

ON CONVERGENCE OF PROJECTED GRADIENT DESCENT FOR MINIMIZING A LARGE-SCALE QUADRATIC OVER THE UNIT SPHERE

Trung Vu, Raviv Raich, and Xiao Fu

School of EECS, Oregon State University, Corvallis, OR 97331-5501 USA

{vutru, raich, xiao.fu}@oregonstate.edu

ABSTRACT

Unit sphere-constrained quadratic optimization has been studied extensively over the past decades. While state-of-art algorithms for solving this problem often rely on relaxation or approximation techniques, there has been little research into scalable first-order methods that tackle the problem in its original form. These first-order methods are often more well-suited for the big data setting. In this paper, we provide a novel analysis of the simple projected gradient descent method for minimizing a quadratic over a sphere. When the gradient step size is sufficiently small, we show that convergence is locally linear and provide a closed-form expression for the rate. Moreover, a careful selection of the step size can stimulate convergence to the global solution while preventing convergence to local minima.

Index Terms— large-scale optimization, unit-norm constraint, quadratic programming, convergence analysis

1. INTRODUCTION

This paper studies the problem of minimizing a quadratic with a norm constraint:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} - \mathbf{b}^T \mathbf{x} \quad \text{subject to } \|\mathbf{x}\| \leq 1, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{b} \in \mathbb{R}^n$ and $\|\cdot\|$ is the Euclidean norm.¹ This optimization problem arises frequently in many machine learning and signal processing applications including contour grouping [1], graph partitioning [2] and seismic inversion [3].

If the global solution \mathbf{x}_* of (1) lies in the interior of the unit sphere, i.e., $\|\mathbf{x}_*\| < 1$, then \mathbf{x}_* is also the solution of the unconstrained problem. Thus, it is more challenging to consider the case when $\|\mathbf{x}_*\| = 1$. To that end, we restrict our

interest to the following problem of minimizing a quadratic over a sphere:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} - \mathbf{b}^T \mathbf{x} \quad \text{subject to } \|\mathbf{x}\|^2 = 1, \quad (2)$$

In this formulation, we assume that \mathbf{A} is symmetric, but not necessarily positive semidefinite. Hence, the objective function is potentially non-convex. Additionally, the norm constraint is non-convex. Both (1) and (2) are instances of quadratic constrained quadratic program with only one constraint (QCQP-1) and they have been extensively studied in the literature. State-of-art methods for solving this type of QCQP-1 problems in polynomial time include semidefinite relaxation (SDR) [4] and Lagrangian relaxation [5]. However, the problem size for these methods often grows quadratically, making them inapplicable to large-scale problems.

From a different standpoint, problems (1) and (2) also arise in linear algebra and optimization as the trust-region subproblem. There have been a few extensions to large-scale settings. In [6], Golub and von Matt leveraged the theory of Gauss quadrature and proposed a method to approximately solve (1) by tridiagonalizing \mathbf{A} using the Lanczos process. In another approach, Sorensen [7] recast the trust-region subproblem in terms of a parameterized eigenvalue problem and developed an implicitly restarted Lanczos method. Related schemes can also be found in [8, 9]. In 2001, Hager introduced sequential subspace method (SSM) [10, 11], carrying out the minimization over a sequence of subspaces that are adjusted after each sequential quadratic programming (SQP) iterate. Similar to the aforementioned methods, SSM relies on Lanczos process to compute the smallest eigenvalue and the corresponding eigenvector of \mathbf{A} .

In this paper, we focus on scalable first-order methods for solving (2) directly. We leverage the use of simple gradient projection and establish convergence results in the non-convex setting of the spherically constrained quadratic minimization problem, where most convergence guarantees in convex optimization start to break. Our analysis provides a

¹This work is partially supported by the National Science Foundation grant CCF-1254218.

¹Generally, we can always assume \mathbf{A} to be symmetric. Otherwise, one can define an equivalent objective function using $\tilde{\mathbf{A}} = \frac{1}{2}(\mathbf{A} + \mathbf{A}^T)$.

novel insight into behaviors of the algorithm in the neighborhoods of the local optima. Understanding these convergence properties enables us to (i) accommodate acceleration near the optimum—e.g., by using optimal step size selection or momentum methods—and (ii) identify ways of enabling convergence to global solution. Finally, we present numerical results that illustrate the theory developed in the paper.

2. SOLUTION PROPERTIES

Consider the Lagrange function

$$\mathcal{L}(\mathbf{x}, \gamma) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} - \mathbf{b}^T \mathbf{x} - \frac{1}{2} \gamma (\|\mathbf{x}\|^2 - 1)$$

where γ is the Lagrange multiplier. The first-order Lagrangian conditions for optimality can be specified as

$$\begin{cases} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \gamma) = \mathbf{A} \mathbf{x} - \mathbf{b} - \gamma \mathbf{x} = \mathbf{0}, \\ \nabla_{\gamma} \mathcal{L}(\mathbf{x}, \gamma) = \|\mathbf{x}\|^2 - 1 = 0. \end{cases}$$

For notational simplicity, we denote the residual by $\mathbf{r} = \mathbf{A} \mathbf{x} - \mathbf{b}$ and the unit sphere by $\mathcal{S}^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| = 1\}$. Formally, these conditions are given in the following lemma.

