

---

# A Richer Theory of Convex Constrained Optimization with Reduced Projections and Improved Rates

---

Tianbao Yang<sup>1</sup> Qihang Lin<sup>1</sup> Lijun Zhang<sup>2</sup>

## Abstract

This paper focuses on convex constrained optimization problems, where the solution is subject to a convex inequality constraint. In particular, we aim at challenging problems for which both projection into the constrained domain and a linear optimization under the inequality constraint are time-consuming, which render both projected gradient methods and conditional gradient methods (a.k.a. the Frank-Wolfe algorithm) expensive. In this paper, we develop projection reduced optimization algorithms for both smooth and non-smooth optimization with improved convergence rates under a certain regularity condition of the constraint function. We first present a general theory of optimization with only one projection. Its application to smooth optimization with only one projection yields  $O(1/\epsilon)$  iteration complexity, which improves over the  $O(1/\epsilon^2)$  iteration complexity established before for non-smooth optimization and can be further reduced under strong convexity. Then we introduce a local error bound condition and develop faster algorithms for non-strongly convex optimization at the price of a logarithmic number of projections. In particular, we achieve an iteration complexity of  $\tilde{O}(1/\epsilon^{2(1-\theta)})$  for non-smooth optimization and  $\tilde{O}(1/\epsilon^{1-\theta})$  for smooth optimization, where  $\theta \in (0, 1]$  appearing the local error bound condition characterizes the functional local growth rate around the optimal solutions. Novel applications in solving the constrained  $\ell_1$  minimization problem and a positive semi-definite constrained distance metric learning problem demonstrate that the proposed algorithms achieve significant speed-up compared with previous algorithms.

---

<sup>1</sup>The University of Iowa, Iowa City, IA 52242, USA <sup>2</sup>National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China. Correspondence to: Tianbao Yang <tianbao-yang@uiowa.edu>.

This is the long version of our paper appearing in the *Proceedings of the 34<sup>th</sup> International Conference on Machine Learning*, Sydney, Australia, PMLR 70, 2017. Copyright 2017 by the author(s).

## 1. Introduction

In this paper, we aim at solving the following convex constrained optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}), \quad \text{s.t. } c(\mathbf{x}) \leq 0, \quad (1)$$

where  $f(\mathbf{x})$  is a smooth or non-smooth convex function and  $c(\mathbf{x})$  is a lower-semicontinuous and convex function. The problem can find applications in machine learning, signal processing, statistics, marketing optimization, and etc. For example, in distance metric learning one needs to learn a positive semi-definite (PSD) matrix such that similar examples are close to each other and dissimilar examples are far from each other (Weinberger et al., 2006; Xing et al., 2003), where the positive semi-definite constraint can be cast into a convex inequality constraint. Another example arising in compressive sensing is to minimize the  $\ell_1$  norm of high-dimensional vector subject to a measurement constraint (Candès & Wakin, 2008). Although general interior-point methods can be applied to solve the problem with linear convergence, they suffer from exceedingly high computational cost per-iteration. Another solution is to employ the projected gradient (PG) method (Nesterov, 2004) or the conditional gradient (CG) method (Frank & Wolfe, 1956), where the PG method needs to compute the projection into the constrained domain at each iteration and CG needs to solve a linear optimization problem under the constraint. However, for many constraints (e.g., PSD, quadratic constraints) both projection into the constrained domain and the linear optimization under the constraint are time-consuming, which restrict their capabilities to solving these problems.

Recently, there emerges a new direction towards addressing the challenge of expensive projection that is to reduce the number of projections. In the seminal paper (Mahdavi et al., 2012), the authors have proposed two algorithms with only one projection at the end of iterations for non-smooth convex and strongly convex optimization, respectively. The idea of both algorithms is to move the constraint function into the objective function and to control the violation of constraint for intermediate solutions. While their developed algorithms enjoy an optimal convergence rate for non-smooth optimization (i.e.,  $O(1/\epsilon^2)$  iteration complexity)

and a close-to-optimal convergence rate for strongly convex optimization (i.e.,  $\tilde{O}(1/\epsilon)^1$ ), *there still lack of theory and algorithms with reduced projections and faster rates for smooth convex optimization and for convex optimization without strong convexity assumptions.*

In this paper, we make significant contributions by developing a richer theory of convex constrained optimization with reduced projections and faster rates. To be specific,

- we develop a general framework and theory of optimization with only one projection, where any favorable smooth or non-smooth convex optimization algorithms can be employed to solve the intermediate augmented unconstrained objective function. We discuss in full details the applicability of the proposed algorithms to problems with polyhedral, quadratic or PSD constraints.
- Applying the general theory to smooth convex optimization<sup>2</sup> with Nesterov’s accelerated gradient methods yields an iteration complexity of  $O(1/\epsilon)$  with only one projection. In addition, when equipped with an optimal algorithm for strongly convex optimization the general theory implies the *optimal* iteration complexity of  $O(1/\epsilon)$  for strongly convex optimization with only one projection. For smooth and strongly convex optimization, the general theory implies an iteration complexity of  $O(1/\epsilon^\beta)$  where  $\beta \in (1/2, 1)$  with only one projection and a sufficiently large number of iterations.
- Building on the general framework and theory, we further develop an improved theory with faster convergence rates for non-strongly convex optimization at the price of a logarithmic number of projections. In particular, we show that under a mild local error bound condition, the iteration complexities can be reduced to  $\tilde{O}(1/\epsilon^{2(1-\theta)})$  for non-smooth optimization and  $\tilde{O}(1/\epsilon^{1-\theta})$  for smooth optimization, where  $\theta \in (0, 1]$  is a constant in the local error bound condition that characterizes the local growth rate of functional values. To our knowledge, these are the best convergence results with only a logarithmic number of projections for non-strongly convex optimization. We also demonstrate their effectiveness for solving compressive sensing and distance metric learning problems.

## 2. Related Work

The issue of high projection cost in projected gradient descent has received increasing attention in recent years. Most studies are based on the Frank-Wolfe technique that eschews the projection in favor of a linear optimization over the constrained domain (Jaggi, 2013; Hazan & Kale, 2012; Lacoste-Julien et al., 2013; Garber & Hazan, 2015). It happens that for many bounded domains (e.g., bounded

balls for vectors and matrices, a PSD constraint with a bounded trace norm) the linear optimization over the constrained domain is much cheaper than projection into the constrained domain (Jaggi, 2013). However, there still exist many constraints that render both projection into the constrained domain and linear optimization under the constraint are comparably expensive. Examples include polyhedral constraints, quadratic constraints and a PSD constraint<sup>3</sup>.

To tackle these complex constraints, the idea of optimization with a reduced number of projections was explored in several studies since (Mahdavi et al., 2012). In a recent paper (Chen et al., 2016), the authors show that for stochastic strongly convex optimization, the optimal convergence rate can be achieved using a logarithmic number of projections. In contrast, the developed theory in this paper implies that only one projection is sufficient to achieve the optimal convergence rate for strongly convex optimization, and a logarithmic number of projections can be used to accelerate convergence rates for non-strongly convex optimization. Cotter et al. (2016) proposed a stochastic algorithm for solving heavily constrained problems with many constraint functions by extending the work of (Mahdavi et al., 2012). Nonetheless, their focus is not to improve the convergence rates. Zhang et al. (2013) studied the smooth and strongly convex optimization and they proposed a stochastic algorithm with  $O(\kappa \log(T))$  projections and proved an  $O(1/T)$  convergence rate, where  $\kappa$  is the condition number and  $T$  is the total number of iterations. Nonetheless, if the condition number is high the number of projections could be very large. In addition, their algorithm utilizes the mini-batch to avoid frequent projections in stochastic optimization, which is different from the present paper.

