Homotopy Smoothing for Non-Smooth Problems with Lower Complexity than $O(1/\epsilon)$

Yi Xu[†] YI-XU@UIOWA.EDU
Yan Yan[‡*] YAN.YAN-3@STUDENT.UTS.EDU.AU
Qihang Lin[♮] QIHANG-LIN@UIOWA.EDU
Tianbao Yang[†✉] TIANBAO-YANG@UIOWA.EDU

[†] *Department of Computer Science, The University of Iowa, Iowa City, IA 52242, USA*
[‡] *QCIS, University of Technology Sydney, NSW 2007, Australia*
[♮] *Department of Management Sciences, The University of Iowa, Iowa City, IA 52242, USA*

Abstract

In this paper, we develop a novel **ho**moto**p**y **s**moothing (HOPS) algorithm for solving a family of non-smooth problems that is composed of a non-smooth term with an explicit max-structure and a smooth term or a simple non-smooth term whose proximal mapping is easy to compute. The best known iteration complexity for solving such non-smooth optimization problems is $O(1/\epsilon)$ without any assumption on the strong convexity. In this work, we will show that the proposed HOPS achieved a lower iteration complexity of $\widetilde{O}(1/\epsilon^{1-\theta})$[1] with $\theta \in (0, 1]$ capturing the local sharpness of the objective function around the optimal solutions. To the best of our knowledge, this is the lowest iteration complexity achieved so far for the considered non-smooth optimization problems without strong convexity assumption. The HOPS algorithm employs Nesterov's smoothing technique and Nesterov's accelerated gradient method and runs in stages, which gradually decreases the smoothing parameter in a stage-wise manner until it yields a sufficiently good approximation of the original function. We show that HOPS enjoys a linear convergence for many well-known non-smooth problems (e.g., empirical risk minimization with a piece-wise linear loss function and ℓ_1 norm regularizer, finding a point in a polyhedron, cone programming, etc). Experimental results verify the effectiveness of HOPS in comparison with Nesterov's smoothing algorithm and the primal-dual style of first-order methods.

1. Introduction

In this paper, we consider the following optimization problem:

$$\min_{x\in\Omega_1} F(x) \triangleq f(x) + g(x) \tag{1}$$

where $g(x)$ is a convex (but not necessarily smooth) function, Ω_1 is a closed convex set and $f(x)$ is a convex but non-smooth function which can be explicitly written as

$$f(x) = \max_{u\in\Omega_2} \langle Ax, u\rangle - \phi(u) \tag{2}$$

* The work of Y. Yan was done when he was a visiting student with T. Yang at Department of Computer Science of the University of Iowa.

1. $\widetilde{O}()$ suppresses a logarithmic factor.

where $\Omega_2 \subset \mathbb{R}^m$ is a closed convex bounded set, $A \in \mathbb{R}^{m\times d}$, $\phi(u)$ is a convex function, and $\langle \cdot, \cdot \rangle$ is a scalar product. This family of non-smooth optimization problems have applications in numerous domains, e.g., machine learning and statistics (Chen et al., 2012), image processing (Chambolle and Pock, 2011), SDP programming (Nesterov, 2007), cone programming (Lan et al., 2011), and etc. Several first-order methods have been developed for solving such non-smooth optimization problems including the primal-dual methods (Nemirovski, 2005; Chambolle and Pock, 2011), Nesterov's smoothing algorithm (Nesterov, 2005a,b) [2], and they can achieve $O(1/\epsilon)$ iteration complexity for finding an ϵ-optimal solution, which is faster than the corresponding black-box lower complexity bounds by an order of magnitude.

In this paper, we propose a novel homotopy smoothing (HOPS) algorithm for solving the problem in (1) that achieves a lower iteration complexity than $O(1/\epsilon)$. In particular, the iteration complexity of HOPS is given by $\widetilde{O}(1/\epsilon^{1-\theta})$, where $\theta \in (0, 1]$ captures the local sharpness (defined shortly) of the objective function around the optimal solutions. The proposed HOPS algorithm builds on the Nesterov's smoothing technique, i.e., approximating the non-smooth function $f(x)$ by a smooth function and optimizing the smoothed function to a desired accuracy level.

The striking difference between HOPS and Nesterov's smoothing algorithm is that Nesterov uses a fixed small smoothing parameter that renders a sufficiently accurate approximation of the non-smooth function $f(x)$, while HOPS adopts a homotopy strategy for setting the value of the smoothing parameter. It starts from a relatively large smoothing parameter and gradually decreases the smoothing parameter in a stage-wise manner until the smoothing parameter reaches a level that gives a sufficiently good approximation of the non-smooth objective function. The benefit of using a homotopy strategy is that a larger smoothing parameter yields a smaller smoothness constant and hence a lower iteration complexity for smoothed problems in earlier stages. For smoothed problems in later stages with larger smoothness constants, warm-start can help reduce the number of iterations to converge. As a result, solving a series of smoothed approximations with a smoothing parameter from large to small and with warm-start is faster than solving one smoothed approximation with a very small smoothing parameter. To the best of our knowledge, **this is the first work** that rigorously analyzes such a homotopy smoothing algorithm and establishes its theoretical guarantee on lower iteration complexities. The keys to our analysis of lower iteration complexity are (i) to leverage a **global error inequality** (Lemma 1) (Yang and Lin, 2016) that bounds the distance of a solution to the ϵ sublevel set by a multiple of the functional distance; and (ii) to explore a local error bound condition to bound the multiplicative factor.