Lemma 1 (Stationary conditions). *The vector \mathbf{x}_* is a stationary point of problem (2) if and only if $\mathbf{x}_* \in \mathcal{S}^{n-1}$ and there exists a constant $\gamma(\mathbf{x}_*)$ such that $\mathbf{r}_* = \mathbf{A} \mathbf{x}_* - \mathbf{b} = \gamma(\mathbf{x}_*) \cdot \mathbf{x}_*$.*

For the rest of this manuscript, we use the shorthand notation γ to refer $\gamma(\mathbf{x}_*)$. Lemma 1 also implies $\gamma = \mathbf{r}_*^T \mathbf{x}_*$. Denote $\mathbf{P}_{\mathbf{x}_*}^\perp = \mathbf{I} - \mathbf{x}_* \mathbf{x}_*^T$. Let $\lambda_n = 0$ be the zero eigenvalue corresponding to the eigenvector \mathbf{x}_* of the matrix $\mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{n-1}$ be the remaining $n - 1$ eigenvalues. Noticeably, the eigenvalues of $\mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp$ can be bounded by

$$\lambda_{\min}(\mathbf{A}) \leq \lambda_{n-1} \leq \dots \leq \lambda_1 \leq \lambda_{\max}(\mathbf{A}),$$

where $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ are the smallest and largest eigenvalues of \mathbf{A} , respectively. Moreover, the relationship among those eigenvalues and the Lagrange multiplier provides necessary and sufficient conditions for determining the type of a stationary point.

Lemma 2. *A stationary point \mathbf{x}_* of problem (2) is a strict local minimum if and only if $\gamma(\mathbf{x}_*) < \lambda_{n-1}(\mathbf{x}_*)$. Furthermore, \mathbf{x}_* is a **global minimizer** of problem (2) if and only if $\gamma(\mathbf{x}_*) \leq \lambda_{\min}(\mathbf{A})$.*

Due to space limitation, we provide proofs of our lemmas at https://trungvietvu.github.io/files/MLSP19_Appendix.pdf.

Example 1. *Figure 1 demonstrates various cases where there are different numbers of stationary points. As an exemplification, let us examine the derivation of the problem in Fig. 1(b):*

$$\min_{x_1, x_2} \frac{1}{2} (4x_1^2 + x_2^2) - 2x_1 \quad \text{s.t.} \quad x_1^2 + x_2^2 = 1.$$

For each stationary point \mathbf{x}_ , the matrix $\mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp$ has one zero eigenvalue corresponding to the eigenvector \mathbf{x}_* , and the other non-zero eigenvalue $\lambda_1 = \lambda_{n-1}$ (since $n = 2$) lies between $\lambda_{\min}(\mathbf{A}) = 1$ and $\lambda_{\max}(\mathbf{A}) = 4$. Omitting the detailed calculation, we list the four stationary points of this problem as follows: (i) a global maximum at $[x_1, x_2] = [-1, 0]$ with $\gamma = 6, \lambda_1 = 1$; (ii) a local maximum at $[x_1, x_2] = [1, 0]$ with $\gamma = 2, \lambda_1 = 1$; and (iii) 2 local (also global) minima at $[x_1, x_2] = [\frac{2}{3}, \pm \frac{\sqrt{5}}{3}]$ with $\gamma = 1, \lambda_1 = \frac{8}{3}$.*

3. THE PROJECTED GRADIENT ALGORITHM

The projected gradient descent approach (see Algorithm 1) starts at an initial point $\mathbf{x}^{(0)}$, then performs the update

$$\mathbf{x}^{(t+1)} = \mathbf{f}_\alpha(\mathbf{x}^{(t)}) = \mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{x}^{(t)} - \alpha(\mathbf{A} \mathbf{x}^{(t)} - \mathbf{b})), \quad (3)$$

where $\alpha > 0$ is the step size and $\mathcal{P}_{\mathcal{S}^{n-1}}(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the spherical projection uniquely given by

$$\mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{x}) = \begin{cases} \frac{\mathbf{x}}{\|\mathbf{x}\|} & \text{if } \mathbf{x} \neq \mathbf{0}, \\ \mathbf{e} & \text{if } \mathbf{x} = \mathbf{0}, \end{cases}$$

with $\mathbf{e} \in \mathcal{S}^{n-1}$ such that \mathbf{e} and $\mathbf{A} \mathbf{e} - \mathbf{b}$ are not collinear. The definition of projection at $\mathbf{0}$ is just for numerical issues when the algorithm encounters the origin at some iteration. In practice, we can choose \mathbf{e} to be one of the natural basis, i.e., $[1, 0, \dots, 0]$. Next, let us consider some important properties associated with Algorithm 1.