We note that several recent works also exploit different forms of error bound conditions to improve the convergence (Wang & Lin, 2014; So, 2013; Hou et al., 2013; Zhou et al., 2015; Yang & Lin, 2016; Xu et al., 2016). Most notably, the technique used in our work is closely related to (Yang & Lin, 2016). However, for constrained optimization problems the methods in (Yang & Lin, 2016) still need to conduct projections at each iteration.

Finally, we comment on the differences between the proposed methods and the classical penalty methods that also move the constraint into the objective using a penalty function (Bertsekas, 1996). The major differences are that (i) the classical penalty methods typically require solving each subproblem exactly while our methods do not require that; and (ii) the classical penalty methods typically guarantee asymptotic convergence while our methods have explicit convergence rates.

<sup>1</sup>where  $\tilde{O}()$  suppresses a logarithmic factor.

<sup>2</sup>where the constraint function is assumed to be smooth.

<sup>3</sup>Indeed, a linear optimization over a PSD constraint is ill-posed because the PSD domain is unbounded.

### 3. Preliminaries

Let  $\Omega = \{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 0\}$  denote the constrained domain,  $\Omega_*$  denote the optimal solution set and  $f_*$  denote the optimal objective value. We denote by  $\nabla f(\mathbf{x})$  the gradient and by  $\partial f(\mathbf{x})$  the subgradient of a smooth or non-smooth function, respectively. When  $f(\mathbf{x})$  is a non-smooth function, we consider the problem as non-smooth constrained optimization. When both  $f(\mathbf{x})$  and  $c(\mathbf{x})$  are smooth, we consider the problem as smooth constrained optimization. A function  $f(\mathbf{x})$  is  $L$ -smooth if it has a Lipschitz continuous gradient, i.e.,  $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$ , where  $\|\cdot\|$  denotes the Euclidean norm. A function  $f(\mathbf{x})$  is  $\mu$ -strongly convex if it satisfies  $f(\mathbf{x}) \geq f(\mathbf{y}) + \partial f(\mathbf{y})^\top(\mathbf{x} - \mathbf{y}) + \frac{\mu}{2}\|\mathbf{x} - \mathbf{y}\|^2$ .

In the sequel,  $\text{dist}(\mathbf{x}, \Omega)$  denotes the distance of  $\mathbf{x}$  to a set  $\Omega$ , i.e.,  $\text{dist}(\mathbf{x}, \Omega) = \min_{\mathbf{u} \in \Omega} \|\mathbf{x} - \mathbf{u}\|$ . Let  $[s]_+$  be a hinge operator that is defined as  $[s]_+ = s$  if  $s \geq 0$ , and  $[s]_+ = 0$  if  $s < 0$ .

Throughout the paper, we make the the following assumptions to facilitate the development of our algorithms and theory.

**Assumption 1.** *For a convex minimization problem (1), we assume (i) there exists a positive value  $\rho > 0$  such that*

$$\min_{\substack{c(\mathbf{x})=0 \\ \mathbf{v} \in \partial c(\mathbf{x}), \mathbf{v} \neq 0}} \|\mathbf{v}\| \geq \rho, \quad (2)$$

or more generally there exists a constant  $\rho > 0$  for any  $\mathbf{x} \in \mathbb{R}^d$ , such that  $\mathbf{x}^\natural = \arg \min_{\mathbf{u} \in \mathbb{R}^d, c(\mathbf{u}) \leq 0} \|\mathbf{u} - \mathbf{x}\|^2$  satisfies

$$\|\mathbf{x}^\natural - \mathbf{x}\| \leq [c(\mathbf{x})]_+ / \rho. \quad (3)$$

(ii) there exists a strictly feasible solution such that  $c(\mathbf{x}) < 0$ ; (iii) both  $f(\mathbf{x})$  and  $c(\mathbf{x})$  are defined everywhere and are Lipschitz continuous with their Lipschitz constants denoted by  $G$  and  $G_c$ , respectively.

We make several remarks about the assumptions. The inequality in (2) is introduced in (Mahdavi et al., 2012), which is to ensure the distance from the final solution before projection to constrained domain  $\Omega$  is not too large. Note that the inequality in (3) is a more general condition than (2) as seen from the following lemma.

**Lemma 1.** *For any  $\mathbf{x} \in \mathbb{R}^d$ , let  $\mathbf{x}^\natural = \arg \min_{c(\mathbf{u}) \leq 0} \|\mathbf{u} - \mathbf{x}\|^2$ . If (2) holds, then (3) holds.*

The above lemma is implicit in the proof of (Mahdavi et al., 2012). We will provide more discussions about Assumption 1(i) - the key assumption, and exhibit the value of  $\rho$  for a number of commonly seen constraints (e.g., polyhedral, quadratic and PSD constraints). To make the presentation more fluent, we postpone these discussions to Section 6. The strict feasibility assumption (ii) allows us to explore the KKT condition of the projection problem shown below.

Assumption (iii) imposes mild Lipschitz continuity conditions on both  $f(\mathbf{x})$  and  $c(\mathbf{x})$ .

Traditional projected gradient descent methods need to solve the following projection at each iteration  $\Pi_\Omega[\mathbf{x}] = \arg \min_{c(\mathbf{u}) \leq 0} \|\mathbf{u} - \mathbf{x}\|^2$ . Conditional gradient methods (a.k.a. the Frank-Wolfe technique) need to solve the following linear optimization at each iteration  $\min_{\mathbf{u} \in \mathbb{R}^d, c(\mathbf{u}) \leq 0} \mathbf{u}^\top \nabla f(\mathbf{x})$ . For many constraint functions (see Section 6), solving the projection problem and the linear optimization could be very expensive.

### 4. A General Theory of Optimization with only one projection

In this section, we extend the idea of only one projection proposed in (Mahdavi et al., 2012) to a general theory, and then present optimization algorithms with only one projection for non-smooth and smooth optimization, respectively. To tackle the constraint, we introduce a penalty function  $h_\gamma(\mathbf{x})$  parameterized by  $\gamma$ , which obeys the following certificate: there exist constants  $C \geq 0$  and  $\lambda > G/\rho$  such that

$$\begin{aligned} h_\gamma(\mathbf{x}) &\geq \lambda[c(\mathbf{x})]_+, \forall \mathbf{x} \\ h_\gamma(\mathbf{x}) &\leq C\gamma, \forall \mathbf{x} \text{ such that } c(\mathbf{x}) \leq 0. \end{aligned} \quad (4)$$

From the above condition, it is clear that  $\gamma \geq 0$ . It is notable that the penalty function  $h_\gamma(\mathbf{x})$  will also depend on  $\lambda$ ; however it will be set to a constant value, thus the dependence on  $\lambda$  is omitted. We will construct such a penalty function  $h_\gamma(\mathbf{x})$  for non-smooth and smooth optimization in next two subsections. We propose to optimize the following augmented objective function

$$\min_{\mathbf{x} \in \mathbb{R}^d} F_\gamma(\mathbf{x}) = f(\mathbf{x}) + h_\gamma(\mathbf{x}). \quad (5)$$

We can employ any applicable optimization algorithms to optimize  $F_\gamma(\mathbf{x})$  pretending that there is no constraint, and finally obtain a solution  $\tilde{\mathbf{x}}_T$  that is not necessarily feasible. In order to obtain a feasible solution, we perform one projection to get  $\tilde{\mathbf{x}}_T = \Pi_\Omega(\tilde{\mathbf{x}}_T)$ . The following theorem allows us to convert the convergence of  $\tilde{\mathbf{x}}_T$  for  $F_\gamma(\mathbf{x})$  to that of  $\tilde{\mathbf{x}}_T$  for  $f(\mathbf{x})$ .

