In [19], the term $\mathcal{L}_{\beta}(x_1, y) - \min_x \mathcal{L}_{\beta}(x, y)$ is bounded by a constant independent of ϵ . We assume
219 a uniform bound for this quantity for every β_k , instead of for one value of β_k as in [19]. Using (26)
220 and Theorem 4.1, we arrive at the following:

221 **Corollary 4.3** *For $b > 1$, let $\beta_k = b^k$ for every k . We assume that*

$$\mathcal{L}_{\beta}(x_1, y) - \min_x \mathcal{L}_{\beta}(x, y) \leq L_u, \quad \forall \beta. \quad (27)$$

222 *If we use the trust region method from [19] for Step 2 of Algorithm 1, the algorithm finds an*
223 *ϵ -second-order stationary point of (1) in T calls to the second-order oracle where*

$$T = \mathcal{O} \left(\frac{L_u Q'^5}{\epsilon^5} \log_b \left(\frac{Q'}{\epsilon} \right) \right) = \tilde{\mathcal{O}} \left(\frac{L_u Q'^5}{\epsilon^5} \right). \quad (28)$$

⁴In this context, homotopy loosely corresponds to the gradual enforcement of the constraints by increasing the penalty weight.

224 **Note:AE: I can't remember: what is Q' and why isn't it defined?** Before closing this section, we
 225 note that the remark after Corollary 4.2 applies here as well.

226 5 Related Work

227 ALM has a long history in the optimization literature, dating back to [31, 52]. In the special case
 228 of (1) with a convex function f and a linear operator A , standard, inexact, and linearized versions of
 229 ALM have been extensively studied [37, 42, 59, 63].

230 Classical works on ALM focused on the general template of (1) with nonconvex f and nonlinear A ,
 231 with arguably stronger assumptions and required exact solutions to the subproblems of the form (8),
 232 which appear in Step 2 of Algorithm 1, see for instance [5].

233 A similar analysis was conducted in [25] for the general template of (1). The authors considered
 234 inexact ALM and proved convergence rates for the outer iterates, under specific assumptions on
 235 the initialization of the dual variable. However, unlike our results, the authors did not analyze how
 236 to solve the subproblems inexactly and they did not provide total complexity results and verifiable
 237 conditions.

238 Problem (1) with similar assumptions to us is also studied in [8] and [20] for first-order and second-
 239 order stationarity, respectively, with explicit iteration complexity analysis. As we have mentioned
 240 in Section 4, our iteration complexity results matches these theoretical algorithms with a simpler
 241 algorithm and a simpler analysis. In addition, these algorithms require setting final accuracies
 242 since they utilize this information in the algorithm. In contrast to [8, 20], Algorithm 1 does not set
 243 accuracies a priori.

244 [18] also considers the same template (1) for first-order stationarity with a penalty-type method
 245 instead of ALM. Even though the authors show $\mathcal{O}(1/\epsilon^2)$ complexity, this result is obtained by
 246 assuming that the penalty parameter remains bounded. We note that such an assumption can also be
 247 used to match our complexity results.

248 [11] studies the general template (1) with specific assumptions involving local error bound conditions
 249 for the (1). These conditions are studied in detail in [10], but their validity for general SDPs (2) has
 250 never been established. This work also lacks the total iteration complexity analysis presented here.

251 Another work [23] focused on solving (1) by adapting the primal-dual method of Chambolle and
 252 Pock [21]. The authors proved the convergence of the method and provided convergence rate by
 253 imposing error bound conditions on the objective function that do not hold for standard SDPs.

254 [15, 16] is the first work that proposes the splitting $X = UU^\top$ for solving SDPs of the form (2).
 255 Following these works, the literature on Burer-Monteiro (BM) splitting for the large part focused on
 256 using ALM for solving the reformulated problem (3).

257 However, this approach has a few drawbacks: First, it requires exact solutions in Step 2 of Algo-
 258 rithm 1 in theory, which in practice is replaced with inexact solutions. Second, their results only
 259 establish convergence without providing the rates. In this sense, our work provides a theoretical
 260 understanding of the BM splitting with inexact solutions to Step 2 of Algorithm 1 and complete
 261 iteration complexities.

262 [7, 49] are among the earliest efforts to show convergence rates for BM splitting, focusing on
 263 the special case of SDPs without any linear constraints. For these specific problems, they prove
 264 the convergence of gradient descent to global optima with convergence rates, assuming favorable
 265 initialization. These results, however, do not apply to general SDPs of the form (2) where the difficulty
 266 arises due to the linear constraints.

267 [6] focused on the quadratic penalty formulation of (1), namely,

$$\min_{X \succeq 0} \langle C, X \rangle + \frac{\mu}{2} \|B(x) - b\|^2, \quad (29)$$

268 which after BM splitting becomes

$$\min_{U \in \mathbb{R}^{d \times r}} \langle C, UU^\top \rangle + \frac{\mu}{2} \|B(UU^\top) - b\|^2, \quad (30)$$

269 for which they study the optimality of the second-order stationary points. These results are for
 270 establishing a connection between the stationary points of (30) and global optima of (29). In contrast,
 271 we focus on the relation of the stationary points of (6) to the constrained problem (1).

272 Another popular method for solving SDPs are due to [13, 12, 14], focusing on the case where the
 273 constraints in (1) can be written as a Riemannian manifold after BM splitting. In this case, the authors
 274 apply the Riemannian gradient descent and Riemannian trust region methods for obtaining first- and
 275 second-order stationary points, respectively. They obtain $\mathcal{O}(1/\epsilon^2)$ complexity for finding first-order
 276 stationary points and $\mathcal{O}(1/\epsilon^3)$ complexity for finding second-order stationary points.

277 While these complexities appear better than ours, the smooth manifold requirement in these works
 278 is indeed restrictive. In particular, this requirement holds for max-cut and generalized eigenvalue
 279 problems, but it is not satisfied for other important SDPs such as quadratic programming (QAP),
 280 optimal power flow and clustering with general affine constraints. In addition, as noted in [12], per
 281 iteration cost of their method for max-cut problem is an astronomical $\mathcal{O}(d^6)$.

282 Lastly, there also exists a line of work for solving SDPs in their original convex formulation, in a
 283 storage efficient way [43, 66, 67]. These works have global optimality guarantees by their virtue of
 284 directly solving the convex formulation. On the downside, these works require the use of eigenvalue
 285 routines and exhibit significantly slower convergence as compared to nonconvex approaches [33].

286 6 Numerical Evidence

287 We first begin with a caveat: It is known that quasi-Newton methods, such as BFGS and LBFGS,
 288 might not converge for nonconvex problems [24, 39]. For this reason, we have used the trust region
 289 method as the second-order solver in our analysis in Section 4, which is well-studied for nonconvex
 290 problems [19]. Empirically, however, BFGS and LBFGS are extremely successful and we have
 291 therefore opted for those solvers in this section since the subroutine does not affect Theorem 4.1 as
 292 long as the subsolver performs well in practice.