2. Related Work

In this section, we review some related work for solving the considered family of non-smooth optimization problems. Traditional first-order methods such as subgradient descent for solving non-smooth optimization suffer from an $O(1/\epsilon^2)$ iteration complexity. Below, we review some related work for solving (1) or its special cases with improved iteration

2. The algorithm in (Nesterov, 2005a) was developed for handling a smooth component $g(x)$, which can be extended to handling a non-smooth component $g(x)$ whose proximal mapping is easy to compute.

complexities. There are two categories of algorithms, one is based on Neterov's smoothing technique and another one is primal-dual style of first-order methods.

In the seminal paper by Nesterov (2005a), he proposed a smoothing technique for a family of structured non-smooth optimization problems as in (1) with $g(x)$ being a smooth function and $f(x)$ given in (2). By adding a strongly convex prox function in terms of u with a smoothing parameter μ into the definition of $f(x)$, one can obtain a smoothed approximation of the original objective function. Then he developed an accelerated gradient method with an $O(1/t^2)$ convergence rate for the smoothed objective function with t being the number of iterations, which implies an $O(1/t)$ convergence rate for the original objective function by setting $\mu \approx c/t$ with c being a constant. Although he only considered a smooth component $g(x)$ in the original paper, the algorithm and theory can be easily generalized to handle a non-smooth component $g(x)$ assuming its proximal mapping is simple to compute by using accelerated proximal gradient methods for composite optimization problems (Nesterov, 2013; Beck and Teboulle, 2009). Later on, Nesterov proposed an excessive gap technique for solving the similar problem (Nesterov, 2005b), which avoids setting the value of the smoothing parameter with the number of iterations or the accuracy given in advance. In (Nesterov, 2005b), Nesterov treats $g(x)$ and $\phi(u)$ symmetrically and simultaneously minimizes the smoothed lower bound and maximizes the smoothed upper bound by updating the primal and dual variables and iteratively reducing the smoothing parameters. He established an $O(1/t)$ convergence rate for linear functions $g(x)$ and $\phi(u)$. He also analyzed the case when $g(x)$ is strongly convex, which gives an improved convergence rate of $O(1/t^2)$. The smoothing technique has been exploited to solving problems in machine learning (Orabona et al., 2012; Ouyang and Gray, 2012; Lin et al., 2014) and statistics (Chen et al., 2012), and cone programming (Lan et al., 2011; Nesterov, 2007).

The primal-dual style of first-order methods treat the problem as a convex-concave minimization problem, i.e.,

$$\min_{x\in\Omega_1} \max_{u\in\Omega_2} g(x) + \langle Ax, u\rangle - \phi(u)$$

Nemirovski (2005) proposed a mirror prox method, which has a convergence rate of $O(1/t)$ by assuming that both $g(x)$ and $\phi(u)$ are smooth functions. Chambolle and Pock (2011) designed first-order primal-dual algorithms, which tackle $g(x)$ and $\phi(u)$ using proximal mapping and achieve the same convergence rate of $O(1/t)$ without assuming smoothness of $g(x)$ and $\phi(u)$. When $g(x)$ or $\phi(u)$ is strongly convex, their algorithms achieve $O(1/t^2)$ convergence rate. The effectiveness of their algorithms was demonstrated on imaging problems. Recently, the primal-dual style of first-order methods have been employed to solve non-smooth optimization problems in machine learning where both the loss function and the regularizer are non-smooth (Yang et al., 2014). Lan et al. (2011) also considered Nemirovski's prox method for solving cone programming problems.

The key condition for us to develop an improved convergence is closely related to local error bounds (LEB) (Pang, 1997) and more generally the Kurdyka-Łojasiewicz property (Łojasiewicz, 1965; Bolte et al., 2006). The LEB characterizes the relationship between the distance of a local solution to the optimal set and the optimality gap of the solution in terms of objective value. The Kurdyka-Łojasiewicz property characterizes that property of a function that whether it can be made "sharp" by some transformation. Recently, these conditions/properties have been explored for feasible descent methods (Luo

and Tseng, 1993), non-smooth optimization (Gilpin et al., 2012), gradient and subgradient methods (Jerome Bolte, 2015; Yang and Lin, 2016). **It is notable** that our local error bound condition is different from the one used in (Luo and Tseng, 1993; Zhou and So, 2015) which bounds the distance of a point to the optimal set by the norm of the projected or proximal gradient at that point instead of the functional distance, consequentially it requires some smoothness assumption about the objective function. By contrast, the local error bound condition in this paper covers a much broad family of functions and thus it is more general. Recent work (Necoara et al., 2015; Zhang, 2016) have shown that the error bound in (Luo and Tseng, 1993; Zhou and So, 2015) is a special case of our considered error bound with $\theta = 1/2$. Two mostly related work leveraging a similar error bound to ours are discussed in order. Gilpin et al. (2012) considered the two-person zero-sum games, which is a special case of (1) with $g(x)$ and $\phi(u)$ being zeros and Ω_1 and Ω_2 being polytopes. The present work is a non-trivial generalization of their work that leads to improved convergence for a much broader family of non-smooth optimization problems. In particular, their result is just a special case of our result when the constant θ that captures the local sharpness is one for problems whose epigraph is a polytope. Recently, Yang and Lin (2016) proposed a restarted subgradient method by exploring the local error bound condition or more generally the Kurdyka-Łojasiewicz property, resulting in an $\widetilde{O}(1/\epsilon^{2(1-\theta)})$ iteration complexity with the same constant of θ. In contrast, our result is an improved iteration complexity of $\widetilde{O}(1/\epsilon^{1-\theta})$.