**Theorem 1.** *Let  $\mathcal{A}$  be any iterative optimization algorithm applied to  $\min_{\mathbf{x}} F_\gamma(\mathbf{x})$  with  $T$  iterations, which starts with  $\mathbf{x}_1$  and returns  $\tilde{\mathbf{x}}_T$  as the final solution, such that the following convergence of  $\tilde{\mathbf{x}}_T$  holds for any  $\mathbf{x} \in \mathbb{R}^d$*

$$F_\gamma(\tilde{\mathbf{x}}_T) - F_\gamma(\mathbf{x}) \leq B_T(\gamma; \mathbf{x}, \mathbf{x}_1), \quad (6)$$

where  $B_T(\gamma; \mathbf{x}, \mathbf{x}_1) \rightarrow 0$  when  $T \rightarrow \infty$ . Suppose that Assumption 1 hold, then

$$f(\tilde{\mathbf{x}}_T) - f(\mathbf{x}_*) \leq \frac{\lambda\rho}{\lambda\rho - G}(C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1)), \quad (7)$$

where  $\tilde{\mathbf{x}}_T = \Pi_\Omega[\tilde{\mathbf{x}}_T]$  and  $\mathbf{x}_*$  is an optimal solution to (1).

**Remark:** It is worth mentioning that we omit some constant factors in the convergence bound  $B_T(\gamma; \mathbf{x}, \mathbf{x}_1)$  that are irrelevant to our discussions. The notation  $B_T(\gamma; \mathbf{x}, \mathbf{x}_1)$  emphasizes that it is a function of  $\gamma$  and depends on  $\mathbf{x}_1$  and a target solution  $\mathbf{x}$  and it will be referred to as  $B_T$ . In the next several subsections, we will see that by carefully choosing the penalty function  $h_\gamma(\mathbf{x})$  we are able to provide nice convergence for smooth and non-smooth optimization with only one projection. In the above theorem, we assume the optimization algorithm  $\mathcal{A}$  is deterministic. However, a similar result can be easily extended to a stochastic optimization algorithm  $\mathcal{A}$ .

*Proof.* First, we consider  $c(\widehat{\mathbf{x}}_T) \leq 0$ , which implies that  $\widehat{\mathbf{x}}_T = \widetilde{\mathbf{x}}_T$ . Due to the certificate of  $h_\gamma(\mathbf{x})$ ,  $F_\gamma(\widetilde{\mathbf{x}}_T) \geq f(\widetilde{\mathbf{x}}_T)$  and  $F_\gamma(\mathbf{x}_*) \leq f(\mathbf{x}_*) + C\gamma$ . Hence  $f(\widetilde{\mathbf{x}}_T) \leq F_\gamma(\widehat{\mathbf{x}}_T) \leq F_\gamma(\mathbf{x}_*) + B_T(\gamma; \mathbf{x}_1, \mathbf{x}_*) \leq f(\mathbf{x}_*) + C\gamma + B_T(\gamma; \mathbf{x}_1, \mathbf{x}_*)$ . Then (7) follows due to  $\lambda\rho/(\lambda\rho - G) \geq 1$ . Next, we assume  $c(\widehat{\mathbf{x}}_T) > 0$ . Inequality (6) implies that

$$f(\widehat{\mathbf{x}}_T) + \lambda[c(\widehat{\mathbf{x}}_T)]_+ \leq f(\mathbf{x}_*) + C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1). \quad (8)$$

By Assumption 1(i), we have  $[c(\widehat{\mathbf{x}}_T)]_+ \geq \rho\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\|$ . Combined with (8) we have

$$\begin{aligned} \lambda\rho\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\| &\leq f(\mathbf{x}_*) - f(\widehat{\mathbf{x}}_T) + C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1) \\ &\leq G\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\| + C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1), \end{aligned}$$

where the last inequality follows that fact  $f(\mathbf{x}_*) - f(\widehat{\mathbf{x}}_T) \leq f(\mathbf{x}_*) - f(\widetilde{\mathbf{x}}_T) + f(\widetilde{\mathbf{x}}_T) - f(\widehat{\mathbf{x}}_T) \leq G\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\|$  because the Lipschitz property and  $f(\mathbf{x}_*) \leq f(\widetilde{\mathbf{x}}_T)$ . Therefore we have

$$\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\| \leq \frac{C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1)}{\lambda\rho - G}.$$

Finally, we obtain

$$\begin{aligned} f(\widehat{\mathbf{x}}_T) - f(\mathbf{x}_*) &\leq f(\widetilde{\mathbf{x}}_T) - f(\widehat{\mathbf{x}}_T) + f(\widehat{\mathbf{x}}_T) - f(\mathbf{x}_*) \\ &\leq G\|\widehat{\mathbf{x}}_T - \widetilde{\mathbf{x}}_T\| + C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1) \\ &\leq \frac{\lambda\rho}{\lambda\rho - G}(C\gamma + B_T(\gamma; \mathbf{x}_*, \mathbf{x}_1)). \end{aligned}$$

□

#### 4.1. Non-smooth Optimization

Since an optimal convergence rate for general non-smooth optimization with only one projection has been attained in (Mahdavi et al., 2012), in this subsection we present an optimal convergence result for strongly convex problems. For non-smooth optimization, we can choose

$$h(\mathbf{x}) = \lambda[c(\mathbf{x})]_+,$$

and hence  $\gamma = 0$ . We will use deterministic subgradient descent as an example to demonstrate the convergence for  $f(\mathbf{x})$ , though many other optimization algorithms designed for non-smooth optimization are applicable (e.g.,

the stochastic subgradient method). The update of subgradient descent method is given by the following

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \partial F(\mathbf{x}_t), \quad t = 1, \dots, T, \quad (9)$$

where  $\eta_t$  is an appropriate step size. If  $f(\mathbf{x})$  is  $\mu$ -strongly convex, the step size can be set as  $\eta_t = 1/(\mu t)$  and the final solution can be computed by the  $\alpha$ -suffix averaging  $\widehat{\mathbf{x}}_T = \frac{1}{\alpha T} \sum_{t=(1-\alpha)T+1}^T \mathbf{x}_t$  with  $\alpha > 0$  (Rakhlin et al., 2012), or by the polynomial decay averaging with  $\widehat{\mathbf{x}}_t = (1 - \frac{s+1}{s+t})\widehat{\mathbf{x}}_{t-1} + \frac{s+1}{s+t}\mathbf{x}_t$  and  $s \geq 1$  (Shamir & Zhang, 2013). Both schemes can attain  $B_T = O(1/(\mu T))$  for the convergence of  $F(\mathbf{x})$  when  $f(\mathbf{x})$  is  $\mu$ -strongly convex. Combining this with Theorem 1, we have the following convergence result with the proof omitted due to its simplicity.

**Corollary 2.** *Suppose that Assumption 1 holds and  $f(\mathbf{x})$  is  $\mu$ -strongly convex. Set  $F(\mathbf{x}) = f(\mathbf{x}) + \lambda[c(\mathbf{x})]_+$  with  $\lambda \geq G/\rho$ . Let (9) run for  $T$  iterations with  $\eta_t = 1/(\mu t)$ . Let  $\widehat{\mathbf{x}}_T$  be computed by  $\alpha$ -suffix averaging or the polynomial decay averaging. Then with only one projection  $\widetilde{\mathbf{x}}_T = \Pi_\Omega(\widehat{\mathbf{x}}_T)$ , we achieve*

$$f(\widetilde{\mathbf{x}}_T) - f_* \leq \frac{\lambda\rho}{\lambda\rho - G} \frac{(G + \lambda G_c)^2 O(1)}{\mu T}.$$

**Remark:** We note that the  $O(1/(\mu T))$  is also achieved for strongly convex optimization in (Zhang et al., 2013; Chen et al., 2016) but with a logarithmic number of projections. In contrast, Corollary 2 implies only one projection is sufficient to achieve the optimal convergence for strongly convex optimization.