293 6.1 Clustering

294 Given data points $\{z_i\}_{i=1}^n$, the entries of the corresponding Euclidean distance matrix $D \in \mathbb{R}^{n \times n}$
 295 are $D_{i,j} = \|z_i - z_j\|^2$. Clustering is then the problem of finding a co-association matrix $Y \in \mathbb{R}^{n \times n}$
 296 such that $Y_{ij} = 1$ if points z_i and z_j are within the same cluster and $Y_{ij} = 0$ otherwise. In [51], the
 297 authors provide a SDP relaxation of the clustering problem, specified as

$$\begin{cases} \min_{Y \in \mathbb{R}^{n \times n}} \text{tr}(DY) \\ Y\mathbf{1} = \mathbf{1}, \text{tr}(Y) = s, Y \succeq 0, Y \geq 0, \end{cases} \quad (31)$$

298 where s is the number of clusters and Y is both positive semidefinite and has nonnegative entries.
 299 Standard SDP solvers do not scale well with the number of data points n , since they often require
 300 projection onto the semidefinite cone with the complexity of $\mathcal{O}(n^3)$. We instead use the BM
 301 factorization to solve (31), sacrificing convexity to reduce the computational complexity. More
 302 specifically, we solve the program

$$\begin{cases} \min_{V \in \mathbb{R}^{n \times r}} \text{tr}(DVV^\top) \\ VV^\top \mathbf{1} = \mathbf{1}, \|V\|_F^2 \leq s, V \geq 0, \end{cases} \quad (32)$$

303 where $\mathbf{1} \in \mathbb{R}^n$ is the vector of all ones. Note that $Y \geq 0$ in (31) is replaced above by the much
 304 stronger but easier-to-enforce constraint $V \geq 0$ in (32), see [36] for the reasoning behind this
 305 relaxation. Now, we can cast (32) as an instance of (1). Indeed, for every $i \leq n$, let $x_i \in \mathbb{R}^r$ denote
 306 the i th row of V . We next form $x \in \mathbb{R}^d$ with $d = nr$ by expanding the factorized variable V , namely,

$$x = [x_1^\top, \dots, x_n^\top]^\top \in \mathbb{R}^d,$$

307 and then set

$$f(x) = \sum_{i,j=1}^n D_{i,j} \langle x_i, x_j \rangle, \quad g = \delta_C,$$




Figure 1: Convergence of different algorithms for k-Means clustering with fashion MNIST dataset. The solution rank for the template (31) is known and it is equal to number of clusters k (Theorem 1. [36]). As discussed in [58], setting rank $r > k$ leads more accurate reconstruction in expense of speed. Therefore, we set the rank to 20.

308

$$A(x) = [x_1^\top \sum_{j=1}^n x_j - 1, \dots, x_n^\top \sum_{j=1}^n x_j - 1]^\top, \quad (33)$$

309 where C is the intersection of the positive orthant in \mathbb{R}^d with the Euclidean ball of radius \sqrt{s} . In
 310 Appendix D, we somewhat informally verify that Theorem 4.1 applies to (1) with f, g, A specified
 311 above.

312 In our simulations, we use two different solvers for Step 2 of Algorithm 1, namely, APGM and
 313 IBFGS. APGM is a solver for nonconvex problems of the form (8) with convergence guarantees
 314 to first-order stationarity, as discussed in Section 4. IBFGS is a limited-memory version of BFGS
 315 algorithm in [26] that approximately leverages the second-order information of the problem. We
 316 compare our approach against the following convex methods:

- 317 • HCGM: Homotopy-based Conditional Gradient Method in [67] which directly solves (31).
- 318 • SDPNAL+: A second-order augmented Lagrangian method for solving SDP's with nonneg-
 319 ativity constraints [65].

320 As for the dataset, our experimental setup is similar to that described by [40]. We use the publicly-
 321 available fashion-MNIST data in [62], which is released as a possible replacement for the MNIST
 322 handwritten digits. Each data point is a 28×28 gray-scale image, associated with a label from ten
 323 classes, labeled from 0 to 9. First, we extract the meaningful features from this dataset using a simple
 324 two-layer neural network with a sigmoid activation function. Then, we apply this neural network to
 325 1000 test samples from the same dataset, which gives us a vector of length 10 for each data point,
 326 where each entry represents the posterior probability for each class. Then, we form the ℓ_2 distance
 327 matrix D from these probability vectors. The results are depicted in Figure 1. We implemented 3
 328 algorithms on MATLAB and used the software package for SDPNAL+ which contains mex files. It is
 329 predictable that the performance of our nonconvex approach would even improve by using mex files.

330 6.2 Basis Pursuit

331 Basis Pursuit (BP) finds sparsest solutions of an under-determined system of linear equations by
 332 solving

$$\begin{cases} \min_z \|z\|_1 \\ Bz = b, \end{cases} \quad (34)$$

333 where $B \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^n$. BP has found many applications in machine learning, statistics and
 334 signal processing [22, 17, 1]. Various primal-dual convex optimization algorithms are available in

335 the literature to solve BP, including [59, 21]. We compare our algorithm against state-of-the-art
 336 primal-dual convex methods for solving (34), namely, Chambolle-Pock [21], ASGARD [60] and
 337 ASGARD-DL [59].

338 Here, we take a different approach and cast (34) as an instance of (1). Note that any $z \in \mathbb{R}^d$
 339 can be decomposed as $z = z^+ - z^-$, where $z^+, z^- \in \mathbb{R}^d$ are the positive and negative parts of
 340 z , respectively. Then consider the change of variables $z^+ = u_1^{\circ 2}$ and $z^- = u_2^{\circ 2} \in \mathbb{R}^d$, where \circ
 341 denotes element-wise power. Next, we concatenate u_1 and u_2 as $x := [u_1^\top, u_2^\top]^\top \in \mathbb{R}^{2d}$ and define
 342 $\bar{B} := [B, -B] \in \mathbb{R}^{n \times 2d}$. Then, (34) is equivalent to (1) with

$$\begin{aligned} f(x) &= \|x\|^2, & g(x) &= 0, \\ A(x) &= \bar{B}x^{\circ 2} - b. \end{aligned} \tag{35}$$

343 In Appendix E, we verify with minimal details that Theorem 4.1 indeed applies to (1) with the above
 344 f, A .

345 We draw the entries of B independently from a zero-mean and unit-variance Gaussian distribution.
 346 For a fixed sparsity level k , the support of $z_* \in \mathbb{R}^d$ and its nonzero amplitudes are also drawn from
 347 the standard Gaussian distribution. Then the measurement vector is created as $b = Bz + \epsilon$, where ϵ
 348 is the noise vector with entries drawn independently from the zero-mean Gaussian distribution with
 349 variance $\sigma^2 = 10^{-6}$.




Figure 2: Convergence with different subsolvers for the aforementioned nonconvex relaxation.