Definition 1. *A fixed point of \mathbf{f}_α is defined as any vector $\bar{\mathbf{x}} \in \mathbb{R}^n$ such that*

$$\mathbf{f}_\alpha(\bar{\mathbf{x}}) = \mathcal{P}_{\mathcal{S}^{n-1}}(\bar{\mathbf{x}} - \alpha(\mathbf{A} \bar{\mathbf{x}} - \mathbf{b})) = \bar{\mathbf{x}}.$$

Lemma 3. *The vector $\bar{\mathbf{x}}$ is a fixed point of \mathbf{f}_α if and only if $\bar{\mathbf{x}} \in \mathcal{S}^{n-1}$ and there exists a constant $\gamma < \frac{1}{\alpha}$ such that $\bar{\mathbf{r}} = \mathbf{A} \bar{\mathbf{x}} - \mathbf{b} = \gamma \bar{\mathbf{x}}$.*

Proof. Since $\bar{\mathbf{x}}$ is a fixed point of \mathbf{f}_α , we have

$$\bar{\mathbf{x}} = \mathbf{f}_\alpha(\bar{\mathbf{x}}) = \mathcal{P}_{\mathcal{S}^{n-1}}(\bar{\mathbf{x}} - \alpha(\mathbf{A} \bar{\mathbf{x}} - \mathbf{b})).$$

Consequently, $\bar{\mathbf{x}} \in \mathcal{S}^{n-1}$. Furthermore, if $\bar{\mathbf{x}} - \alpha(\mathbf{A} \bar{\mathbf{x}} - \mathbf{b}) = \mathbf{0}$, then $\bar{\mathbf{x}} = \mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{0}) = \mathbf{e}$. But this contradicts with the non-collinearity of \mathbf{e} and $\mathbf{A} \mathbf{e} - \mathbf{b}$. Thus, it must be the case that $\bar{\mathbf{x}} - \alpha(\mathbf{A} \bar{\mathbf{x}} - \mathbf{b}) \neq \mathbf{0}$, and hence $\bar{\mathbf{x}} = \frac{\bar{\mathbf{x}} - \alpha \bar{\mathbf{r}}}{\|\bar{\mathbf{x}} - \alpha \bar{\mathbf{r}}\|}$. There exists a constant γ such that $\bar{\mathbf{r}} = \gamma \bar{\mathbf{x}}$. Substituting back into the fixed-point equation yields

$$\bar{\mathbf{x}} = \frac{(1 - \alpha\gamma) \bar{\mathbf{x}}}{|1 - \alpha\gamma| \|\bar{\mathbf{x}}\|} = \frac{1 - \alpha\gamma}{|1 - \alpha\gamma|} \frac{\bar{\mathbf{x}}}{\|\bar{\mathbf{x}}\|} = \text{sign}(1 - \alpha\gamma) \cdot \bar{\mathbf{x}}.$$

Therefore, the sign of $1 - \alpha\gamma$ must be 1 or $1 - \alpha\gamma > 0$. \square

From Lemma 2 and 3, we can establish the necessary and sufficient conditions for a fixed point of \mathbf{f}_α to be a stationary point as follows.

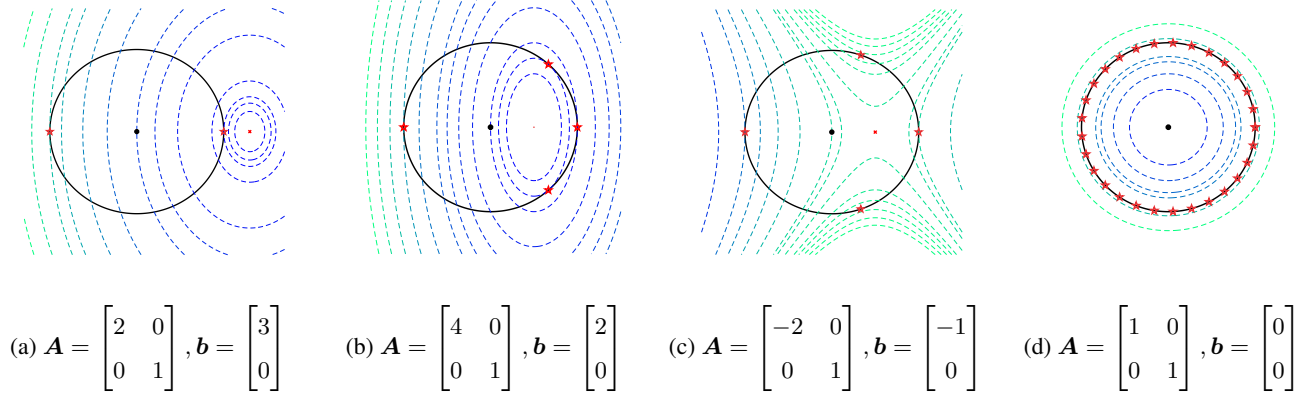


Fig. 1: Examples of minimizing a quadratic over a sphere. Stationary points are given in red stars. In 2D scenario, they can be either local minima or local maxima (a-c). It is also possible that a local optimum lies in a continuum of optima (d).