#### 4.2. Smooth Optimization

For smooth optimization, we consider both  $f(\mathbf{x})$  and  $c(\mathbf{x})$  to be smooth<sup>4</sup>. Let the smoothness parameter of  $f(\mathbf{x})$  and  $c(\mathbf{x})$  be  $L_f$  and  $L_c$ , respectively. In order to ensure the augmented function  $F_\gamma(\mathbf{x})$  to be still a smooth function, we construct the following penalty function

$$h_\gamma(\mathbf{x}) = \gamma \ln(1 + \exp(\lambda c(\mathbf{x})/\gamma)). \quad (10)$$

The following proposition shows that  $h_\gamma(\mathbf{x})$  is a smooth function and obeys the condition in (4).

**Proposition 1.** *Suppose  $c(\mathbf{x})$  is  $L_c$ -smooth and  $G_c$ -Lipschitz continuous. The penalty function in (10) is a  $(\lambda L_c + \frac{\lambda^2 G_c^2}{4\gamma})$ -smooth function and satisfies (i)  $h_\gamma(\mathbf{x}) \geq \lambda[c(\mathbf{x})]_+$  and (ii)  $h_\gamma(\mathbf{x}) \leq \gamma \ln 2$ ,  $\forall \mathbf{x}$  such that  $c(\mathbf{x}) \leq 0$ .*

Then  $F_\gamma(\mathbf{x})$  is a smooth function and its smoothness parameter is given by  $L_F = L_f + \lambda L_c + \frac{\lambda^2 G_c^2}{4\gamma}$ . Next, we will

<sup>4</sup>it can be extended to when  $f(\mathbf{x})$  is non-smooth but its proximal mapping can be easily solved.

establish the convergence for  $f(\mathbf{x})$  using Nesterov's optimal accelerated gradient (NAG) methods. The update of one variant of NAG can be written as follows

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{y}_t - \nabla F_\gamma(\mathbf{y}_t)/L_F \\ \mathbf{y}_{t+1} &= \mathbf{x}_{t+1} + \beta_{t+1}(\mathbf{x}_{t+1} - \mathbf{x}_t), \end{aligned} \quad (11)$$

where the value of  $\beta_t$  can be set to different values depending on whether  $f(\mathbf{x})$  is strongly convex or not (see Corollary 3). Previous work have established the convergence of  $\tilde{\mathbf{x}}_T = \mathbf{x}_T$  for  $F_\gamma(\mathbf{x})$ , in particular  $B_T = O(\frac{L_F}{T^{\frac{1}{2}}})$  for smooth non-strongly convex optimization and  $B_T = O\left(L_F \exp(-T\sqrt{\frac{\mu}{L_F}})\right)$  for smooth and strongly convex optimization. By combining these results with Theorem 1 and appropriately setting  $\gamma$ , we can achieve the following convergence of  $\tilde{\mathbf{x}}_T$  for  $f(\mathbf{x})$ .

**Corollary 3.** *Suppose that Assumption 1 holds,  $\text{dist}(\mathbf{y}_0, \Omega_*) \leq D$ ,  $f(\mathbf{x})$  is  $L_f$ -smooth and  $c(\mathbf{x})$  is  $L_c$ -smooth. Set  $F_\gamma(\mathbf{x}) = f(\mathbf{x}) + h_\gamma(\mathbf{x})$  with  $\lambda > G/\rho$  and  $h_\gamma(\mathbf{x})$  being (10). Let (11) run for  $T$  iterations and  $\tilde{\mathbf{x}}_T = \Pi_\Omega(\mathbf{x}_T)$ .*

- If  $f(\mathbf{x})$  is convex, we can set  $\gamma = \frac{\lambda G_c D}{(T+1)\sqrt{2\ln 2}}$ ,  $\beta_t = \frac{\tau_{t-1}-1}{\tau_t}$ , where  $\tau_t = \frac{1+\sqrt{1+4\tau_{t-1}^2}}{2}$  with  $\tau_0 = 1$ , and achieve  $f(\tilde{\mathbf{x}}_T) - f_* \leq \frac{\lambda\rho}{\lambda\rho - G} \left[ \frac{\lambda G_c D \sqrt{2\ln 2}}{T+1} + \frac{(L_f + \lambda L_c) D^2}{(T+1)^2} \right]$
- If  $f(\mathbf{x})$  is  $\mu$ -strongly convex, we can set  $\gamma = \frac{1}{T^{2\alpha}}$  with  $\alpha \in (1/2, 1)$  and  $\beta_t = \frac{\sqrt{L_F} - \sqrt{\mu}}{\sqrt{L_F} + \sqrt{\mu}}$ , and achieve

$$f(\tilde{\mathbf{x}}_T) - f_* \leq O\left(\frac{1}{T^{2\alpha}} + \frac{1}{T^{4\alpha}}\right),$$

as long as  $T \geq \left(\frac{L_f + \lambda L_c + \lambda^2 G_c^2 / 4}{\mu}\right)^{\frac{1}{2(1-\alpha)}} (4\alpha \ln T)^{\frac{1}{1-\alpha}}$ .

**Remark:** The convergence results above indicate an  $O(1/\epsilon)$  iteration complexity for smooth optimization and  $O(1/\epsilon^{1/(2\alpha)})$  with  $\alpha \in (1/2, 1)$  for smooth and strongly convex optimization with only one projection. All omitted proofs can be found in (Yang et al., 2017).

## 5. Improved Convergence for Non-strongly Convex Optimization

In this section, we will develop improved convergence for non-strongly convex optimization at a price of a logarithmic number of projections by considering an additional condition on the target problem. To facilitate the presentation, we first introduce some notations. The  $\epsilon$ -sublevel set  $\mathcal{S}_\epsilon$  and  $\epsilon$ -level set  $\mathcal{L}_\epsilon$  of the problem (1) are denoted by  $\mathcal{S}_\epsilon = \{\mathbf{x} \in \Omega : f(\mathbf{x}) \leq f_* + \epsilon\}$ , and  $\mathcal{L}_\epsilon = \{\mathbf{x} \in \Omega : f(\mathbf{x}) = f_* + \epsilon\}$ , respectively. Let  $\mathbf{x}_\epsilon^\dagger$  denote the closest point in the  $\epsilon$ -sublevel set  $\mathcal{S}_\epsilon$  to  $\mathbf{x} \in \Omega$ , i.e.,

$$\mathbf{x}_\epsilon^\dagger = \arg \min_{\mathbf{u} \in \Omega} \|\mathbf{u} - \mathbf{x}\|^2, \quad \text{s.t.} \quad f(\mathbf{u}) \leq f_* + \epsilon. \quad (12)$$

Let  $\mathbf{x}^*$  denote the closest optimal solution in  $\Omega_*$  to  $\mathbf{x}$ , i.e.,  $\mathbf{x}^* = \arg \min_{\mathbf{u} \in \Omega_*} \|\mathbf{u} - \mathbf{x}\|^2$ .

In this section, we will make the following additional assumption about the problem (1).