350 The results are compiled in Figure 2. Clearly, the performance of Algorithm 1 with a second-order
 351 solver for BP is comparable to the rest. It is, indeed, interesting to see that these type of nonconvex
 352 relaxations gives the solution of convex one and first order methods succeed.

353 **Discussion:** The true potential of our reformulation is in dealing with more structured norms rather
 354 than ℓ_1 , where computing the proximal operator is often intractable. One such case is the latent group
 355 lasso norm [47], defined as

$$\|z\|_{\Omega} = \sum_{i=1}^I \|z_{\Omega_i}\|,$$

356 where $\{\Omega_i\}_{i=1}^I$ are (not necessarily disjoint) index sets of $\{1, \dots, d\}$. Although not studied here, we
 357 believe that the nonconvex framework presented in this paper can serve to solve more complicated
 358 problems, such as the latent group lasso. We leave this research direction for future work.

359 References

- 360 [1] S. Arora, M. Khodak, N. Saunshi, and K. Vodrahalli. A compressed sensing view of unsuper-
 361 vised text embeddings, bag-of-n-grams, and lstms. 2018.
- 362 [2] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps.
 363 *Discrete & Computational Geometry*, 13(2):189–202, 1995.

- 364 [3] D. P. Bertsekas. On penalty and multiplier methods for constrained minimization. *SIAM Journal*
365 *on Control and Optimization*, 14(2):216–235, 1976.
- 366 [4] D. P. Bertsekas. Constrained optimization and lagrange multiplier methods. *Computer Science*
367 *and Applied Mathematics, Boston: Academic Press, 1982, 1982.*
- 368 [5] D. P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press,
369 2014.
- 370 [6] S. Bhojanapalli, N. Boumal, P. Jain, and P. Netrapalli. Smoothed analysis for low-rank solutions
371 to semidefinite programs in quadratic penalty form. *arXiv preprint arXiv:1803.00186*, 2018.
- 372 [7] S. Bhojanapalli, A. Kyriillidis, and S. Sanghavi. Dropping convexity for faster semi-definite
373 optimization. In *Conference on Learning Theory*, pages 530–582, 2016.
- 374 [8] E. G. Birgin, J. Gardenghi, J. M. Martinez, S. Santos, and P. L. Toint. Evaluation complexity
375 for nonlinear constrained optimization using unscaled kkt conditions and high-order models.
376 *SIAM Journal on Optimization*, 26(2):951–967, 2016.
- 377 [9] E. G. Birgin and J. M. Mart_nez. *Practical augmented Lagrangian methods for constrained*
378 *optimization*, volume 10. SIAM, 2014.
- 379 [10] J. Bolte, T. P. Nguyen, J. Peypouquet, and B. W. Suter. From error bounds to the complexity of
380 first-order descent methods for convex functions. *Mathematical Programming*, 165(2):471–507,
381 2017.
- 382 [11] J. Bolte, S. Sabach, and M. Teboulle. Nonconvex lagrangian-based optimization: monitoring
383 schemes and global convergence. *Mathematics of Operations Research*, 2018.
- 384 [12] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization
385 on manifolds. *arXiv preprint arXiv:1605.08101*, 2016.
- 386 [13] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a matlab toolbox for optimization
387 on manifolds. *The Journal of Machine Learning Research*, 15(1):1455–1459, 2014.
- 388 [14] N. Boumal, V. Voroninski, and A. Bandeira. The non-convex burer-monteiro approach works on
389 smooth semidefinite programs. In *Advances in Neural Information Processing Systems*, pages
390 2757–2765, 2016.
- 391 [15] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite
392 programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- 393 [16] S. Burer and R. D. Monteiro. Local minima and convergence in low-rank semidefinite program-
394 ming. *Mathematical Programming*, 103(3):427–444, 2005.
- 395 [17] E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *IEEE signal processing*
396 *magazine*, 25(2):21–30, 2008.
- 397 [18] C. Cartis, N. I. Gould, and P. L. Toint. On the evaluation complexity of composite func-
398 tion minimization with applications to nonconvex nonlinear programming. *SIAM Journal on*
399 *Optimization*, 21(4):1721–1739, 2011.
- 400 [19] C. Cartis, N. I. Gould, and P. L. Toint. Complexity bounds for second-order optimality in
401 unconstrained optimization. *Journal of Complexity*, 28(1):93–108, 2012.
- 402 [20] C. Cartis, N. I. Gould, and P. L. Toint. Optimality of orders one to three and beyond: characteri-
403 zation and evaluation complexity in constrained nonconvex optimization. *Journal of Complexity*,
404 2018.
- 405 [21] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with
406 applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- 407 [22] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM*
408 *review*, 43(1):129–159, 2001.

- 409 [23] C. Clason, S. Mazurenko, and T. Valkonen. Acceleration and global convergence of a first-order
410 primal–dual method for nonconvex problems. *arXiv preprint arXiv:1802.03347*, 2018.
- 411 [24] Y.-H. Dai. Convergence properties of the bfgs algorithm. *SIAM Journal on Optimization*,
412 13(3):693–701, 2002.
- 413 [25] D. Fernandez and M. V. Solodov. Local convergence of exact and inexact augmented lagrangian
414 methods under the second-order sufficient optimality condition. *SIAM Journal on Optimization*,
415 22(2):384–407, 2012.
- 416 [26] R. Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2013.
- 417 [27] F. Flores-Bazán, F. Flores-Bazán, and C. Vera. A complete characterization of strong duality in
418 nonconvex optimization with a single constraint. *Journal of Global Optimization*, 53(2):185–
419 201, 2012.
- 420 [28] R. Ge, C. Jin, P. Netrapalli, A. Sidford, et al. Efficient algorithms for large-scale generalized
421 eigenvector computation and canonical correlation analysis. In *International Conference on*
422 *Machine Learning*, pages 2741–2750, 2016.
- 423 [29] S. Ghadimi and G. Lan. Accelerated gradient methods for nonconvex nonlinear and stochastic
424 programming. *Mathematical Programming*, 156(1-2):59–99, 2016.
- 425 [30] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville,
426 and Y. Bengio. Generative Adversarial Networks. *ArXiv e-prints*, June 2014.
- 427 [31] M. R. Hestenes. Multiplier and gradient methods. *Journal of optimization theory and applica-*
428 *tions*, 4(5):303–320, 1969.
- 429 [32] A. Ilyas, A. Jalal, E. Asteri, C. Daskalakis, and A. G. Dimakis. The Robust Manifold Defense:
430 Adversarial Training using Generative Models. *arXiv e-prints*, page arXiv:1712.09196, Dec.
431 2017.
- 432 [33] M. Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*,
433 pages 427–435, 2013.
- 434 [34] H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient
435 methods under the polyak-jojasiewicz condition. In *Joint European Conference on Machine*
436 *Learning and Knowledge Discovery in Databases*, pages 795–811. Springer, 2016.
- 437 [35] S. Khot and A. Naor. Grothendieck-type inequalities in combinatorial optimization. *arXiv*
438 *preprint arXiv:1108.2464*, 2011.
- 439 [36] B. Kulis, A. C. Surendran, and J. C. Platt. Fast low-rank semidefinite programming for
440 embedding and clustering. In *Artificial Intelligence and Statistics*, pages 235–242, 2007.
- 441 [37] G. Lan and R. D. Monteiro. Iteration-complexity of first-order augmented lagrangian methods
442 for convex programming. *Mathematical Programming*, 155(1-2):511–547, 2016.
- 443 [38] L. Lovász. Semidefinite programs and combinatorial optimization. In *Recent advances in*
444 *algorithms and combinatorics*, pages 137–194. Springer, 2003.
- 445 [39] W. F. Mascarenhas. The bfgs method with exact line searches fails for non-convex objective
446 functions. *Mathematical Programming*, 99(1):49–61, 2004.
- 447 [40] D. G. Mixon, S. Villar, and R. Ward. Clustering subgaussian mixtures by semidefinite program-
448 ming. *arXiv preprint arXiv:1602.06612*, 2016.
- 449 [41] E. Mossel, J. Neeman, and A. Sly. Consistency thresholds for the planted bisection model. In
450 *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 69–75.
451 ACM, 2015.
- 452 [42] V. Nedelcu, I. Necoara, and Q. Tran-Dinh. Computational complexity of inexact gradient
453 augmented lagrangian methods: application to constrained mpc. *SIAM Journal on Control and*
454 *Optimization*, 52(5):3109–3134, 2014.