Algorithm 1: Projected Gradient Descent (PGD)

- 1: Initialize $\mathbf{x}^{(0)} \in \mathcal{S}^{n-1}$
 - 2: **for** $t = 0, 1, \dots$ **do**
 - 3: $\mathbf{z}^{(t+1)} = \mathbf{x}^{(t)} - \alpha(\mathbf{A}\mathbf{x}^{(t)} - \mathbf{b})$
 - 4: $\mathbf{x}^{(t+1)} = \mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{z}^{(t+1)})$
-

Corollary 1. *The vector \mathbf{x}_* is a stationary point of problem (2) if and only if there exists $\alpha > 0$ such that \mathbf{x}_* is a fixed point of \mathbf{f}_α .*

Example 2. *Continued from Example 1, we illustrate fixed points with different step sizes in Fig. 2. When α is small enough, all stationary points can be fixed points. As α increases, only stationary points with the multiplier $\gamma < 1/\alpha$ remains to be fixed points of \mathbf{f}_α . Interestingly, while any convergence point of Algorithm 1 with step size α is a fixed point of the iterated function \mathbf{f}_α , the vice versa is not true: the global maximum at $[x_1, x_2] = [-1, 0]$ is a fixed point of \mathbf{f}_α , for $\alpha < 1/6$, but as can be seen in the next section, it is not a convergence point of the algorithm.*

4. CONVERGENCE ANALYSIS

In this section, we present our result on the local uniform convergence of Algorithm 1 with a certain choice of step size to a strict local optimum. The convergence is shown to be linear and the asymptotic rate is given in a closed-form expression. The challenges come the non-convexity of the norm constraint and (potentially) the negative curvature of the objective function. Let us begin with the analysis of the projection operator.

Lemma 4 (Taylor series expansion of the projection). *Let $\mathbf{x} \in \mathbb{R}^n$ be a nonzero vector and δ be a small perturbation*

such that $\|\delta\| \ll \|\mathbf{x}\|$. Then,

$$\mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{x} + \delta) = \mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{x}) + \frac{1}{\|\mathbf{x}\|} \left(\mathbf{I} - \frac{\mathbf{x}\mathbf{x}^T}{\|\mathbf{x}\|^2} \right) \delta + O(\|\delta\|^2).$$

Now, considering the convergence of Algorithm 1 in the region near a strict local minimum \mathbf{x}_* where $\mathbf{r}_* = \mathbf{A}\mathbf{x}_* - \mathbf{b} = \gamma\mathbf{x}_*$. Denote $\delta^{(t)} = \mathbf{x}^{(t)} - \mathbf{x}_*$ and reorganize the update equation (3) as

$$\begin{aligned} \delta^{(t+1)} &= \mathcal{P}_{\mathcal{S}^{n-1}}(\mathbf{x}_* - \alpha(\mathbf{A}\mathbf{x}_* - \mathbf{b}) + (\mathbf{I} - \alpha\mathbf{A})\delta^{(t)}) - \mathbf{x}_* \\ &= \mathcal{P}_{\mathcal{S}^{n-1}}((1 - \alpha\gamma)\mathbf{x}_* + (\mathbf{I} - \alpha\mathbf{A})\delta^{(t)}) - \mathcal{P}_{\mathcal{S}^{n-1}}((1 - \alpha\gamma)\mathbf{x}_*). \end{aligned}$$

Substituting $\mathbf{x} = (1 - \alpha\gamma)\mathbf{x}_*$ and $\delta = (\mathbf{I} - \alpha\mathbf{A})\delta^{(t)}$ into Lemma 4 yields

$$\begin{aligned} \delta^{(t+1)} &= \frac{1}{\|(1 - \alpha\gamma)\mathbf{x}_*\|} \left(\mathbf{I} - \frac{(1 - \alpha\gamma)\mathbf{x}_*(1 - \alpha\gamma)\mathbf{x}_*^T}{\|(1 - \alpha\gamma)\mathbf{x}_*\|^2} \right) \\ &\quad \cdot (\mathbf{I} - \alpha\mathbf{A})\delta^{(t)} + O(\|\delta^{(t)}\|^2). \end{aligned}$$

Assume the step size is chosen such that $\alpha\gamma < 1$, and recall that $\mathbf{P}_{\mathbf{x}_*}^\perp = \mathbf{I} - \mathbf{x}_*\mathbf{x}_*^T$ for $\|\mathbf{x}_*\| = 1$. The recursion can be rewritten as

$$\delta^{(t+1)} = \frac{1}{1 - \alpha\gamma} \mathbf{P}_{\mathbf{x}_*}^\perp (\mathbf{I} - \alpha\mathbf{A})\delta^{(t)} + O(\|\delta^{(t)}\|^2). \quad (4)$$

The stability of a general nonlinear difference equation of the form $x^{k+1} = Tx^k + o(\|x^k\|)$ has been well-studied in [12, 13]. In particular, let ρ_α be the spectral radius of $(1 - \alpha\gamma)^{-1} \mathbf{P}_{\mathbf{x}_*}^\perp (\mathbf{I} - \alpha\mathbf{A})$, i.e., the largest absolute value of its eigenvalues. If $\rho_\alpha < 1$, then the series $\{\delta^{(t)}\}$ approaches zeros with sufficiently small $\delta^{(0)}$, where

$$\|\delta^{(t)}\| \leq K \|\delta^{(0)}\| (\rho_\alpha + o(1))^t.$$

Therefore, to prove the local convergence of the PGD algorithm, it is sufficient to show that $\rho_\alpha < 1$. We present our main result on the local uniform convergence to a strict local minimum as follows.