**Assumption 2.** *For a convex minimization problem (1), we assume (i) there exist  $\mathbf{x}_0 \in \Omega$  and  $\epsilon_0 \geq 0$  such that  $f(\mathbf{x}_0) - \min_{\mathbf{x} \in \Omega} f(\mathbf{x}) \leq \epsilon_0$ ; (ii)  $\Omega_*$  is a non-empty convex compact set; (iii) the optimization problem (1) satisfies a local error bound condition, i.e., there exist  $\theta \in (0, 1]$  and  $\sigma > 0$  such that for any  $\mathbf{x} \in \mathcal{S}_\epsilon$  we have  $\text{dist}(\mathbf{x}, \Omega_*) \leq \sigma(f(\mathbf{x}) - f_*)^\theta$  where  $\Omega_*$  denotes the optimal set and  $f_*$  denotes the optimal value.*

**Remark:** we would like to remark that the new assumption only imposes mild conditions on the problem. In particular, Assumption 2 (i) supposes there is a lower bound of the optimal value  $f_*$ , which usually holds in machine learning applications where the objective function is non-negative; Assumption 2 (ii) ensures that  $\mathcal{S}_\epsilon$  is also bounded (Rockafellar, 1970), therefore the  $\sigma$  in the local error bound is finite, which can be easily satisfied for a norm regularized or constraint problems; the local error bound condition holds for a broad family of functions (e.g., semi-algebraic functions or real subanalytic functions (Jerome Bolte, 2015; Yang & Lin, 2016)). In Section 7, we will also demonstrate several applications of the improved algorithms proposed in this section by establishing the local error bound condition.

Although the local error bound condition is much weaker than the strong convexity assumption, below we will propose novel algorithms leveraging this condition with faster convergence and only a logarithmic number of projections.

### 5.1. Non-smooth Optimization

To establish an improved convergence for non-smooth optimization, we develop a new algorithm shown in Algorithm 1 based on subgradient descent (GD) method, to which we refer as LoPGD. The algorithm runs for  $K$  epochs and each epoch employs GD for minimizing  $F(\mathbf{x}) = f(\mathbf{x}) + \lambda[c(\mathbf{x})]_+$  with a feasible solution  $\mathbf{x}_{k-1} \in \Omega$  as a starting point and  $t$  iterations of updates. At the end of each epoch, the averaged solution  $\tilde{\mathbf{x}}_k$  is projected into the constrained domain  $\Omega$  and the solution  $\mathbf{x}_k$  will be used as the starting point for next epoch. The step size  $\eta_k$  is decreased by half every epoch starting from a given value  $\eta_1$ . The theorem below establishes the iteration complexity of LoPGD and also exhibits the values of  $K$ ,  $t$  and  $\eta_1$ . To simplify notations, we let  $p = \frac{\lambda\rho}{\lambda\rho - G}$  and  $\tilde{G} = G + \lambda G_c$ .

**Theorem 4.** *Suppose Assumptions 1 and 2 hold. Let  $\eta_1 = \frac{\epsilon_0}{2p\tilde{G}^2}$ ,  $K = \lceil \log_2(\epsilon_0/\epsilon) \rceil$  and  $t = \frac{4\sigma^2 p^2 \tilde{G}^2}{\epsilon^2(1-\theta)}$  in Algorithm 1, where  $\theta$  and  $\sigma$  are constants appearing in the local error bound condition. Then  $f(\mathbf{x}_K) - f_* \leq 2\epsilon$ .*

**Remark:** Since the projection is only conducted at the end of each epoch and the total number of epochs is at

**Algorithm 1** LoPGD

---

```

1: INPUT:  $K \in \mathbb{N}^+, t \in \mathbb{N}^+, \eta_1$ 
2: Initialization:  $\mathbf{x}_0 \in \Omega, \epsilon_0$ 
3: for  $k = 1, 2, \dots, K$  do
4:   Let  $\mathbf{x}_1^k = \mathbf{x}_{k-1}$ 
5:   for  $s = 1, 2, \dots, t - 1$  do
6:     Update  $\mathbf{x}_{s+1}^k = \mathbf{x}_s^k - \eta_k \partial F(\mathbf{x}_s^k)$ 
7:   end for
8:   Let  $\widehat{\mathbf{x}}_k = \sum_{s=1}^t \mathbf{x}_s^k / t$ 
9:   Let  $\mathbf{x}_k = \Pi_{\Omega}[\widehat{\mathbf{x}}_k]$  and  $\eta_{k+1} = \eta_k / 2$ 
10: end for
    
```

---

**Algorithm 2** LoPNAG

---

```

1: INPUT:  $K \in \mathbb{N}^+, t_1, \dots, t_K \in \mathbb{N}^+, \gamma_1$ 
2: Initialization:  $\mathbf{x}_0 \in \Omega, \epsilon_0$ 
3: for  $k = 1, 2, \dots, K$  do
4:   Let  $\mathbf{y}_0^k = \mathbf{x}_{k-1}$ 
5:   for  $s = 0, 1, 2, \dots, t_k - 1$  do
6:     Update  $\mathbf{x}_{s+1}^k = \mathbf{y}_s^k - \frac{1}{L_k} \nabla F_{\gamma_k}(\mathbf{x}_s^k)$ 
7:     Update  $\mathbf{y}_{s+1}^k = \mathbf{x}_{s+1}^k + \beta_{s+1}(\mathbf{x}_{s+1}^k - \mathbf{x}_s^k)$ 
8:   end for
9:   Let  $\widehat{\mathbf{x}}_k = \mathbf{x}_{t_k}^k, \mathbf{x}_k = \Pi_{\Omega}[\widehat{\mathbf{x}}_k]$  and  $\gamma_{k+1} = \gamma_k / 2$ 
10: end for
    
```

---

most  $K = \lceil \log_2(\epsilon_0/\epsilon) \rceil$ , so the total number of projections is only a logarithmic number  $K$ . The iteration complexity in Theorem 4 is  $\widetilde{O}(1/\epsilon^{2(1-\theta)})$  that improves the standard result of  $O(1/\epsilon^2)$  without strong convexity. With  $\theta = 1/2$ , we can achieve  $\widetilde{O}(1/\epsilon)$  iteration complexity with only  $O(\log(1/\epsilon))$  projections.

## 5.2. Smooth Optimization

Similar to non-smooth optimization, we also develop a new algorithm based on NAG shown in Algorithm 2, where  $F_{\gamma}(\mathbf{x})$  is defined using  $h_{\gamma}(\mathbf{x})$  in (10),  $L_k = L_{F_{\gamma_k}}$  is the smoothness parameter of  $F_{\gamma_k}$  and  $\beta_s = \frac{\tau_{s-1}-1}{\tau_s}, s = 1, \dots$ , is a sequence with  $\tau_s$  updated as in Corollary 3. We refer to this algorithm as LoPNAG. The key idea is to use to a sequence of reducing values for  $\gamma_k$  instead of using a small value as in Corollary 3, and solve each augmented unconstrained problem  $F_{\gamma_k}(\mathbf{x})$  approximately with one projection. The theorem below exhibits the iteration complexity of LoPNAG and reveals the values of  $K, \gamma_1$  and  $t_1, \dots, t_K$ . To simplify notations, we let  $\bar{L} = L_f + \lambda L_c$ .

**Theorem 5.** *Suppose Assumptions 1 and 2 hold and  $f(\mathbf{x})$  is  $L_f$ -smooth and  $c(\mathbf{x})$  is  $L_c$ -smooth. Let  $\gamma_1 = \frac{\epsilon_0}{6p \ln 2}, K = \lceil \log_2(\epsilon_0/\epsilon) \rceil$  and  $t_k = \frac{\sigma}{\epsilon^{1-\theta}} \max\{\lambda G_c p \sqrt{18 \ln 2}, \sqrt{12(L_f + \lambda L_c)\epsilon_0/2^{k-1}}\}$  in Algorithm 2, where  $\theta$  and  $\sigma$  are constants appearing in the local error bound condition. Then  $f(\mathbf{x}_K) - f_* \leq 2\epsilon$ .*

**Remark:** It is not difficult to show that the total number of iterations is bounded by  $\widetilde{O}(1/\epsilon^{1-\theta})$ , which improves the one in Corollary 3 without strong convexity. If  $f(\mathbf{x})$  is a

simple non-smooth function whose proximal mapping can be easily computed (e.g.,  $\ell_1$  norm), we can replace step 6 in Algorithm 2 by a proximal mapping to handle  $f(\mathbf{x})$ , which gives the same convergence result in Theorem 5. An example is presented in Section 7 for compressive sensing with  $\theta = 1/2$ .