- 455 [43] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*,
456 120(1):221–259, 2009.
- 457 [44] Y. E. Nesterov. A method for solving the convex programming problem with convergence rate
458 $o(1/k^2)$. In *Dokl. Akad. Nauk SSSR*, volume 269, pages 543–547, 1983.
- 459 [45] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and
460 Financial Engineering. Springer New York, 2006.
- 461 [46] M. Nouiehed, J. D. Lee, and M. Razaviyayn. Convergence to second-order stationarity for
462 constrained non-convex optimization. *arXiv preprint arXiv:1810.02024*, 2018.
- 463 [47] G. Obozinski, L. Jacob, and J.-P. Vert. Group lasso with overlaps: the latent group lasso
464 approach. *arXiv preprint arXiv:1110.0413*, 2011.
- 465 [48] N. Parikh, S. Boyd, et al. Proximal algorithms. *Foundations and Trends in Optimization*,
466 1(3):127–239, 2014.
- 467 [49] D. Park, A. Kyrillidis, S. Bhojanapalli, C. Caramanis, and S. Sanghavi. Provable burer-monteiro
468 factorization for a class of norm-constrained matrix problems. *arXiv preprint arXiv:1606.01316*,
469 2016.
- 470 [50] G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of
471 optimal eigenvalues. *Mathematics of operations research*, 23(2):339–358, 1998.
- 472 [51] J. Peng and Y. Wei. Approximating K-means-type clustering via semidefinite programming.
473 *SIAM J. Optim.*, 18(1):186–205, 2007.
- 474 [52] M. J. Powell. A method for nonlinear constraints in minimization problems. *Optimization*,
475 pages 283–298, 1969.
- 476 [53] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings*
477 *of the fortieth annual ACM symposium on Theory of computing*, pages 245–254. ACM, 2008.
- 478 [54] R. T. Rockafellar. Lagrange multipliers and optimality. *SIAM review*, 35(2):183–238, 1993.
- 479 [55] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied*
480 *and computational harmonic analysis*, 30(1):20, 2011.
- 481 [56] A. Singer and Y. Shkolnisky. Three-dimensional structure determination from common lines in
482 cryo-em by eigenvectors and semidefinite programming. *SIAM journal on imaging sciences*,
483 4(2):543–572, 2011.
- 484 [57] L. Song, A. Smola, A. Gretton, and K. M. Borgwardt. A dependence maximization view of
485 clustering. In *Proceedings of the 24th international conference on Machine learning*, pages
486 815–822. ACM, 2007.
- 487 [58] M. Tepper, A. M. Sengupta, and D. Chklovskii. Clustering is semidefinitely not that hard:
488 Nonnegative sdp for manifold disentangling. *Journal of Machine Learning Research*, 19(82),
489 2018.
- 490 [59] Q. Tran-Dinh, A. Alacaoglu, O. Fercoq, and V. Cevher. An adaptive primal-dual framework for
491 nonsmooth convex minimization. *arXiv preprint arXiv:1808.04648*, 2018.
- 492 [60] Q. Tran-Dinh, O. Fercoq, and V. Cevher. A smooth primal-dual optimization framework for
493 nonsmooth composite convex minimization. *SIAM Journal on Optimization*, 28(1):96–134,
494 2018.
- 495 [61] I. Waldspurger and A. Waters. Rank optimality for the burer-monteiro factorization. *arXiv*
496 *preprint arXiv:1812.03046*, 2018.
- 497 [62] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking
498 machine learning algorithms, 2017.

- 499 [63] Y. Xu. Iteration complexity of inexact augmented lagrangian methods for constrained convex
500 programming. *arXiv preprint arXiv:1711.05812v2*, 2017.
- 501 [64] Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block
502 coordinate update. *Journal of Scientific Computing*, 72(2):700–734, 2017.
- 503 [65] L. Yang, D. Sun, and K.-C. Toh. Sdpnal+: a majorized semismooth newton-cg augmented
504 lagrangian method for semidefinite programming with nonnegative constraints. *Mathematical*
505 *Programming Computation*, 7(3):331–366, 2015.
- 506 [66] A. Yurtsever, Q. T. Dinh, and V. Cevher. A universal primal-dual convex optimization framework.
507 In *Advances in Neural Information Processing Systems*, pages 3150–3158, 2015.
- 508 [67] A. Yurtsever, O. Fercoq, F. Locatello, and V. Cevher. A conditional gradient framework for
509 composite convex minimization with applications to semidefinite programming. *arXiv preprint*
510 *arXiv:1804.08544*, 2018.
- 511 [68] Q. Zhao, S. E. Karisch, F. Rendl, and H. Wolkowicz. Semidefinite programming relaxations for
512 the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2(1):71–109, 1998.