It is worth emphasizing that the proposed homotopy smoothing technique is different from recently proposed homotopy methods for sparse learning (e.g., ℓ_1 regularized least-squares problem (Xiao and Zhang, 2013)), though a homotopy strategy on an involved parameter is also employed to boost the convergence. In particular, the involved parameter in the homotopy methods for sparse learning is the regularization parameter before the ℓ_1 regularization, while the parameter in the present work is the introduced smoothing parameter. In addition, the benefit of starting from a relatively large regularization parameter in sparse learning is the sparsity of the solution, which makes it possible to explore the restricted strong convexity for proving faster convergence. We do not make such assumption of the data and we are mostly interested in that when both $f(x)$ and $g(x)$ are non-smooth.

Lastly, we discuss several closely related recent work that also employ homotopy strategies on the smoothing parameter, but are different in how to decrease the smoothing parameter and in the iteration complexities from the proposed HOPS. Tran-Dinh (2015) proposed an adaptive smoothing algorithm by combining Nesterov's accelerated proximal gradient method and a homotopy strategy for smoothing parameter. Different from Nesterov's smoothing, their algorithm and analysis require that the added prox function is not only μ-strongly convex but also smooth with a smoothness parameter $\beta \geq \mu$. In addition, the smoothing parameter is decreased iteratively in the order of $O(1/t)$ where t is the iteration number. In terms of iteration complexity, when $\beta > \mu$, their method has an iteration complexity of $\widetilde{O}(1/\epsilon)$, and when $\beta = \mu$ it achieves the same iteration complexity of $O(1/\epsilon)$ as Nesterov's smoothing method (Nesterov, 2005a). In contrast, the proposed HOPS employs a different homotopy strategy that decreases the smoothing parameter geometrically in a stage-wise manner and has a better iteration complexity.

In (Freund and Lu, 2015), the authors introduced a new smooth approximation algorithm by leveraging the strict lower bound of the objective function and a function growth

condition. The function growth condition is an inequality that the distance of a point to the optimal solution set is less than a growth constant multiple of the difference between the objective value at the point and the strict lower bound. Their algorithm also has two loops where the outer loop decreases the smoothing parameter according to the difference between the objective value at current solution and the strict lower bound and the inner loop exploits Nesterov's accelerated gradient method to solve the intermediate smoothed problem until the relative improvement in terms of the strict lower bound is above 0.8. In terms of the iteration complexity, their algorithm has an $O\left(\frac{\log H}{\sqrt{\epsilon'}} + \frac{1}{\epsilon'}\right)$ iteration complexity for obtaining an ϵ'-relative optimal solution, i.e., $F(x_t) - F_* \leq \epsilon'(F_* - F_{slb})$ [3], where $H = \frac{F(x_0)-F_*}{F_*-F_{slb}}$. As discussed in (Freund and Lu, 2015), their algorithm is favorable when the distance of the initial solution to the optimal set is sufficiently large as their iteration complexity has a logarithmic dependence on the initial solution. However, their algorithm still suffers $O(1/\epsilon)$ iteration complexity in the worst case. In contrast, HOPS leverages the LEB condition instead of the function growth condition so that it not only enjoys a logarithmic dependence on the initial solution but also a reduced iteration complexity.

More recently, Allen-Zhu and Hazan (2016) proposed black-box reduction methods for convex optimization by reducing the objective function to a β-smooth and μ-strongly convex function and employing an algorithm which satisfies the homotopy objective decrease (HOOD) property (defined shortly). For a non-smooth and non-strongly convex objective function, they propose to add a μ-strongly convex regularization term $\frac{\mu}{2}\|x - x_0\|^2$, where x_0 is a starting point, and smooth the non-smooth function in the finite sum form $\sum_i f_i(x)$ by applying Nesterov's smoothing technique to its Fenchel conjugate, which results in a smooth and strongly convex function. By assuming each $f_i(x)$ is G-Lipschitz continuous and $\|x_0 - x_*\| \leq D$, their method enjoys an $O(GD/\epsilon)$ iteration complexity by employing accelerated gradient descent method for solving the resulted smooth and strongly convex optimization problems. Their reduction method uses a similar homotopy strategy on the smoothing parameter as HOPS. Nonetheless, we emphasize that HOPS is fundamentally different from their method, in particular the inner loop in their method is to ensure the HOOD property of the black-box algorithm for minimizing the intermediate smooth and strongly convex function $f(x)$, i.e., for any starting point x_0, it produces an output x' satisfying $f(x') - \min f(x) \leq \frac{f(x_0)-\min f(x)}{4}$; in contrast, the inner loop of HOPS is to solve the intermediate smoothed problem to an accuracy that matches the oder of the current smoothing parameter. By leveraging the local error bound condition, we are able to achieve an iteration complexity with a better dependence on the accuracy level ϵ and the distance of the initial solution to the optimal set.

Finally, we note that a similar homotopy strategy is employed in Nesterov's smoothing algorithm for solving an ℓ_1 norm minimization problem subject to a constraint for recovering a sparse solution (Becker et al., 2011). However, we would like to draw readers' attention to that they did not provide any theoretical guarantee on the iteration complexity of the homotopy strategy and consequentially their implementation is ad-hoc without guidance from theory. More importantly, our developed algorithms and theory apply to a much broader family of problems.

3. F_* is the optimal objective value and F_{slb} is its strict lower bound

3. Preliminaries

We present some preliminaries in this section. Let $\|x\|$ denote the Euclidean norm on the primal variable x. A function $h(x)$ is L-smooth in terms of $\|\cdot\|$, if

$$\|\nabla h(x) - \nabla h(y)\| \leq L\|x - y\|$$

We remark here that the generalization to a smoothness definition with respect to a p-norm $\|\cdot\|_p$ with $p \in (1, 2]$ is mostly straightforward. We defer the discussion to the supplement. Let $\|u\|_+$ denote a norm on the dual variable, which is not necessarily the Euclidean norm. Denote by $\omega_+(u)$ a 1-strongly convex function of u in terms of $\|\cdot\|_+$.