## 6. Discussion of Assumption 1 (i)

One might note that a key condition for developing the theory with reduced projections is Assumption 1 (i). Although Mahdavi et al. (2012) has briefly mentioned that the condition can be satisfied for a PSD cone or a Polytope (a bounded polyhedron), their discussion lacks of details in particular on the value of  $\rho$  in (2) or (3). Below, we discuss the condition in details about three types of constraints.

**Polyhedral constraints.** First, we show that when  $c(\mathbf{x})$  is a polyhedral function, i.e., its epigraph is a polyhedron (not necessarily bounded), the inequality (3) is satisfied. To this end, we explore the polyhedral error bound (PEB) condition (Gilpin et al., 2012; Yang & Lin, 2016). In particular, if we consider an optimization problem,  $\min_{\mathbf{x} \in \mathbb{R}^d} h(\mathbf{x})$ , where the epigraph of  $h(\mathbf{x})$  is polyhedron. Let  $\mathcal{H}_*$  denote the optimal set and  $h_*$  denote the optimal value of the problem above. The PEB says that there exists  $\rho > 0$  such that for any  $\mathbf{x} \in \mathbb{R}^d$

$$\text{dist}(\mathbf{x}, \mathcal{H}_*) \leq (h(\mathbf{x}) - h_*)/\rho. \quad (13)$$

To show that the inequality (3) holds for a polyhedral function  $c(\cdot)$ , we can consider the optimization problem  $\min_{\mathbf{x} \in \mathbb{R}^d} [c(\mathbf{x})]_+$ . The optimal set of the above problem is given by  $\mathcal{H}_* = \{\mathbf{x} \in \mathbb{R}^d : c(\mathbf{x}) \leq 0\}$ . For any  $\mathbf{x}$  such that  $c(\mathbf{x}) > 0$ , let  $\mathbf{x}^{\natural} = \arg \min_{c(\mathbf{u}) \leq 0} \|\mathbf{u} - \mathbf{x}\|^2$  be the closest point in the optimal set to  $\mathbf{x}$ . Therefore if  $c(\cdot)$  is a polyhedral function so does  $[c(\mathbf{x})]_+$ , by the PEB condition (13) there exists a  $\rho > 0$  such that  $\|\mathbf{x} - \mathbf{x}^{\natural}\| \leq ([c(\mathbf{x})]_+ - \min_{\mathbf{x}} [c(\mathbf{x})]_+)/\rho = [c(\mathbf{x})]_+/\rho$ . Let us consider a concrete example, where the problem has a set of affine inequalities  $\mathbf{c}_i^{\top} \mathbf{x} - b_i \leq 0, i = 1, \dots, m$ . There are two methods to encode this into a single constraint function  $c(\mathbf{x}) \leq 0$ . The first method is to use  $c(\mathbf{x}) = \max_{1 \leq i \leq m} \mathbf{c}_i^{\top} \mathbf{x} - b_i$ , which is a polyhedral function and therefore satisfies (3). The second method is to use  $c(\mathbf{x}) = \|[C\mathbf{x} - \mathbf{b}]_+\|$ , where  $[\mathbf{a}]_+ = \max(0, \mathbf{a})$  and  $C = (\mathbf{c}_1, \dots, \mathbf{c}_m)^{\top}$ . Thus  $[c(\mathbf{x})]_+ = \|[C\mathbf{x} - \mathbf{b}]_+\|$ . The inequality (3) is then guaranteed by Hoffman's bound and the parameter  $\rho$  is given by the minimum non-zero eigenvalue of  $C^{\top}C$  (Wang & Lin, 2014). Note that the projection onto a polyhedron is a linear constrained quadratic programming problem, and the linear optimization over a polyhedron is a linear programming problem. Both have polynomial time complexity that would be high if  $m$  and  $d$  are large (Karmarkar, 1984; Kozlov et al., 1980).

**Quadratic constraint.** A quadratic constraint can take the form of  $\|A\mathbf{x} - \mathbf{y}\|^2 \leq \tau$ , where  $A \in \mathbb{R}^{m \times d}$  and  $\mathbf{y} \in \mathbb{R}^m$ . Such a constraint appears in compressive sensing (Candès & Wakin, 2008)<sup>5</sup>, where the goal is to reconstruct a sparse high-dimensional vector  $\mathbf{x}$  from a small number of noisy measurements  $\mathbf{y} = A\mathbf{x} + \varepsilon \in \mathbb{R}^m$  with  $m \ll d$ . The corresponding optimization problem is

$$\min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{x}\|_1, \quad \text{s.t.} \quad \|A\mathbf{x} - \mathbf{y}\|^2 \leq \tau. \quad (14)$$

where  $\tau \geq \|\varepsilon\|^2$  is an upper bound on the magnitude of the noise. To check the Assumption 1(i), we note that  $c(\mathbf{x}) = \|A\mathbf{x} - \mathbf{y}\|^2 - \tau$  and  $\nabla c(\mathbf{x}) = A^\top(A\mathbf{x} - \mathbf{y})$ . Let us consider that  $A$  has a full row rank<sup>6</sup> and denote by  $\mathbf{v} = A\mathbf{x} - \mathbf{y}$ , then on the boundary  $c(\mathbf{x}) = 0$  we have  $\|\mathbf{v}\| = \sqrt{\tau}$  and  $\|A^\top \mathbf{v}\| \geq \sqrt{\tau} \lambda_{\min}(AA^\top)$ , where  $\lambda_{\min}(AA^\top) > 0$  is the minimum eigenvalue of  $AA^\top \in \mathbb{R}^{m \times m}$ . Therefore the Assumption 1(i) is satisfied with  $\rho = \sqrt{\tau} \lambda_{\min}(AA^\top)$ . It is notable that the projection and the linear optimization under the quadratic constraint require solving a quadratic programming problem and therefore could be expensive.

**PSD constraint.** A PSD constraint  $X \succeq 0$  for  $X \in \mathbb{R}^{d \times d}$  can be written as an inequality constraint  $-\lambda_{\min}(X) \leq 0$ , where  $\lambda_{\min}(X)$  denotes the minimum eigen-value of  $X$ . The subgradient of  $c(X) = -\lambda_{\min}(X)$  when  $\lambda_{\min}(X) = 0$  is given by  $\text{Conv}\{-\mathbf{u}\mathbf{u}^\top \mid \|\mathbf{u}\| = 1, X\mathbf{u} = 0\}$ , i.e., the convex hull of the outer products of normalized vectors in the null space of the matrix  $X$ . In (Yang et al., 2017), we show that if the dimension of the null space of  $X$  is  $r$  with  $1 \leq r \leq d$ , the norm of the subgradient of  $c(X)$  on the boundary  $c(X) = 0$  is lower bounded by  $\rho = \frac{1}{\sqrt{r}} \geq \frac{1}{\sqrt{d}}$ . Finally, we note that computing a subgradient of  $[c(X)]_+$  only needs to compute one eigen-vector corresponding to the smallest eigen-value. In contrast, both projection and linear optimization under a PSD constraint could be very expensive for high-dimensional problems. In particular, the projection onto a PSD domain needs to conduct a singular value decomposition. The linear optimization over a PSD cone is ill-posed due to that PSD cone is not compact (the solution is either 0 or infinity). One may add an artificial constraint on the upper bound of the eigen-values. According to (Jaggi, 2013), the time complexity for solving this linear optimization problem approximately up to an accuracy level  $\epsilon'$  is  $O(Nd^{1.5}/\epsilon'^{2.5})$  with  $N$  being the number of non-zeros in the gradient and  $\epsilon'$  decreasing iteratively required in the Frank-Wolfe method, which could be much more expensive especially for high-dimensional problems and in later iterations than computing the first eigen-pairs at each iteration in our methods.