513 **A Proof of Theorem 4.1**

514 For every $k \geq 2$, recall from (7) and Step 2 of Algorithm 1 that x_k satisfies

$$\begin{aligned} & \text{dist}(-\nabla f(x_k) - DA(x_k)^\top y_{k-1} \\ & \quad - \beta_{k-1} DA(x_k)^\top A(x_k), \partial g(x_k)) \\ & = \text{dist}(-\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1}), \partial g(x_k)) \leq \epsilon_k. \end{aligned} \quad (36)$$

515 With an application of the triangle inequality, it follows that

$$\begin{aligned} & \text{dist}(-\beta_{k-1} DA(x_k)^\top A(x_k), \partial g(x_k)) \\ & \leq \|\nabla f(x_k)\| + \|DA(x_k)^\top y_{k-1}\| + \epsilon_k, \end{aligned} \quad (37)$$

516 which in turn implies that

$$\begin{aligned} & \text{dist}(-DA(x_k)^\top A(x_k), \partial g(x_k)/\beta_{k-1}) \\ & \leq \frac{\|\nabla f(x_k)\|}{\beta_{k-1}} + \frac{\|DA(x_k)^\top y_{k-1}\|}{\beta_{k-1}} + \frac{\epsilon_k}{\beta_{k-1}} \\ & \leq \frac{\lambda'_f + \lambda'_A \|y_{k-1}\| + \epsilon_k}{\beta_{k-1}}, \end{aligned} \quad (38)$$

517 where λ'_f, λ'_A were defined in (19). We next translate (38) into a bound on the feasibility gap $\|A(x_k)\|$.

518 Using the regularity condition (20), the left-hand side of (38) can be bounded below as

$$\text{dist}(-DA(x_k)^\top A(x_k), \partial g(x_k)/\beta_{k-1}) \geq \nu \|A(x_k)\|. \quad (\text{see (20)}) \quad (39)$$

519 By substituting (39) back into (38), we find that

$$\|A(x_k)\| \leq \frac{\lambda'_f + \lambda'_A \|y_{k-1}\| + \epsilon_k}{\nu \beta_{k-1}}. \quad (40)$$

520 In words, the feasibility gap is directly controlled by the dual sequence $\{y_k\}_k$. We next establish that
521 the dual sequence is bounded. Indeed, for every $k \in K$, note that

$$\begin{aligned} \|y_k\| & = \|y_0 + \sum_{i=1}^k \sigma_i A(x_i)\| \quad (\text{Step 5 of Algorithm 1}) \\ & \leq \|y_0\| + \sum_{i=1}^k \sigma_i \|A(x_i)\| \quad (\text{triangle inequality}) \\ & \leq \|y_0\| + \sum_{i=1}^k \frac{\|A(x_1)\| \log^2 2}{k \log^2(k+1)} \quad (\text{Step 4}) \\ & \leq \|y_0\| + c \|A(x_1)\| \log^2 2 =: y_{\max}, \end{aligned} \quad (41)$$

522 where

$$c \geq \sum_{i=1}^{\infty} \frac{1}{k \log^2(k+1)}. \quad (42)$$

523 Substituting (41) back into (40), we reach

$$\begin{aligned} \|A(x_k)\| & \leq \frac{\lambda'_f + \lambda'_A y_{\max} + \epsilon_k}{\nu \beta_{k-1}} \\ & \leq \frac{2\lambda'_f + 2\lambda'_A y_{\max}}{\nu \beta_{k-1}}, \end{aligned} \quad (43)$$

524 where the second line above holds if k_0 is large enough, which would in turn guarantees that
525 $\epsilon_k = 1/\beta_{k-1}$ is sufficiently small since $\{\beta_k\}_k$ is increasing and unbounded. It remains to control

526 the first term in (12). To that end, after recalling Step 2 of Algorithm 1 and applying the triangle
 527 inequality, we can write that

$$\begin{aligned} & \text{dist}(-\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k), \partial g(x_k)) \\ & \leq \text{dist}(-\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1}), \partial g(x_k)) \\ & \quad + \|\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k) - \nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1})\|. \end{aligned} \quad (44)$$

528 The first term on the right-hand side above is bounded by ϵ_k , by Step 5 of Algorithm 1. For the
 529 second term on the right-hand side of (44), we write that

$$\begin{aligned} & \|\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k) - \nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1})\| \\ & = \|DA(x_k)^\top (y_k - y_{k-1})\| \quad (\text{see (7)}) \\ & \leq \lambda'_A \|y_k - y_{k-1}\| \quad (\text{see (19)}) \\ & = \lambda'_A \sigma_k \|A(x_k)\| \quad (\text{see Step 5 of Algorithm 1}) \\ & \leq \frac{2\lambda'_A \sigma_k}{\nu \beta_{k-1}} (\lambda'_f + \lambda'_A y_{\max}). \quad (\text{see (43)}) \end{aligned} \quad (45)$$

530 By combining (44,45), we find that

$$\begin{aligned} & \text{dist}(\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k), \partial g(x_k)) \\ & \leq \frac{2\lambda'_A \sigma_k}{\nu \beta_{k-1}} (\lambda'_f + \lambda'_A y_{\max}) + \epsilon_k. \end{aligned} \quad (46)$$

531 By combining (43,46), we find that

$$\begin{aligned} & \text{dist}(-\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k), \partial g(x_k)) + \|A(x_k)\| \\ & \leq \left(\frac{2\lambda'_A \sigma_k}{\nu \beta_{k-1}} (\lambda'_f + \lambda'_A y_{\max}) + \epsilon_k \right) \\ & \quad + 2 \left(\frac{\lambda'_f + \lambda'_A y_{\max}}{\nu \beta_{k-1}} \right). \end{aligned} \quad (47)$$

532 Applying $\sigma_k \leq \sigma_1$, we find that

$$\begin{aligned} & \text{dist}(-\nabla_x \mathcal{L}_{\beta_{k-1}}(x_k, y_k), \partial g(x_k)) + \|A(x_k)\| \\ & \leq \frac{2\lambda'_A \sigma_1 + 2}{\nu \beta_{k-1}} (\lambda'_f + \lambda'_A y_{\max}) + \epsilon_k. \end{aligned} \quad (48)$$

533 For the second part of the theorem, we use the Weyl's inequality and Step 5 of Algorithm 1 to write

$$\begin{aligned} & \lambda_{\min}(\nabla_{xx} \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1})) \geq \lambda_{\min}(\nabla_{xx} \mathcal{L}_{\beta_{k-1}}(x_k, y_k)) \\ & \quad - \sigma_k \left\| \sum_{i=1}^m A_i(x_k) \nabla^2 A_i(x_k) \right\|. \end{aligned} \quad (49)$$

534 The first term on the right-hand side is lower bounded by $-\epsilon_{k-1}$ by Step 2 of Algorithm 1. We next
 535 bound the second term on the right-hand side above as

$$\begin{aligned} & \sigma_k \left\| \sum_{i=1}^m A_i(x_k) \nabla^2 A_i(x_k) \right\| \\ & \leq \sigma_k \sqrt{m} \max_i \|A_i(x_k)\| \|\nabla^2 A_i(x_k)\| \\ & \leq \sigma_k \sqrt{m} \lambda_A \frac{2\lambda'_f + 2\lambda'_A y_{\max}}{\nu \beta_{k-1}}, \end{aligned}$$

536 where the last inequality is due to (5,43). Plugging into (49) gives

$$\begin{aligned} & \lambda_{\min}(\nabla_{xx} \mathcal{L}_{\beta_{k-1}}(x_k, y_{k-1})) \\ & \geq -\epsilon_{k-1} - \sigma_k \sqrt{m} \lambda_A \frac{2\lambda'_f + 2\lambda'_A y_{\max}}{\nu \beta_{k-1}}, \end{aligned}$$

537 which completes the proof of Theorem 4.1.