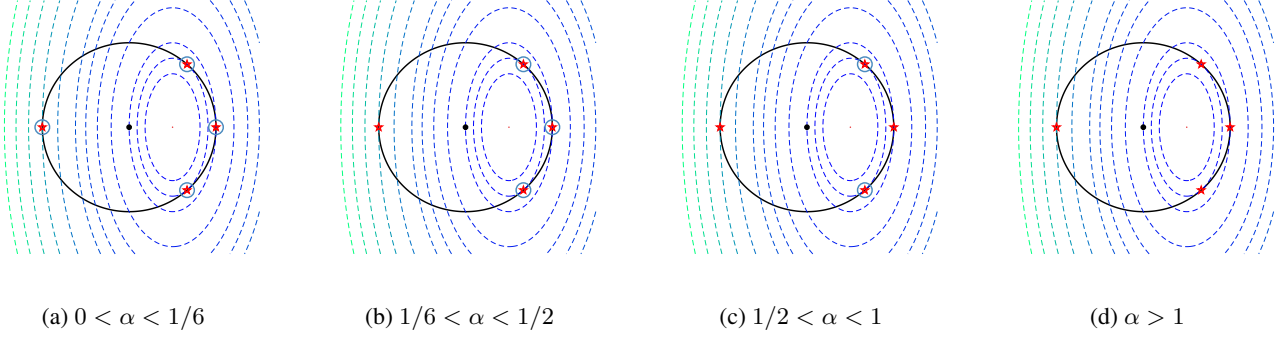


Fig. 2: Stationary points (red stars) versus fixed points (blue circles) with different step size α in optimizing the quadratic objective $\frac{1}{2}(4x_1^2 + x_2^2) - 2x_1$ over the unit circle. Dashed lines are the contour levels of the objective value.

Definition 2. Algorithm 1 with step size α converges **locally uniformly** to \mathbf{x}_* if and only if there exists a constant ϵ such that for any $\mathbf{x}^{(0)}$ satisfying $\|\mathbf{x}^{(0)} - \mathbf{x}_*\| \leq \epsilon$, we have $\|\mathbf{x}^{(t)} - \mathbf{x}_*\| \leq \epsilon$, $\forall t = 0, 1, \dots$ and $\lim_{t \rightarrow \infty} \|\mathbf{x}^{(t)} - \mathbf{x}_*\| = 0$.

Theorem 1. The vector \mathbf{x}_* is a strict local minimum of problem (2), i.e. $\gamma < \lambda_{n-1}$, if and only if there exists $\alpha > 0$ such that Algorithm 1 with step size α converges locally uniformly to \mathbf{x}_* . Furthermore, for any step size $\alpha > 0$ such that $\alpha(\lambda_1 + \gamma) < 2$, the sequence $\{\mathbf{x}^{(t)}\}$ satisfies

$$\|\mathbf{x}^{(t)} - \mathbf{x}_*\| \leq K \|\mathbf{x}^{(0)} - \mathbf{x}_*\| (\rho_\alpha + o(1))^t,$$

for some constant $K > 0$ and $\rho_\alpha = \max_{1 \leq i \leq n-1} \frac{|1 - \alpha \lambda_i|}{1 - \alpha \gamma}$.

The proof of Theorem 1 is given in the appendix. The theorem reveals PGD converges to a local minimum at an asymptotic linear rate ρ_α . Note that in our problem, \mathbf{A} is not necessarily PSD, meaning λ_i could be negative. To facilitate acceleration, one can speed up the convergence by optimizing over the step size α .

Lemma 5. The optimal rate of local convergence and the optimal step size for Algorithm 1 are given by

$$\begin{cases} \rho_* = \frac{\lambda_1 - \lambda_{n-1}}{\lambda_1 + \lambda_{n-1} - 2\gamma}, & \alpha_* = \frac{2}{\lambda_1 + \lambda_{n-1}} & \text{if } \lambda_1 + \lambda_{n-1} > 0 \\ \rho_* = \frac{\lambda_{n-1}}{\gamma}, & \alpha_* = \infty & \text{otherwise.} \end{cases}$$

Example 3. Continued from Example 2, Theorem 1 states that the PGD algorithm only converges locally uniformly to the two local minima at $[2/3, \pm\sqrt{(5)/3}]$ with $\lambda_1 = \frac{8}{3}$. Notice that these points are fixed points of \mathbf{f}_α for $\alpha < 1$. Since $\lambda_1 = \lambda_{n-1}$, the optimal rate is $\rho_* = 0$ with step size $\alpha_* = \frac{3}{8}$. In this case, the convergence is quadratic due to the residual term in (4).

Global convergence of the PGD algorithm. Theorem 1 also implies that for any step size α satisfying $\rho_\alpha > 1$, the algorithm tends to move away from the local minimum \mathbf{x}_* . This

intuition leads us to the following strategy for step size selection: choosing α large enough such that $g(\alpha, \mathbf{x}_*) \geq 2$, where $g(\alpha, \mathbf{x}_*) \triangleq \alpha(\lambda_1(\mathbf{x}_*) + \gamma(\mathbf{x}_*))$, for all strict local minimum \mathbf{x}_* except the global minimum \mathbf{x}_* .