<sup>5</sup>Here we use the square constraint to make it a smooth function so that the proposed algorithms for smooth optimization are applicable by using proximal gradient mapping to handle the  $\ell_1$  norm.

<sup>6</sup>which is reasonable because  $m \ll d$ .

## 7. Applications

### 7.1. Compressive Sensing

We first consider a compressive sensing problem in (14). Becker et al. (2011) proposed an optimization algorithm based on the Nesterov's smoothing and the Nesterov's optimal method for the smoothed problem, known as NESTA. It needs to perform the projection into the domain  $\|A\mathbf{x} - \mathbf{y}\|^2 \leq \tau$  at every iteration and has an iteration complexity of  $O(1/\epsilon)$ . In contrast, the presented algorithm with *only one projection* in Section 4.2 using Nesterov's accelerated proximal gradient method (Beck & Teboulle, 2009) to solve the unconstrained problem enjoys an iteration complexity of  $O(1/\epsilon)$ . Moreover, we present a theorem below showing that the problem (14) satisfies the local error bound condition with  $\theta = 1/2$ , and hence the presented LoPNAG enjoys an  $\tilde{O}(1/\sqrt{\epsilon})$  iteration complexity with only a logarithmic number of projections.

**Theorem 6.** *Let  $f(\mathbf{x}) = \|\mathbf{x}\|_1$ ,  $c(\mathbf{x}) = \|A\mathbf{x} - \mathbf{y}\|^2 - \tau$ ,  $\Omega_*$  denote the optimal set and  $f_*$  be the optimal solution to (14). Assume that there exists  $\mathbf{x}_0$  such that  $\|A\mathbf{x}_0 - \mathbf{y}\|^2 < \tau$  and  $0 \notin \Omega_*$ . Then for any  $\epsilon > 0$ ,  $\mathbf{x} \in \mathbb{R}^d$  such that  $c(\mathbf{x}) \leq 0$  and  $f(\mathbf{x}) \leq f_* + \epsilon$ , there exists  $0 < \sigma < \infty$  such that  $\text{dist}(\mathbf{x}, \Omega_*) \leq \sigma(f(\mathbf{x}) - f_*)^{1/2}$ . Hence, LoPNAG can have an iteration complexity of  $\tilde{O}(1/\sqrt{\epsilon})$  with only  $O(\log(1/\epsilon))$  projections.*

Next, we demonstrate the effectiveness of the LoPNAG for solving the compressive sensing problem in (14) by comparing with NESTA. We generate a synthetic data for testing. In particular, we generate a random measurement matrix  $A \in \mathbb{R}^{m \times d}$  with  $m = 1000$  and  $d = 5000$ . The entries of the matrix  $A$  are generated independently with the uniform distribution over the interval  $[-1, +1]$ . The vector  $\mathbf{x}_* \in \mathbb{R}^d$  is generated with the same distribution at 100 randomly chosen coordinates. The noise  $\varepsilon \in \mathbb{R}^m$  is a dense vector with independent random entries with the uniform distribution over the interval  $[-\zeta, \zeta]$ , where  $\zeta$  is the noise magnitude and is set to 0.01. Finally the vector  $\mathbf{y}$  was obtained as  $\mathbf{y} = A\mathbf{x}_* + \varepsilon$ .

We use the Matlab package of NESTA<sup>7</sup>. For fair comparison, we also use the projection code in the NESTA package for conducting projection. To handle the unknown smoothness parameter in the proposed algorithm, we use the backtracking technique (Beck & Teboulle, 2009). The parameter  $\gamma$  is initially set to 0.001 and decreased by half every 5000 iterations after a projection and the target smoothing parameter in NESTA is set to  $10^{-5}$ . For the value of  $\lambda$  in LoPNAG, we tune it from its theoretical value to several smaller values and choose the one that yields the fastest convergence. We report the results in Table 7.1, which include different number of iterations, the corresponding

<sup>7</sup><http://statweb.stanford.edu/~candes/nesta/>

Table 1. LoPNAG vs. NESTA for solving the compressive sensing problem.

LoPNAG				NESTA			
Iters - Projs	Rec. Err.	Objective	Time (s)	Iters - Projs	Rec. Err.	Objective	Time (s)
5000 - 1	0.018017	52.042878	18.04	1000 - 2000	0.137798	52.703275	48.49
10000 - 2	0.018038	52.042418	35.88	3000 - 6000	0.018669	52.050051	93.84
15000 - 3	0.018043	52.042358	53.09	5000 - 10000	0.018659	52.050046	245.23
20000 - 4	0.018043	52.042358	70.24	8000 - 16000	0.018657	52.050045	404.72
25000 - 5	0.018043	52.042358	87.32	10000 - 20000	0.018657	52.050044	501.65

Table 2. LoPGD vs. OPGD and PGD for solving the considered distance metric learning problem.

LoPGD			OPGD			PGD		
Iters - Projs	Objective	Time (h)	Iters - Projs	Objective	Time (h)	Iters - Projs	Objective	Time (h)
1000 - 1	0.0953	0.22	1000 - 1	0.1707	0.20	1000 - 1000	0.1491	7.97
2000 - 2	0.0695	0.43	2000 - 1	0.1583	0.40	2000 - 2000	0.1278	15.46
4000 - 4	0.0494	0.87	4000 - 1	0.1469	0.80	4000 - 4000	0.1072	29.39
6000 - 6	0.0428	1.33	6000 - 1	0.1398	1.22	6000 - 6000	0.0957	43.36
8000 - 8	0.0405	1.89	8000 - 1	0.1343	1.64	8000 - 8000	0.0879	57.43

number of projections, the recovery error of the found solution compared to the underlying true sparse solution, the objective value (i.e., the  $\ell_1$  norm of the found solution) and the running time. Note that each iteration of NESTA requires two projections because it maintains two extra sequence of solutions. From the results, we can see that LoPNAG converges significantly faster than NESTA. Even with only one projection, we are able to obtain a better solution than that of NESTA after running 10000 iterations.

## 7.2. High-dimensional Distance Metric Learning

Consider the following distance metric learning problem:

$$\min_{A \succeq 0} \frac{1}{2|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} (1 - y_{ij} - \|\mathbf{x}_i - \mathbf{x}_j\|_A^2)^2 + \tau \|A\|_1^{\text{off}}, \quad (15)$$

where  $\mathcal{E}$  denotes all pairs of training examples,  $y_{ij} = 1$  indicates  $\mathbf{x}_i, \mathbf{x}_j$  belong to the same class and  $y_{ij} = -1$  indicates they belong to different classes,  $\|\mathbf{z}\|_A^2 = \mathbf{z}^\top A \mathbf{z}$  and  $\|A\|_1^{\text{off}} = \sum_{i \neq j} |A_{ij}|$ . We note that such a formulation is useful for high dimensional problems due to the  $\ell_1$  regularizer. A similar formulation with different forms of loss function has been adopted in literature (Qi et al., 2009). We consider the square loss because it gives us faster convergence with a logarithmic number of projections by LoPGD. Due to the presence of the non-smooth PSD constraint and the  $\ell_1$  regularizer, Nesterov’s accelerated proximal gradient methods can not be applied efficiently to solving (15) and the augmented unconstrained problem. Nevertheless, we can apply the proposed LoPGD method for solving the problem with a logarithmic number of projections. Regarding the constant  $\theta$  in the local error bound condition for (15), it still remains an open problem. Nonetheless, a local error bound condition with  $\theta = 0.5$  might be established under certain regularity condition of the problem (Zhou & So, 2015; Cui et al., 2017). For example, Cui et al. (2017) provided a direct analysis of a local error bound condition with  $\theta = 0.5$  for a class of constrained convex symmetric matrix optimization problems

regularized by nonsmooth spectral functions (including the indicator function of a PSD constraint). They established sufficient conditions (Theorem 16) for a local error bound condition with  $\theta = 0.5$  to hold, which reduces to a regularity condition for (15) depending on the optimal solutions of the problem. A thorough analysis of the regularity condition is much more involved and left as an open problem.