538 **B Proof of Corollary 4.2**

539 Let K denote the number of (outer) iterations of Algorithm 1 and let ϵ_f denote the desired accuracy
 540 of Algorithm 1, see (11). Recalling Theorem 4.1, we can then write that

$$\epsilon_f = \frac{Q}{\beta_K}, \quad (50)$$

541 or, equivalently, $\beta_K = Q/\epsilon_f$. We now count the number of total (inner) iterations T of Algorithm 1
 542 to reach the accuracy ϵ_f . From (17) and for sufficiently large k , recall that $\lambda_{\beta_k} \leq \lambda''\beta_k$ is the
 543 smoothness parameter of the augmented Lagrangian. Then, from (24) ad by summing over the outer
 544 iterations, we bound the total number of (inner) iterations of Algorithm 1 as

$$\begin{aligned} T &= \sum_{k=1}^K \mathcal{O}\left(\frac{\lambda_{\beta_{k-1}}^2 \rho^2}{\epsilon_k}\right) \\ &= \sum_{k=1}^K \mathcal{O}(\beta_{k-1}^3 \rho^2) \quad (\text{Step 1 of Algorithm 1}) \\ &\leq \mathcal{O}(K \beta_{K-1}^3 \rho^2) \quad (\{\beta_k\}_k \text{ is increasing}) \\ &\leq \mathcal{O}\left(\frac{K Q^3 \rho^2}{\epsilon_f^3}\right). \quad (\text{see (50)}) \end{aligned} \quad (51)$$

545 In addition, if we specify $\beta_k = b^k$ for all k , we can further refine T . Indeed,

$$\beta_K = b^K \implies K = \log_b\left(\frac{Q}{\epsilon_f}\right), \quad (52)$$

546 which, after substituting into (51) gives the final bound in Corollary 4.2.

547 **C Proof of Lemma 2.1**

548 Note that

$$\mathcal{L}_\beta(x, y) = f(x) + \sum_{i=1}^m y_i A_i(x) + \frac{\beta}{2} \sum_{i=1}^m (A_i(x))^2, \quad (53)$$

549 which implies that

$$\begin{aligned} &\nabla_x \mathcal{L}_\beta(x, y) \\ &= \nabla f(x) + \sum_{i=1}^m y_i \nabla A_i(x) + \frac{\beta}{2} \sum_{i=1}^m A_i(x) \nabla A_i(x) \\ &= \nabla f(x) + DA(x)^\top y + \beta DA(x)^\top A(x), \end{aligned} \quad (54)$$

550 where $DA(x)$ is the Jacobian of A at x . By taking another derivative with respect to x , we reach

$$\begin{aligned} \nabla_x^2 \mathcal{L}_\beta(x, y) &= \nabla^2 f(x) + \sum_{i=1}^m (y_i + \beta A_i(x)) \nabla^2 A_i(x) \\ &\quad + \beta \sum_{i=1}^m \nabla A_i(x) \nabla A_i(x)^\top. \end{aligned} \quad (55)$$

551 It follows that

$$\begin{aligned} &\|\nabla_x^2 \mathcal{L}_\beta(x, y)\| \\ &\leq \|\nabla^2 f(x)\| + \max_i \|\nabla^2 A_i(x)\| (\|y\|_1 + \beta \|A(x)\|_1) \\ &\quad + \beta \sum_{i=1}^m \|\nabla A_i(x)\|^2 \\ &\leq \lambda_h + \sqrt{m} \lambda_A (\|y\| + \beta \|A(x)\|) + \beta \|DA(x)\|_F^2. \end{aligned} \quad (56)$$

552 For every x such that $\|x\| \leq \rho$ and $\|A(x)\| \leq \rho$, we conclude that

$$\|\nabla_x^2 \mathcal{L}_\beta(x, y)\| \leq \lambda_f + \sqrt{m} \lambda_A (\|y\| + \beta \rho') + \beta \max_{\|x\| \leq \rho} \|DA(x)\|_F^2, \quad (57)$$

553 which completes the proof of Lemma 2.1.

554 D Clustering

555 We only verify the condition in (20) here. Note that

$$A(x) = VV^\top \mathbf{1} - \mathbf{1}, \quad (58)$$

556

$$\begin{aligned} DA(x) &= \begin{bmatrix} w_{1,1}x_1^\top & \cdots & w_{1,n}x_1^\top \\ \vdots & & \\ w_{n,1}x_n^\top & \cdots & w_{n,n}x_n^\top \end{bmatrix} \\ &= [V \quad \cdots \quad V] + \begin{bmatrix} x_1^\top & & \\ & \ddots & \\ & & x_n^\top \end{bmatrix}, \end{aligned} \quad (59)$$

557 where $w_{i,i} = 2$ and $w_{i,j} = 1$ for $i \neq j$. In the last line above, n copies of V appear and the last
558 matrix above is block-diagonal. For x_k , define V_k accordingly and let $x_{k,i}$ be the i th row of V_k .
559 Consequently,

$$\begin{aligned} DA(x_k)^\top A(x_k) &= \begin{bmatrix} (V_k^\top V_k - I_n)V_k^\top \mathbf{1} \\ \vdots \\ (V_k^\top V_k - I_n)V_k^\top \mathbf{1} \end{bmatrix} \\ &\quad + \begin{bmatrix} x_{k,1}(V_k V_k^\top \mathbf{1} - \mathbf{1})_1 \\ \vdots \\ x_{k,n}(V_k V_k^\top \mathbf{1} - \mathbf{1})_n \end{bmatrix}, \end{aligned} \quad (60)$$

560 where $I_n \in \mathbb{R}^{n \times n}$ is the identity matrix. Let us make a number of simplifying assumptions. First, we
561 assume that $\|x_k\| < \sqrt{s}$ (which can be enforced in the iterates by replacing C with $(1 - \epsilon)C$ for a
562 small positive ϵ in the subproblems). Under this assumption, it follows that

$$(\partial g(x_k))_i = \begin{cases} 0 & (x_k)_i > 0 \\ \{a : a \leq 0\} & (x_k)_i = 0, \end{cases} \quad i \leq d. \quad (61)$$