Remark 1. Assume that there exists sufficiently large α satisfying $g(\alpha, \mathbf{x}_*) < 2$ for any global minimum \mathbf{x}_* and $g(\alpha, \mathbf{x}_*) \geq 2$ for any strict local minimum \mathbf{x}_* . Then Algorithm 1 with step size α converges to one of the optimal solutions \mathbf{x}_* at an asymptotic geometric rate of $\rho_\alpha(\mathbf{x}_*)$.

5. NUMERICAL RESULTS

Motivating applications of problem (1) are the well-known trust-region subproblem in nonlinear optimization [14], the variational problem in structural limit analysis [15], and the optimizing precoder method in transmitter based CDMA optimization [16]. For the purpose of demonstrating the theoretical analysis, we will focus on numerical results for local convergence of Algorithm 1 with different step sizes, and empirical evidence for our conjecture about the global convergence with an appropriate step size. In our experiment, we first generate a random symmetric matrix \mathbf{A} of size $n = 1000$ such that the smallest eigenvalue is far away from the other eigenvalues. Then, we choose one multiplier for the global solution $\gamma(\mathbf{x}_*) < \lambda_{\min}(\mathbf{A})$ and one multiplier for the local solution $\gamma(\mathbf{x}_*) > \lambda_{\min}(\mathbf{A})$. Next, the coefficient vector \mathbf{b} is chosen such that $\mathbf{b}^T(\mathbf{A} - \gamma(\mathbf{x}_*)\mathbf{I})^{-2}\mathbf{b} = \mathbf{b}^T(\mathbf{A} - \gamma(\mathbf{x}_*)\mathbf{I})^{-2}\mathbf{b} = 1$. Finally, we compute $\mathbf{x}_* = (\mathbf{A} - \gamma(\mathbf{x}_*)\mathbf{I})^{-1}\mathbf{b}$ and $\mathbf{x}_* = (\mathbf{A} - \gamma(\mathbf{x}_*)\mathbf{I})^{-1}\mathbf{b}$.

For local convergence, starting at an initial point $\mathbf{x}^{(0)}$ close to the local minimum \mathbf{x}_* , we examine four PGD algorithms with different step sizes:

(i) Commonly used step size: $\alpha = 1/L$, where L is the spectral radius of \mathbf{A} . This step size selection is often used in many classic proofs of convergence in convex optimization, where an L -smooth objective function can be guaranteed to monotonically decrease through PGD iterations.




Fig. 3: (a) Local convergence of the projected gradient methods with different step sizes for solving a unit-constrained least squares. (b-c) Empirical evidence of the convergence to global optimum where $\mathbf{x}^{(0)}$ is initialized near a local optimum and α is chosen according to Remark 1. Dashed lines are added as an illustration for the theoretical bounds for the convergence rate of fixed step size methods (up to a constant).

(ii) Optimal step size: $\alpha = \alpha_*$ as in Lemma 5. We choose \mathbf{A} and \mathbf{x}_* such that $\lambda_1 + \lambda_{n-1} > 0$, hence $\alpha_* = 2/(\lambda_1 + \lambda_{n-1})$.
(iii) Projected backtracking line search: rewriting the PGD update as generalized gradient step $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} - \alpha_t \mathbf{G}_{\alpha_t}(\mathbf{x}^{(t)})$, where $\mathbf{G}_\alpha(\mathbf{x}) = \frac{1}{\alpha}(\mathbf{x} - \mathcal{P}_{S^{n-1}}(\mathbf{x} - \alpha(\mathbf{A}\mathbf{x} - \mathbf{b})))$. Denote the quadratic objective by $q(\mathbf{x})$. Starting with $\alpha = 1$, we shrink $\alpha = \beta\alpha$, for $0 < \beta < 1$, while

$$q(\mathbf{x} - \alpha \mathbf{G}_\alpha(\mathbf{x})) > q(\mathbf{x}) - \alpha(\mathbf{A}\mathbf{x} - \mathbf{b})^T \mathbf{G}_\alpha(\mathbf{x}) + \frac{\alpha}{2} \|\mathbf{G}_\alpha(\mathbf{x})\|^2.$$

In our case, this backtracking condition can be simplified to

$$\mathbf{G}_\alpha(\mathbf{x})^T \mathbf{A} \mathbf{G}_\alpha(\mathbf{x}) > \frac{1}{\alpha} \|\mathbf{G}_\alpha(\mathbf{x})\|^2.$$

(iv) Exact line search: finding the step size that maximizes the decrease in objective function

$$\alpha_{\min} = \min_{\alpha > 0} q(\mathcal{P}_{S^{n-1}}(\mathbf{x} - \alpha(\mathbf{A}\mathbf{x} - \mathbf{b}))).$$

As can be seen from Fig. 3(a), the convergence of PGD with step size $\alpha = 1/L$ (blue) is the slowest among the considered methods, followed by the one with optimal step size $\alpha = 2/(\lambda_1 + \lambda_{n-1})$ (red). Note that they both match the asymptotic rate predicted in theory. The adaptive schemes, namely projected backtracking (yellow) and exact line search (magenta) perform slightly better than the optimal fixed step size scheme.