For the optimization problem in (1), we let Ω_*, F_* denote the set of optimal solutions and optimal value, respectively, and make the following assumption throughout the paper.

Assumption 1 *For a convex minimization problem (1), we assume (i) there exist $x_0 \in \Omega_1$ and $\epsilon_0 \geq 0$ such that $F(x_0) - \min_{x\in\Omega_1} F(x) \leq \epsilon_0$; (ii) $f(x)$ is characterized as in (2), where $\phi(u)$ is a convex function; (iii) There exists a constant D such that $\max_{u\in\Omega_2} \omega_+(u) \leq D^2/2$; (iv) Ω_* is a non-empty convex compact set.*

Note that: 1) Assumption 1(i) assumes that the objective function is lower bounded; 2) Assumption 1(iii) assumes that Ω_2 is a bounded set, which is also required in (Nesterov, 2005a).

In addition, for brevity we assume that $g(x)$ is simple enough [4] such that the proximal mapping defined below is easy to compute similar to (Chambolle and Pock, 2011):

$$P_{\lambda g}(x) = \min_{z\in\Omega_1} \frac{1}{2}\|z - x\|^2 + \lambda g(z) \tag{3}$$

Relying on the proximal mapping, the key updates in the optimization algorithms presented below take the following form:

$$\Pi^c_{v,\lambda g}(x) = \arg\min_{z\in\Omega_1} \frac{c}{2}\|z - x\|^2 + \langle v, z\rangle + \lambda g(z) \tag{4}$$

For any $x \in \Omega_1$, let x^* denote the closest optimal solution in Ω_* to x measured in terms of norm $\|\cdot\|$, i.e., $x^* = \arg\min_{z\in\Omega_*} \|z - x\|^2$, which is unique because Ω_* is a non-empty convex compact set (Hou et al., 2013). We denote by $\mathcal{L}_\epsilon$ the ϵ-level set of $F(x)$ and by $\mathcal{S}_\epsilon$ the ϵ-sublevel set of $F(x)$, respectively, i.e.,

$$\mathcal{L}_\epsilon = \{x \in \Omega_1 : F(x) = F_* + \epsilon\}, \quad \mathcal{S}_\epsilon = \{x \in \Omega_1 : F(x) \leq F_* + \epsilon\}$$

It follows from (Rockafellar, 1970, Corollary 8.7.1) that the sublevel set $\mathcal{S}_\epsilon$ is bounded for any $\epsilon \geq 0$ and so as the level set $\mathcal{L}_\epsilon$ due to that Ω_* is bounded. Define $dist(\mathcal{L}_\epsilon, \Omega_*)$ to be the maximum distance of points on the level set $\mathcal{L}_\epsilon$ to the optimal set Ω_*, i.e.,

$$dist(\mathcal{L}_\epsilon, \Omega_*) = \max_{x\in\mathcal{L}_\epsilon} \left[dist(x, \Omega_*) \triangleq \min_{z\in\Omega_*} \|x - z\| \right]. \tag{5}$$

4. If $g(x)$ is smooth, this assumption can be relaxed. We will defer the discussion and result on a smooth function $g(x)$ to Section 4.4.

Due to that $\mathcal{L}_\epsilon$ and Ω_* are bounded, $dist(\mathcal{L}_\epsilon, \Omega_*)$ is also bounded. Let $x_\epsilon^\dagger$ denote the closest point in the ϵ-sublevel set to x, i.e.,

$$x_\epsilon^\dagger = \arg\min_{z \in \mathcal{S}_\epsilon} \|z - x\|^2 \tag{6}$$

It is easy to show that $x_\epsilon^\dagger \in \mathcal{L}_\epsilon$ when $x \notin \mathcal{S}_\epsilon$ (using the KKT condition).

4. Homotopy Smoothing

In this section, we first describe Nesterov's smoothing technique, and then present the HOPS algorithm and its convergence analysis. Next, we will discuss the local error bound and its application. Finally, we will discuss the HOPS when g is smooth, and also extend the HOPS to general p-norm.

4.1. Nesterov's Smoothing

We first present the Nesterov's smoothing technique and accelerated proximal gradient methods for solving the smoothed problem due to that the proposed algorithm builds upon these techniques. The idea of smoothing is to construct a smooth function $f_\mu(x)$ that well approximates $f(x)$. Nesterov considered the following function

$$f_\mu(x) = \max_{u \in \Omega_2} \langle Ax, u\rangle - \phi(u) - \mu\omega_+(u)$$

It was shown in (Nesterov, 2005a) that $f_\mu(x)$ is smooth w.r.t $\|\cdot\|$ and its smoothness parameter is given by $L_\mu = \frac{1}{\mu}\|A\|^2$ where $\|A\|$ is defined by $\|A\| = \max_{\|x\|\leq 1} \max_{\|u\|_+ \leq 1} \langle Ax, u\rangle$. Denote by

$$u_\mu(x) = \arg\max_{u \in \Omega_2} \langle Ax, u\rangle - \phi(u) - \mu\omega_+(u)$$