563 Second, we assume that V_k has nearly orthonormal columns, namely, $V_k^\top V_k \approx I_n$. This can also be
564 enforced in each iterate of Algorithm 1 and naturally corresponds to well-separated clusters. While a
565 more fine-tuned argument can remove these assumptions, they will help us simplify the presentation
566 here. Under these assumptions, the (squared) right-hand side of (20) becomes

$$\begin{aligned} &\text{dist} \left(-DA(x_k)^\top A(x_k), \frac{\partial g(x_k)}{\beta_{k-1}} \right)^2 \\ &= \left\| (-DA(x_k)^\top A(x_k))_+ \right\|^2 \quad (a_+ = \max(a, 0)) \\ &= \left\| \begin{bmatrix} x_{k,1}(V_k V_k^\top \mathbf{1} - \mathbf{1})_1 \\ \vdots \\ x_{k,n}(V_k V_k^\top \mathbf{1} - \mathbf{1})_n \end{bmatrix} \right\|^2 \quad (x_k \in C \Rightarrow x_k \geq 0) \\ &= \sum_{i=1}^n \|x_{k,i}\|^2 (V_k V_k^\top \mathbf{1} - \mathbf{1})_i^2 \\ &\geq \min_i \|x_{k,i}\|^2 \cdot \sum_{i=1}^n (V_k V_k^\top \mathbf{1} - \mathbf{1})_i^2 \\ &= \min_i \|x_{k,i}\|^2 \cdot \|V_k V_k^\top \mathbf{1} - \mathbf{1}\|^2. \end{aligned} \quad (62)$$

567 Therefore, given a prescribed ν , ensuring $\min_i \|x_{k,i}\| \geq \nu$ guarantees (20). When the algorithm
 568 is initialized close enough to the constraint set, there is indeed no need to separately enforce (62).
 569 In practice, often n exceeds the number of true clusters and a more intricate analysis is required to
 570 establish (20) by restricting the argument to a particular subspace of \mathbb{R}^n .

571 E Basis Pursuit

572 We only verify the regularity condition in (20) for (1) with f, A, g specified in (35). Note that

$$DA(x) = 2\overline{B}\text{diag}(x), \quad (63)$$

573 where $\text{diag}(x) \in \mathbb{R}^{2d \times 2d}$ is the diagonal matrix formed by x . The left-hand side of (20) then reads as

$$\begin{aligned} & \text{dist} \left(-DA(x_k)^\top A(x_k), \frac{\partial g(x_k)}{\beta_{k-1}} \right) \\ &= \text{dist} \left(-DA(x_k)^\top A(x_k), \{0\} \right) \quad (g \equiv 0) \\ &= \|DA(x_k)^\top A(x_k)\| \\ &= 2\|\text{diag}(x_k)\overline{B}^\top (\overline{B}x_k^{\circ 2} - b)\|. \quad (\text{see (63)}) \end{aligned} \quad (64)$$

574 To bound the last line above, let x_* be a solution of (1) and note that $\overline{B}x_*^{\circ 2} = b$ by definition. Let also
 575 $z_k, z_* \in \mathbb{R}^d$ denote the vectors corresponding to x_k, x_* . Corresponding to x_k , also define $u_{k,1}, u_{k,2}$
 576 naturally and let $|z_k| = u_{k,1}^{\circ 2} + u_{k,2}^{\circ 2} \in \mathbb{R}^d$ be the vector of amplitudes of z_k . To simplify matters, let
 577 us assume also that B is full-rank. We then rewrite the norm in the last line of (64) as

$$\begin{aligned} & \|\text{diag}(x_k)\overline{B}^\top (\overline{B}x_k^{\circ 2} - b)\|^2 \\ &= \|\text{diag}(x_k)\overline{B}^\top \overline{B}(x_k^{\circ 2} - x_*^{\circ 2})\|^2 \quad (\overline{B}x_*^{\circ 2} = b) \\ &= \|\text{diag}(x_k)\overline{B}^\top B(x_k - x_*)\|^2 \\ &= \|\text{diag}(u_{k,1})B^\top B(z_k - z_*)\|^2 \\ & \quad + \|\text{diag}(u_{k,2})B^\top B(z_k - z_*)\|^2 \\ &= \|\text{diag}(u_{k,1}^{\circ 2} + u_{k,2}^{\circ 2})B^\top B(z_k - z_*)\|^2 \\ &= \|\text{diag}(|z_k|)B^\top B(z_k - z_*)\|^2 \\ &\geq \eta_n(B\text{diag}(|z_k|))^2 \|B(z_k - z_*)\|^2 \\ &= \eta_n(B\text{diag}(|z_k|))^2 \|Bz_k - b\|^2 \quad (Bz_* = \overline{B}x_*^{\circ 2} = b) \\ &\geq \min_{|T|=n} \eta_n(B_T) \cdot |z_{k,(n)}|^2 \|Bz_k - b\|^2, \end{aligned} \quad (65)$$

578 where $\eta_n(\cdot)$ returns the n th largest singular value of its argument. In the last line above, B_T is the
 579 restriction of B to the columns indexed by T of size n . Moreover, $z_{k,(n)}$ is the n th largest entry of z
 580 in magnitude. Given a prescribed ν , (20) therefore holds if

$$|z_{k,(n)}| \geq \frac{\nu}{2\sqrt{\min_{|T|=n} \eta_n(B_T)}}, \quad (66)$$

581 for every iteration k . If Algorithm 1 is initialized close enough to the solution z^* and the entries of
 582 z^* are sufficiently large in magnitude, there will be no need to directly enforce (66).

583 E.1 ℓ_∞ Denoising with a Generative Prior

584 (mfs): We need Fabian's input here.

585 The authors of [32] have proposed to project onto the range of a Generative Adversarial network
 586 (GAN) [30], as a way to defend against adversarial examples. For a given noisy observation $x^* + \eta$,
 587 they consider a projection in the ℓ_2 norm. We instead propose to use our augmented Lagrangian
 588 method to denoise in the ℓ_∞ norm, a much harder task:

$$\begin{aligned} & \min_{x,z} \|x^* + \eta - x\|_\infty \\ & \text{s.t.} \quad x = G(z). \end{aligned} \quad (67)$$




Figure 3: Augmented Lagrangian vs Adam for ℓ_∞ denoising (left). ℓ_2 vs ℓ_∞ denoising as defense against adversarial examples

589 We use a pretrained generator for the MNIST dataset, given by a standard deconvolutional neural
 590 network architecture. We compare the successful optimizer Adam against our method. Our algorithm
 591 involves two forward/backward passes through the network, as opposed to Adam that requires only
 592 one. For this reason we let our algorithm run for 4000 iterations, and Adam for 8000 iterations.
 593 For a particular example, we plot the objective value vs iteration count in figure E.1. Our method
 594 successfully minimizes the objective value, while Adam does not succeed.