For global convergence, we purposely initialize the algorithm at the same $\mathbf{x}^{(0)}$ that is close to the **local** minimum \mathbf{x}_* , and run the PGD algorithm with step size $\alpha = \frac{1}{2}(\frac{2}{\lambda_1(\mathbf{x}_*) + \gamma(\mathbf{x}_*)} + \frac{2}{\lambda_{n-1}(\mathbf{x}_*) + \gamma(\mathbf{x}_*)})$. It is easy to verify that α satisfies the condition in Remark 1: $g(\alpha, \mathbf{x}_*) < 2 < g(\alpha, \mathbf{x}_*)$. Figure 3 demonstrates the convergence of the algorithm to the global minimizer \mathbf{x}_* , in terms of the distance to the solution (b) and the decrease in the objective value (c). In the first 5000 iterations, the algorithm tries to escape from the

local minimum. Then it experiences a period of fluctuation before getting attracted by the global minimum. Notice that when it reaches the neighborhood of \mathbf{x}_* , (monotonic) linear convergence is observed.²

6. CONCLUSIONS AND FUTURE WORK

We analyzed the projected gradient descent approach to minimizing a quadratic over a sphere. We showed that the algorithm always converges linearly to a strict local minimum in its neighborhood. Further, we provided the closed-form expression for convergence rate and identified ways of achieving optimal rate of convergence near the optimum. Our analysis can be extended in the following directions: (i) minimizing a quadratic over an ellipsoid; (ii) acceleration of gradient projection using momentum; and (iii) analysis of convergence to a continuum of optima.

7. APPENDIX

The proof of Theorem 1 is given as follows.

[\Rightarrow] First, if Algorithm 1 converges locally uniformly to \mathbf{x}_* , our goal is to prove $\gamma < \lambda_{n-1}$. By contradiction, assume that $\gamma > \lambda_{n-1}$ ³. Then choosing $\mathbf{x}^{(0)} = \mathbf{x}_* + \epsilon \mathbf{u}_{n-1}$, where \mathbf{u}_{n-1} is the eigenvector corresponding to λ_{n-1} , leads to

$$\begin{aligned} \delta^{(1)} &= \frac{|1 - \alpha\lambda_{n-1}|}{1 - \alpha\gamma} \delta^{(0)} = \frac{1 - \alpha\lambda_{n-1}}{1 - \alpha\gamma} \delta^{(0)} \\ \Rightarrow \|\mathbf{x}^{(1)} - \mathbf{x}_*\| &= \|\delta^{(1)}\| = \left| \frac{1 - \alpha\lambda_{n-1}}{1 - \alpha\gamma} \right| \|\epsilon \mathbf{u}_{n-1}\| \\ &= \frac{1 - \alpha\lambda_{n-1}}{1 - \alpha\gamma} \epsilon > \epsilon. \quad (\text{since } \alpha\lambda_{n-1} < \alpha\gamma < 1) \end{aligned}$$

²This result is provided merely as an illustration of a typical run, not to be considered as an empirical proof. In our experiments, we re-ran simulations multiple times with various problem sizes and always observed convergence.

³The case $\gamma = \lambda_{n-1}$ leads to convergence to a continuum which we leave as a future work.

This contradicts with the assumption that the sequence $\{\mathbf{x}^{(t)}\}$ lies inside the ϵ -vicinity of \mathbf{x}_* .

[\Leftarrow] Conversely, we will show that if \mathbf{x}_* is a strict local minimum, then for any $\alpha > 0$ such that $\alpha(\lambda_1 + \gamma) < 2$, Algorithm 1 with step size α converges locally uniformly to \mathbf{x}_* . By the same argument in [13], to prove the local stability of equation (4), it is sufficient to consider the linear equation without the quadratic residual

$$\delta^{(t+1)} = \frac{1}{1 - \alpha\gamma} \mathbf{P}_{\mathbf{x}_*}^\perp (\mathbf{I} - \alpha\mathbf{A}) \delta^{(t)}.$$

The above equation implies $\delta^{(t)} = \mathbf{P}_{\mathbf{x}_*}^\perp \delta^{(t)}$ for $t = 1, 2, \dots$. Thus, we have

$$\begin{aligned} \delta^{(t+1)} &= \frac{\mathbf{P}_{\mathbf{x}_*}^\perp (\mathbf{I} - \alpha\mathbf{A}) \mathbf{P}_{\mathbf{x}_*}^\perp}{1 - \alpha\gamma} \delta^{(t)} = \frac{\mathbf{P}_{\mathbf{x}_*}^\perp - \alpha \mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp}{1 - \alpha\gamma} \delta^{(t)} \\ &= \frac{(\mathbf{I} - \alpha \mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp) \mathbf{P}_{\mathbf{x}_*}^\perp}{1 - \alpha\gamma} \delta^{(t)} = \frac{\mathbf{I} - \alpha \mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp}{1 - \alpha\gamma} \delta^{(t)}. \end{aligned}$$