The gradient of $f_\mu(x)$ is computed by $\nabla f_\mu(x) = A^\top u_\mu(x)$. Then

$$f_\mu(x) \leq f(x) \leq f_\mu(x) + \mu D^2/2 \tag{7}$$

From the inequality above, we can see that when μ is very small, $f_\mu(x)$ gives a good approximation of $f(x)$. This motivates us to solve the following composite optimization problem

$$\min_{x \in \Omega_1} F_\mu(x) \triangleq f_\mu(x) + g(x)$$

Many works have studied such an optimization problem (Beck and Teboulle, 2009; Tseng, 2008) and the best convergence rate is given by $O(L_\mu/t^2)$, where t is the total number of iterations. We present a variant of accelerated proximal gradient (APG) methods in Algorithm 1 that works even with $\|x\|$ replaced with a general norm as long as its square is strongly convex. We make several remarks about Algorithm 1: (i) the variant here is similar to Algorithm 3 in (Tseng, 2008) and the algorithm proposed in (Nesterov, 2005a) except that the prox function $d(x)$ is replaced by $\|x - x_0\|^2/2$ in updating the sequence of z_k, which is assumed to be σ_1-strongly convex w.r.t $\|\cdot\|$; (ii) If $\|\cdot\|$ is simply the Euclidean

Algorithm 1 An Accelerated Proximal Gradient Method: APG(x_0, t, L_μ)

Input: the number of iterations t, the initial solution x_0, and the smoothness constant L_μ
Let $\theta_0 = 1$, $V_{-1} = 0$, $\Gamma_{-1} = 0$, $z_0 = x_0$
Let α_k and θ_k be two sequences given in Theorem 2.
for $k = 0, \ldots, t-1$ **do**
  Compute $y_k = (1-\theta_k)x_k + \theta_k z_k$
  Compute $v_k = \nabla f_\mu(y_k)$, $V_k = V_{k-1} + \frac{v_k}{\alpha_k}$, and $\Gamma_k = \Gamma_{k-1} + \frac{1}{\alpha_k}$
  Compute $z_{k+1} = \Pi^{L_\mu/\sigma_1}_{V_k, \Gamma_k g}(x_0)$ and $x_{k+1} = \Pi^{L_\mu}_{v_k, g}(y_k)$
end for
Output: x_t

norm, a simplified algorithm with only one update in (4) can be used (e.g., FISTA (Beck and Teboulle, 2009)); (iii) if L_μ is difficult to compute, we can use the backtracking trick (see (Beck and Teboulle, 2009; Tseng, 2008)).

The following theorem states the convergence result for APG.

Theorem 2 *(Nesterov, 2005a; Tseng, 2008) Let* $\theta_k = \frac{2}{k+2}$, $\alpha_k = \frac{2}{k+1}, k \geq 0$ *or* $\alpha_{k+1} = \theta_{k+1} = \frac{\sqrt{\theta_k^4 + 4\theta_k^2} - \theta_k^2}{2}, k \geq 0$. *For any* $x \in \Omega_1$, *we have*

$$F_\mu(x_t) - F_\mu(x) \leq \frac{2L_\mu \|x - x_0\|^2}{t^2} \tag{8}$$

Combining the above convergence result with the relation in (7), we can establish the iteration complexity of Nesterov's smoothing algorithm for solving the original problem (1).

Corollary 3 *(Nesterov, 2005a) For any* $x \in \Omega_1$, *we have*

$$F(x_t) - F(x) \leq \mu D^2/2 + \frac{2L_\mu \|x - x_0\|^2}{t^2} \tag{9}$$

In particular in order to have $F(x_t) \leq F_* + \epsilon$, *it suffices to set* $\mu \leq \frac{\epsilon}{D^2}$ *and* $t \geq \frac{2D\|A\|\|x_0 - x_*\|}{\epsilon}$, *where* x_* *is an optimal solution to (1).*

4.2. Homotopy Smoothing

From the convergence result in (9), we can see that in order to obtain a very accurate solution, we have to set μ - the smoothing parameter - to be a very small value, which will cause the blow-up of the second term because $L_\mu \propto 1/\mu$. On the other hand, if μ is set to be a relatively large value, then t can be set to be a relatively small value to match the first term, which may lead to a not sufficiently accurate solution. It seems that the $O(1/\epsilon)$ is unbeatable. However, if we adopt a homotopy strategy, i.e., starting from a relatively large value μ and optimizing the smoothed function with a certain number of iterations t such that the second term in (9) matches the first term, which will give $F(x_t) - F(x_*) \leq O(\mu)$. Then we can reduce the value of μ by a constant factor $b > 1$ and

warm-start the optimization process again from x_t. The key observation is that although μ decreases and L_μ increases, the other term $\|x_* - x_t\|$ is also reduced compared to $\|x_* - x_0\|$, which could cancel the blow-up effect caused by increased L_μ. As a result, we expect to use the same number of iterations to optimize the smoothed function with a smaller μ such that $F(x_{2t}) - F(x_*) \leq O(\mu/b)$.

To formalize our observation, we first present the following key lemma below.

Lemma 1 (Yang and Lin (2016)) *For any $x \in \Omega_1$ and $\epsilon > 0$, we have*

$$\|x - x_\epsilon^\dagger\| \leq \frac{dist(x_\epsilon^\dagger, \Omega_*)}{\epsilon}(F(x) - F(x_\epsilon^\dagger))$$

where $x_\epsilon^\dagger \in \mathcal{S}_\epsilon$ is the closest point in the ϵ-sublevel set to x as defined in (6).

The lemma is proved in (Yang and Lin, 2016). We include its proof in Appendix. If we apply the above bound into (9), we will see in the proof of the main theorem (Theorem 4) that the number of iterations t for solving each smoothed problem is roughly $O(\frac{dist(\mathcal{L}_\epsilon, \Omega_*)}{\epsilon})$, which will be lower than $O(\frac{1}{\epsilon})$ in light of the local error bound condition given below.