595 E.2 Generalized Eigenvalue Problem

596 Generalized eigenvalue problem has extensive applications in machine learning, statistics and data
 597 analysis [28]. The well-known nonconvex formulation of the problem is [14] given by

$$\begin{cases} \min_{x \in \mathbb{R}^n} x^\top C x \\ x^\top B x = 1, \end{cases} \quad (68)$$

598 where $B, C \in \mathbb{R}^{n \times n}$ are symmetric matrices and B is positive definite, namely, $B \succ 0$. The
 599 generalized eigenvector computation is equivalent to performing principal component analysis (PCA)
 600 of C in the norm B . It is also equivalent to computing the top eigenvector of symmetric matrix
 601 $S = B^{-1/2} C B^{1/2}$ and multiplying the resulting vector by $B^{-1/2}$. However, for large values of n ,
 602 computing $B^{-1/2}$ is extremely expensive. The natural convex SDP relaxation for (68) involves lifting
 603 $Y = x x^\top$ and removing the nonconvex $\text{rank}(Y) = 1$ constraint, namely,

$$\begin{cases} \min_{Y \in \mathbb{R}^{n \times n}} \text{tr}(C Y) \\ \text{tr}(B Y) = 1, \quad X \succeq 0. \end{cases} \quad (69)$$

604 Here, however, we opt to directly solve (68) because it fits into our template with

$$\begin{aligned} f(x) &= x^\top C x, & g(x) &= 0, \\ A(x) &= x^\top B x - 1. \end{aligned} \quad (70)$$

605 We compare our approach against three different methods: manifold based Riemannian gradient
 606 descent and Riemannian trust region methods in [12] and the linear system solver in [28], abbreviated
 607 as GenELin. We have used Manopt software package in [?] for the manifold based methods. For
 608 GenELin, we have utilized Matlab’s backslash operator as the linear solver. The results are compiled
 609 in Figure 4.

610 Here, we verify the regularity condition in (20) for problem (68). Note that

$$DA(x) = (2Bx)^\top. \quad (71)$$

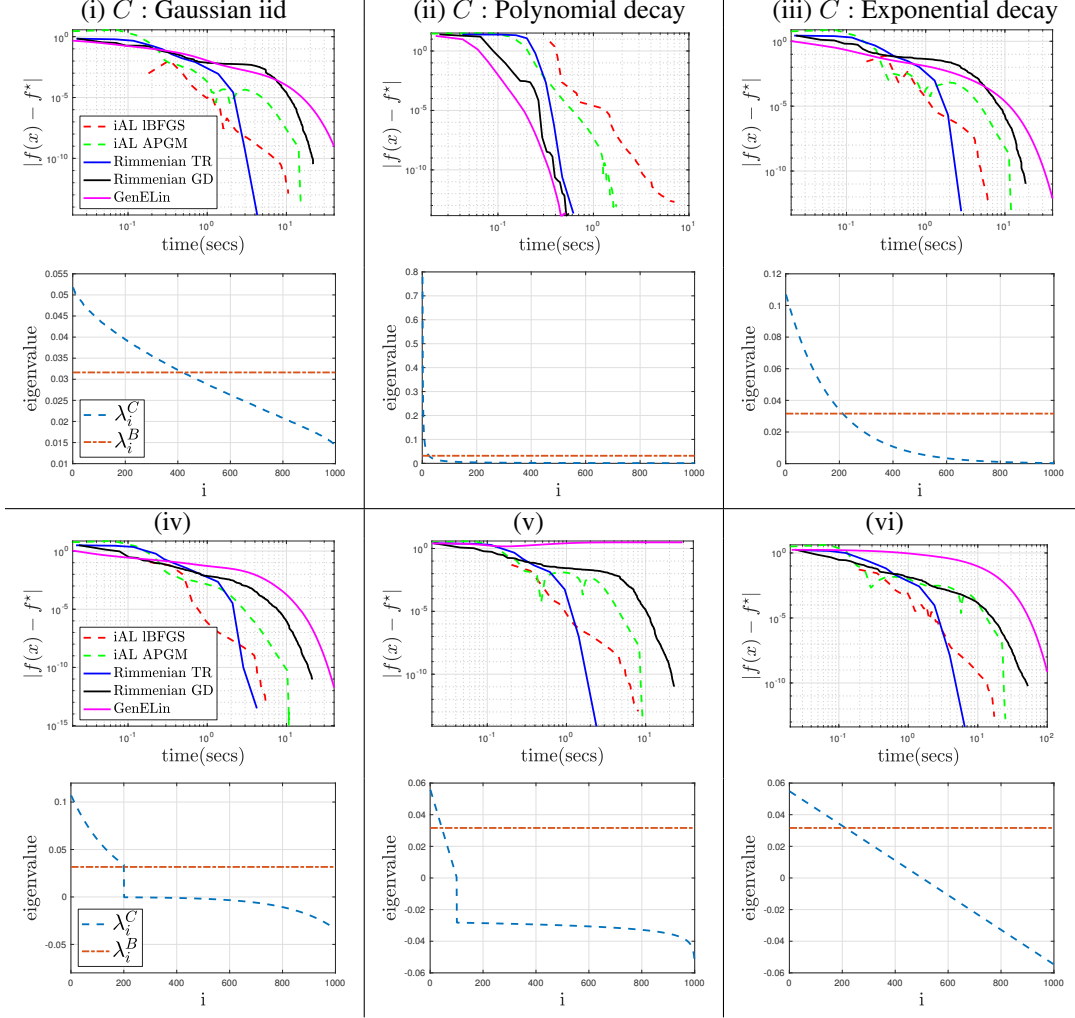


Figure 4: (Top) Objective convergence for calculating top generalized eigenvalue and eigenvector of B and C . (Bottom) Eigenvalue structure of the matrices. For (i),(ii) and (iii), C is positive semidefinite; for (iv), (v) and (vi), C contains negative eigenvalues. [(i): Generated by taking symmetric part of iid Gaussian matrix. (ii): Generated by randomly rotating $\text{diag}(1^{-p}, 2^{-p}, \dots, 1000^{-p})(p = 1)$. (iii): Generated by randomly rotating $\text{diag}(10^{-p}, 10^{-2p}, \dots, 10^{-1000p})(p = 0.0025)$.]

611 Therefore,

$$\begin{aligned}
 \text{dist} \left(-DA(x_k)^\top A(x_k), \frac{\partial g(x_k)}{\beta_{k-1}} \right)^2 &= \text{dist} (-DA(x_k)^\top A(x_k), \{0\})^2 \quad (g \equiv 0) \\
 &= \|DA(x_k)^\top A(x_k)\|^2 \\
 &= \|2Bx_k(x_k^\top Bx_k - 1)\|^2 \quad (\text{see (71)}) \\
 &= 4(x_k^\top Bx_k - 1)^2 \|Bx_k\|^2 \\
 &= 4\|Bx_k\|^2 \|A(x_k)\|^2 \quad (\text{see (70)}) \\
 &\geq \eta_{\min}(B)^2 \|x_k\|^2 \|A(x_k)\|^2, \quad (72)
 \end{aligned}$$

612 where $\eta_{\min}(B)$ is the smallest eigenvalue of the positive definite matrix B . Therefore, for a prescribed
 613 ν , the regularity condition in (20) holds with $\|x_k\| \geq \nu/\eta_{\min}$ for every k . If the algorithm is initialized
 614 close enough to the constraint set, there will be again no need to directly enforce this latter condition.