Now consider the matrix $\mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp$. There exists an eigenvalue decomposition $\mathbf{P}_{\mathbf{x}_*}^\perp \mathbf{A} \mathbf{P}_{\mathbf{x}_*}^\perp = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$ where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{n-1}, \mathbf{x}_*]$ is an orthogonal matrix and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{n-1}, 0)$. Let $\mathbf{y}^{(t)} = \mathbf{U}^T \delta^{(t)}$. Then

$$\mathbf{y}^{(t+1)} = \mathbf{U}^T \delta^{(t+1)} = \frac{\mathbf{I} - \alpha \mathbf{\Lambda}}{1 - \alpha\gamma} \mathbf{y}^{(t)} = \left(\frac{\mathbf{I} - \alpha \mathbf{\Lambda}}{1 - \alpha\gamma} \right)^t \mathbf{y}^{(1)}. \quad (5)$$

In addition, since the last column of \mathbf{U} is \mathbf{x}_* , we can compute the last element of $\mathbf{y}^{(1)}$ by

$$y_n^{(1)} = \mathbf{x}_*^T \delta^{(1)} = \mathbf{x}_*^T (\mathbf{I} - \mathbf{x}_* \mathbf{x}_*^T) \delta^{(0)} = 0. \quad (6)$$

From (5) and (6), we obtain

$$\|\mathbf{y}^{(t+1)}\| \leq \max_{1 \leq i \leq n-1} \left| \frac{1 - \alpha\lambda_i}{1 - \alpha\gamma} \right|^t \cdot \|\mathbf{y}^{(1)}\|.$$

Since \mathbf{x}_* is a strict local minimum, it follows from Lemma 2 that $\gamma < \lambda_{n-1} \leq \lambda_1$. Combining with the condition $\alpha\lambda_1 + \alpha\gamma < 2$, we obtain $\alpha\gamma < 1$. In order to show convergence, it remains to prove the inequality

$$\max_{1 \leq i \leq n-1} \frac{|1 - \alpha\lambda_i|}{1 - \alpha\gamma} < 1 \Leftrightarrow |1 - \alpha\lambda_i| < 1 - \alpha\gamma, \forall 1 \leq i \leq n-1.$$

Indeed, this inequality stems from the fact that $\alpha\lambda_1 + \alpha\gamma < 2$ and $\gamma < \lambda_{n-1} \leq \dots \leq \lambda_1$.

8. REFERENCES

- [1] Jean Gallier, *Quadratic Optimization and Contour Grouping*, pp. 439–457, Springer New York, New York, NY, 2011.
- [2] William W Hager and Yaroslav Krylyuk, “Graph partitioning and continuous quadratic programming,” *SIAM Journal on Discrete Mathematics*, vol. 12, no. 4, pp. 500–523, 1999.
- [3] William Menke, *Geophysical data analysis: Discrete inverse theory*, Academic press, 2018.
- [4] Zhi-Quan Luo, Nicholas D Sidiropoulos, Paul Tseng, and Shuzhong Zhang, “Approximation bounds for quadratic optimization with homogeneous quadratic constraints,” *SIAM Journal on optimization*, vol. 18, no. 1, pp. 1–28, 2007.
- [5] Jaehyun Park and Stephen Boyd, “General heuristics for nonconvex quadratically constrained quadratic programming,” *arXiv preprint arXiv:1703.07870*, 2017.
- [6] Gene H Golub and Urs Von Matt, “Quadratically constrained least squares and quadratic problems,” *Numerische Mathematik*, vol. 59, no. 1, pp. 561–580, 1991.
- [7] Danny C Sorensen, “Minimization of a large-scale quadratic function subject to a spherical constraint,” *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 141–161, 1997.
- [8] Franz Rendl and Henry Wolkowicz, “A semidefinite framework for trust region subproblems with applications to large scale minimization,” *Mathematical Programming*, vol. 77, no. 1, pp. 273–299, 1997.
- [9] Marielba Rojas, *A large-scale trust-region approach to the regularization of discrete ill-posed problems*, Ph.D. thesis, Rice University, 1999.
- [10] William W Hager, “Minimizing a quadratic over a sphere,” *SIAM Journal on Optimization*, vol. 12, no. 1, pp. 188–208, 2001.
- [11] William Hager and Soonchul Park, “Global convergence of SSM for minimizing a quadratic over a sphere,” *Mathematics of Computation*, vol. 74, no. 251, pp. 1413–1423, 2005.
- [12] Richard Bellman, *Stability theory of differential equations*, Courier Corporation, 2008.
- [13] B. Polyak, “Some methods of speeding up the convergence of iteration methods,” *USSR Computational Mathematics and Mathematical Physics*, vol. 4, no. 5, pp. 1–17, 1964.
- [14] Andrew R Conn, Nicholas IM Gould, and Ph L Toint, *Trust region methods*, vol. 1, Siam, 2000.
- [15] David Yang Gao, *Duality principles in nonconvex systems: theory, methods and applications*, vol. 39, Springer Science & Business Media, 2013.
- [16] Erik Hons, *Constrained quadratic optimization: Theory and application for wireless communication systems*, Ph.D. thesis, University of Waterloo, 2001.