Definition 2 (Local error bound (LEB)) *A function $F(x)$ is said to satisfy a local error bound condition if there exist $\theta \in (0, 1]$ and $c > 0$ such that for any $x \in \mathcal{S}_\epsilon$*

$$dist(x, \Omega_*) \leq c(F(x) - F_*)^\theta \tag{10}$$

Remark: In next subsection, we will discuss the relationship with other types of conditions and show that a broad family of non-smooth functions (including almost all commonly seen functions in machine learning) obey the local error bound condition. The exponent constant θ can be considered as a local sharpness measure of the function. Figure 1 illustrates the sharpness of $F(x) = |x|^p$ for $p = 1, 1.5$, and 2 around the optimal solutions and their corresponding θ.

With the local error bound condition, we can see that $dist(\mathcal{L}_\epsilon, \Omega_*) \leq c\epsilon^\theta, \theta \in (0, 1]$. Now, we are ready to present the homotopy smoothing algorithm and its convergence guarantee under the local error bound condition. The HOPS algorithm is presented in Algorithm 2, which starts from a relatively large smoothing parameter $\mu = \mu_1$ and gradually reduces μ by a factor of $b > 1$ after running a number t of iterations of APG with warm-start. The iteration complexity of HOPS is established below.

Theorem 4 *Suppose Assumption 1 holds and $F(x)$ obeys the local error bound condition. Let HOPS run with $t = O(\frac{2bcD\|A\|}{\epsilon^{1-\theta}}) \geq \frac{2bcD\|A\|}{\epsilon^{1-\theta}}$ iterations for each stage, and $m = \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$. Then*

$$F(x_m) - F_* \leq 2\epsilon.$$

Hence, the iteration complexity for achieving an 2ϵ-optimal solution is $\frac{2bcD\|A\|}{\epsilon^{1-\theta}} \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$ in the worst-case.

Algorithm 2 Homotopy Smoothing (HOPS) for solving (1)

1: **Input**: the number of stages m and the number of iterations t per-stage, and the initial solution $x_0 \in \Omega_1$ and a parameter $b > 1$.
2: Let $\mu_1 = \epsilon_0/(bD^2)$
3: **for** $s = 1, \ldots, m$ **do**
4: Let $x_s = \text{APG}(x_{s-1}, t, L_{\mu_s})$
5: Update $\mu_{s+1} = \mu_s/b$
6: **end for**
7: **Output**: x_m

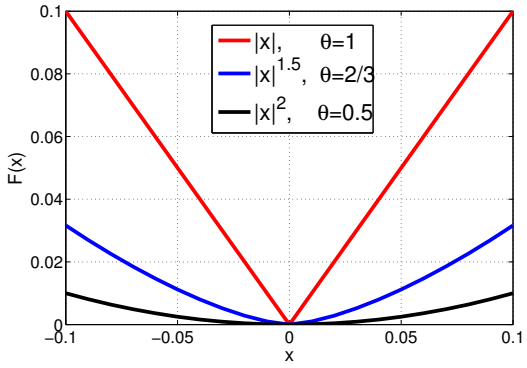

Figure 1: Illustration of local sharpness of three functions and the corresponding θ in the LEB condition.

Proof Let $x^\dagger_{s,\epsilon}$ denote the closest point to x_s in the ϵ sublevel set. Define $\epsilon_s \triangleq \frac{\epsilon_0}{b^s}$. Note that $\mu_s = \epsilon_s/D^2$. We will show by induction that $F(x_s) - F_* \leq \epsilon_s + \epsilon$ for $s = 0, 1, \ldots$ which leads to our conclusion when $s = m$. The inequality holds obviously for $s = 0$. Assuming $F(x_{s-1}) - F_* \leq \epsilon_{s-1} + \epsilon$, we need to show that $F(x_s) - F_* \leq \epsilon_s + \epsilon$. We apply Corollary 3 to the s-th epoch of Algorithm 2 and get

$$F(x_s) - F(x^\dagger_{s-1,\epsilon}) \leq \frac{D^2\mu_s}{2} + \frac{2\|A\|^2\|x_{s-1} - x^\dagger_{s-1,\epsilon}\|^2}{\mu_s t^2} \tag{11}$$

First, we assume $F(x_{s-1}) - F_* \leq \epsilon$, i.e. $x_{s-1} \in \mathcal{S}_\epsilon$. Then we have $x^\dagger_{s-1,\epsilon} = x_{s-1}$ and

$$F(x_s) - F(x^\dagger_{s-1,\epsilon}) \leq \frac{D^2\mu_s}{2} \leq \frac{\epsilon_s}{2}$$

As a result,

$$F(x_s) - F_* \leq F(x^\dagger_{s-1,\epsilon}) - F_* + \frac{\epsilon_s}{2} \leq \epsilon + \epsilon_s$$

Next, we consider $F(x_{s-1}) - F_* > \epsilon$, i.e. $x_{s-1} \notin \S_\epsilon$. Then we have $F(x^\dagger_{s-1,\epsilon}) - F_* = \epsilon$. By Lemma 1, we have

$$\begin{aligned}\|x_{s-1} - x^\dagger_{s-1,\epsilon}\| &\leq \frac{dist(x^\dagger_{s-1,\epsilon}, \Omega_*)}{\epsilon}(F(x_{s-1}) - F(x^\dagger_{s-1,\epsilon})) \\ &\leq \frac{dist(x^\dagger_{s-1,\epsilon}, \Omega_*)}{\epsilon}[\epsilon_{s-1} + \epsilon - \epsilon] = \frac{dist(x^\dagger_{s-1,\epsilon}, \Omega_*)\epsilon_{s-1}}{\epsilon} \\ &\leq \frac{c(F(x^\dagger_{s-1,\epsilon}) - F_*)^\theta \epsilon_{s-1}}{\epsilon} \\ &\leq \frac{c(\epsilon)^\theta \epsilon_{s-1}}{\epsilon} = \frac{c\epsilon_{s-1}}{\epsilon^{1-\theta}} \end{aligned} \tag{12}$$