Next, we demonstrate the empirical performance of LoPGD for solving (15). We use the colon-cancer data available on libsvm web portal, which has 2000 features and 62 examples. Forty examples are used as training examples to generate 780 pairs to learn the distance metric. The regularization parameter is set to  $\tau = 0.001$ . We compare LoPGD, gradient descent method with only one projection (referred to as OPGD), and standard projected GD (referred to as PGD). The step size in PGD and OPGD is set to  $\eta_0/\sqrt{t}$ , where  $t$  is the iteration index. We use the same tuned initial step size for all algorithms. The number of iterations per-epoch in LoPGD is set to 1000. The penalization parameter  $\lambda$  in both OPGD and LoPGD is tuned and set to 10. In Table 2, we report the objective values, the #of iterations/projections, and running time across the first 8000 iterations. We can see that LoPGD converges dramatically faster than PGD and also much faster than OPGD.

## 8. Conclusion

We have developed a general theory of optimization with only one projection for a family of inequality constrained convex optimization problems. It yields an improved iteration complexity for smooth optimization compared with non-smooth optimization. By exploring the local error bound condition, we further develop new algorithms with a logarithmic number of projections and achieve better convergence for both smooth and non-smooth optimization without strong convexity assumption. Applications in compressive sensing and distance metric learning demonstrate the effectiveness of the proposed improved algorithms.



## Acknowledgements

We are grateful to all anonymous reviewers for their helpful comments. T. Yang is partially supported by National Science Foundation (IIS-1463988, IIS-1545995). L. Zhang thanks the support from NSFC (61603177) and JiangsuSF (BK20160658).

## References

- Beck, Amir and Teboulle, Marc. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2:183–202, 2009.
- Becker, Stephen, Bobin, Jérôme, and Candès, Emmanuel J. NESTA: A fast and accurate first-order method for sparse recovery. *SIAM J. Img. Sci.*, 4:1–39, 2011. ISSN 1936-4954.
- Bertsekas, Dimitri P. *Constrained Optimization and Lagrange Multiplier Methods (Optimization and Neural Computation Series)*. Athena Scientific, 1 edition, 1996. ISBN 1886529043.
- Candès, Emmanuel J. and Wakin, Michael B. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008.
- Chen, Jianhui, Yang, Tianbao, Lin, Qihang, Zhang, Lijun, and Chang, Yi. Optimal stochastic strongly convex optimization with a logarithmic number of projections. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Cotter, Andrew, Gupta, Maya R., and Pfeifer, Jan. A light touch for heavily constrained SGD. In *Proceedings of the 29th Conference on Learning Theory (COLT)*, pp. 729–771, 2016.
- Cui, Ying, Ding, Chao, and Zhao, Xinyuan. Quadratic growth conditions for convex matrix optimization problems associated with spectral functions. *CoRR*, 2017.
- Frank, Marguerite and Wolfe, Philip. An algorithm for quadratic programming. *Naval Research Logistics (NRL)*, 3:149–154, 1956.
- Garber, Dan and Hazan, Elad. Faster rates for the frank-wolfe method over strongly-convex sets. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 541–549, 2015.
- Gilpin, Andrew, Peña, Javier, and Sandholm, Tuomas. First-order algorithm with  $\log(1/\epsilon)$  convergence for  $\epsilon$ -equilibrium in two-person zero-sum games. *Math. Program.*, 133(1-2):279–298, 2012.
- Hazan, Elad and Kale, Satyen. Projection-free online learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.
- Hou, Ke, Zhou, Zirui, So, Anthony Man-Cho, and Luo, Zhi-Quan. On the linear convergence of the proximal gradient method for trace norm regularization. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 710–718, 2013.
- Jaggi, Martin. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 427–435, 2013.
- Jerome Bolte, Trong Phong Nguyen, Juan Peypouquet Bruce Suter. From error bounds to the complexity of first-order descent methods for convex functions. *CoRR*, abs/1510.08234, 2015.
- Karmarkar, N. A new polynomial-time algorithm for linear programming. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pp. 302–311, 1984.
- Kozlov, M.K., Tarasov, S.P., and Khachiyan, L.G. Polynomiale Loesbarkeit der konvexen quadratischen Programmierung. *Zh. Vychisl. Mat. Mat. Fiz.*, 20:1319–1323, 1980. ISSN 0044-4669.
- Lacoste-Julien, Simon, Jaggi, Martin, Schmidt, Mark, and Pletscher, Patrick. Block-coordinate frank-wolfe optimization for structural svms. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 53–61, 2013.
- Mahdavi, M., Yang, T., Jin, R., and Zhu, S. Stochastic gradient descent with only one projection. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 503–511, 2012.
- Nesterov, Yurii. *Introductory lectures on convex optimization: a basic course*, volume 87 of *Applied optimization*. Kluwer Academic Publishers, 2004.
- Qi, Guo-Jun, Tang, Jinhui, Zha, Zheng-Jun, Chua, Tat-Seng, and Zhang, Hong-Jiang. An efficient sparse metric learning in high-dimensional space via  $l_1$ -penalized log-determinant regularization. In *Proceedings of the 26th international conference on Machine learning (ICML)*, pp. 106, 2009.
- Rakhlin, Alexander, Shamir, Ohad, and Sridharan, Karthik. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th international conference on Machine learning (ICML)*, 2012.

- Rockafellar, R.T. *Convex Analysis*. Princeton mathematical series. Princeton University Press, 1970.
- Shamir, Ohad and Zhang, Tong. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pp. 71–79, 2013.
- So, Anthony Man-Cho. Non-asymptotic convergence analysis of inexact gradient methods for machine learning without strong convexity. *CoRR*, abs/1309.0113, 2013.
- Wang, Po-Wei and Lin, Chih-Jen. Iteration complexity of feasible descent methods for convex optimization. *Journal of Machine Learning Research*, 15(1):1523–1548, 2014.
- Weinberger, Kilian Q., Blitzer, John, and Saul, Lawrence K. Distance metric learning for large margin nearest neighbor classification. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1473–1480, 2006.
- Xing, E., Ng, A., Jordan, M., and Russell, S. Distance metric learning with application to clustering with side-information. In *Advances in Neural Information Processing Systems (NIPS)*, volume 15, pp. 505–512, 2003.
- Xu, Yi, Yan, Yan, Lin, Qihang, and Yang, Tianbao. Homotopy smoothing for non-smooth problems with lower complexity than  $o(1/\epsilon)$ . In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1208–1216, 2016.
- Yang, Tianbao and Lin, Qihang. Rsg: Beating subgradient method without smoothness and strong convexity. *CoRR*, abs/1512.03107, 2016.
- Yang, Tianbao, Lin, Qihang, and Zhang, Lijun. A richer theory of convex constrained optimization with reduced projections and improved rates (the long version of icml paper). *CoRR*, arXiv:1608.03487, 2017.
- Zhang, Lijun, Yang, Tianbao, Jin, Rong, and He, Xiaofei.  $O(\log t)$  projections for stochastic optimization of smooth and strongly convex functions. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1121–1129, 2013.
- Zhou, Zirui and So, Anthony Man-Cho. A unified approach to error bounds for structured convex optimization problems. *arXiv:1512.03518*, 2015.
- Zhou, Zirui, Zhang, Qi, and So, Anthony Man-Cho. L1p-norm regularization: Error bounds and convergence rate analysis of first-order methods. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 1501–1510, 2015.