Combining (11) and (12) and using the fact that $\mu_s = \dfrac{\epsilon_s}{D^2}$ and $t \geq \dfrac{2bcD\|A\|}{\epsilon^{1-\theta}}$, we have

$$F(x_s) - F(x^\dagger_{s-1,\epsilon}) \leq \frac{\epsilon_s}{2} + \frac{\epsilon^2_{s-1}}{2\epsilon_s b^2} = \epsilon_s$$

which together with the fact that $F(x^\dagger_{s-1,\epsilon}) = F_* + \epsilon$ implies

$$F(x_s) - F_* \leq \epsilon + \epsilon_s$$

Therefore by induction, we have

$$F(x_m) - F_* \leq \epsilon_m + \epsilon = \frac{\epsilon_0}{b^m} + \epsilon \leq 2\epsilon$$

where the last inequality is due to the value of m.

■

4.3. Local Error Bounds and Applications

In this subsection, we discuss the local error bound condition and its application in non-smooth optimization problems.

The Hoffman's bound and finding a point in a polyhedron. A polyhedron can be expressed as $\mathcal{P} = \{x \in \mathbb{R}^d; B_1 x \leq b_1, B_2 x = b_2\}$. The Hoffman's bound Pang (1997) is expressed as

$$dist(x, \mathcal{P}) \leq c(\|(B_1 x - b_1)_+\| + \|B_2 x - b_2\|), \exists c > 0 \tag{13}$$

where $[s]_+ = \max(0, s)$. This can be considered as the error bound for the polyhedron feasibility problem, i.e., finding a $x \in \mathcal{P}$, which is equivalent to

$$\min_{x \in \mathbb{R}^d} F(x) \triangleq \left[\|(B_1 x - b_1)_+\| + \|B_2 x - b_2\| = \max_{u \in \Omega_2} \langle B_1 x - b_1, u_1\rangle + \langle B_2 x - b_2, u_2\rangle\right]$$

where $u = (u_1^\top, u_2^\top)^\top$ and $\Omega_2 = \{u | u_1 \succeq 0, \|u_1\| \leq 1, \|u_2\| \leq 1\}$. If there exists a $x \in \mathcal{P}$, then $F_* = 0$. Thus the Hoffman's bound in (13) implies a local error bound (10) with $\theta = 1$. Therefore, the HOPS has a linear convergence for finding a feasible solution in a polyhedron. If we let $\omega_+(u) = \frac{1}{2}\|u\|^2$ then $D^2 = 2$ so that the iteration complexity is $2\sqrt{2}bc \max(\|B_1\|, \|B_2\|)\lceil \log_b(\frac{\epsilon_0}{\epsilon})\rceil$.

Cone programming. Let U, V denote two vector spaces. Given a linear opearator $\mathcal{E}: U \to V^*$ [4], a closed convex set $\Omega \subseteq U$, and a vector $e \in V^*$, and a closed convex cone $\mathcal{K} \subseteq V$, the general constrained cone linear system (cone programing) consists of finding a vector $x \in \Omega$ such that $\mathcal{E}x - e \in \mathcal{K}^*$. Lan et al. Lan et al. (2011) have considered Nesterov's smoothing algorithm for solving the cone programming problem with $O(1/\epsilon)$ iteration complexity. The problem can be cast into a non-smooth optimization problem:

$$\min_{x\in\Omega} F(x) \triangleq \left[dist(\mathcal{E}x - e, \mathcal{K}^*) = \max_{\|u\|\leq 1, u\in -\mathcal{K}} \langle \mathcal{E}x - e, u\rangle\right]$$

Assume that $e \in Range(\mathcal{E}) - \mathcal{K}^*$, then $F_* = 0$. Burke et al. Burke and Tseng (1996) have considered the error bound for such problems and their results imply that there exists $c > 0$ such that $dist(x, \Omega_*) \leq c(F(x) - F_*)$ as long as $\exists x \in \Omega, \text{s.t. } \mathcal{E}x - e \in int(\mathcal{K}^*)$, where Ω_* denotes the optimal solution set. Therefore, the HOPS also has a linear convergence for cone programming. Considering that both U and V are Euclidean spaces, we set $\omega_+(u) = \frac{1}{2}\|u\|^2$ then $D^2 = 1$. Thus, the iteraction complexity of HOPS for finding an 2ϵ-solution is $2bc\|\mathcal{E}\|\lceil\log_b(\frac{\epsilon_0}{\epsilon})\rceil$.

Non-smooth regularized empirical loss (REL) minimization in Machine Learning The REL consists of a sum of loss functions on the training data and a regularizer, i.e.,

$$\min_{x\in\mathbb{R}^d} F(x) \triangleq \frac{1}{n}\sum_{i=1}^{n} \ell(x^\top a_i, y_i) + \lambda g(x)$$

where $(a_i, y_i), i = 1, \ldots, n$ denote pairs of a feature vector and a label of training data. Non-smooth loss functions include hinge loss $\ell(z, y) = \max(0, 1 - yz)$, absolute loss $\ell(z, y) = |z - y|$, which can be written as the max structure in (2). Non-smooth regularizers include e.g., $g(x) = \|x\|_1$, $g(x) = \|x\|_\infty$. These loss functions and regularizers are essentially piecewise linear functions, whose epigraph is a polyhedron. The error bound condition has been developed for such kind of problems (Yang and Lin, 2016). In particular, if $F(x)$ has a polyhedral epigraph, then there exists $c > 0$ such that $dist(x, \Omega_*) \leq c(F(x) - F_*)$ for any $x \in \mathbb{R}^d$. It then implies HOPS has an $O(\log(\epsilon_0/\epsilon))$ iteration complexity for solving a non-smooth REL minimization with a polyhedral epigraph. Yang et al. (2014) has also considered such non-smooth problems, but they only have $O(1/\epsilon)$ iteration complexity.

When $F(x)$ is essentially locally strongly convex (Goebel and Rockafellar, 2008) in terms of $\|\cdot\|$ such that [5]

$$dist^2(x, \Omega_*) \leq \frac{2}{\sigma}(F(x) - F_*), \forall x \in \mathcal{S}_\epsilon \tag{14}$$

then we can see that the local error bound holds with $\theta = 1/2$, which implies the iteration complexity of HOPS is $\widetilde{O}(\frac{1}{\sqrt{\epsilon}})$, which is up to a logarithmic factor the same as the result in (Chambolle and Pock, 2011) for a strongly convex function. However, here only local

4. V^* represents the dual space of V. The notations and descriptions are adopted from (Lan et al., 2011).
5. This is true if $g(x)$ is strongly convex or locally strongly convex.

strong convexity is sufficient and there is no need to develop a different algorithm and different analysis from the non-strongly convex case as done in (Chambolle and Pock, 2011). For example, one can consider $F(x) = \|Ax - y\|_p^p = \sum_{i=1}^n |a_i^\top x - y_i|^p, p \in (1, 2)$, which satisfies (14) according to (Yang and Lin, 2016).

The Kurdyka-Łojasiewicz (KL) property. The definition of KL property is given below.

Definition 3 *Let's define the subdifferential of F at x as $\partial F(x) = \{u : F(y) \geq F(x) + \langle u, y - x\rangle$ for all $y \in \Omega_1\}$. The function $F(x)$ is said to have the KL property at $x_* \in \Omega_*$ if there exist $\eta \in (0, \infty]$, a neighborhood U of x_* and a continuous concave function $\varphi : [0, \eta) \to \mathbb{R}_+$ such that i) $\varphi(0) = 0$, φ is continuous on $(0, \eta)$, ii) for all $s \in (0, \eta)$, $\varphi'(s) > 0$, iii) and for all $x \in U \cup \{x : F(x_*) < F(x) < F(x_*) + \eta\}$, the KL inequality $\varphi'(F(x) - F(x_*))\|\partial F(x)\| \geq 1$ holds.*

The function φ is called the desingularizing function of F at x_*, which makes the function $F(x)$ sharp by reparameterization. An important desingularizing function is in the form of $\varphi(s) = cs^{1-\beta}$ for some $c > 0$ and $\beta \in [0, 1)$, which gives the KL inequality $\|\partial F(x)\| \geq \frac{1}{c(1-\beta)}(F(x) - F(x_*))^\beta$. It has been established that the KL property is satisfied by a wide class of non-smooth functions definable in an o-minimal structure (Bolte et al., 2006). Semialgebraic functions and (globally) subanalytic functions are for instance definable in their respective classes. While the definition of KL property involves a neighborhood U and a constant η, in practice many convex functions satisfy the above property with $U = \mathbb{R}^d$ and $\eta = \infty$ (Attouch et al., 2010). The proposition below shows that a function with the KL property with a desingularizing function $\varphi(s) = cs^{1-\beta}$ obeys the local error bound condition in (10) with $\theta = 1 - \beta \in (0, 1]$, which implies an iteration complexity of $\widetilde{O}(1/\epsilon^\theta)$ of HOPS for optimizing such a function.

Proposition 1 *(Jerome Bolte, 2015, Theorem 5) Let $F(x)$ be a proper, convex and lower-semicontinuous function that satisfies KL property at x_* and U be a neighborhood of x_*. For all $x \in U \cap \{x : F(x_*) < F(x) < F(x_*) + \eta\}$, if $\|\partial F(x)\| \geq \frac{1}{c(1-\beta)}(F(x) - F(x_*))^\beta$, then $dist(x, \Omega_*) \leq c(F(x) - F_*)^{1-\beta}$.*

Remark: In order to apply the KL property to our method, we usually need to assume the KL property is satisfied at every x_* with U containing $\mathcal{S}_\epsilon$ in Proposition 1, i.e. assume $\mathcal{S}_\epsilon \subset \bigcup_{x_* \in \Omega_*} U_{x_*}$, where U_{x_*} is the neighborhood U of a particular x_*. However, as we mentioned above, in practice many convex functions satisfy the KL property with $U = \mathbb{R}^d$ and $\eta = \infty$ (Attouch et al., 2010) so that above assumption holds.

4.4. HOPS for a smooth $g(x)$

In the preliminaries section, we assume that $g(z)$ is simple enough such that the proximal mapping defined below is easy to compute:

$$P_{\lambda g}(x) = \min_{z \in \Omega_1} \frac{1}{2}\|z - x\|^2 + \lambda g(z) \tag{15}$$

We claimed that if $g(z)$ is smooth, this assumption can be relaxed. In this section, we present the discussion and result for a smooth function $g(x)$ without assuming that its proximal

Algorithm 3 An Accelerated Proximal Gradient Method (g is smooth): APG(x_0, t, L_μ)

1: **Input**: the number of iterations t, the initial solution x_0, and the smoothness constant L_μ
2: Let $\theta_0 = 1$, $U_{-1} = 0$, $z_0 = x_0$
3: Let α_k and θ_k be two sequences given in Theorem 2.
4: **for** $k = 0, \ldots, t-1$ **do**
5: Compute $y_k = (1-\theta_k)x_k + \theta_k z_k$
6: Compute $u_k = \nabla f_\mu(y_k) + \nabla g(y_k)$, $U_k = U_{k-1} + \frac{u_k}{\alpha_k}$
7: Compute $z_{k+1} = \widetilde{\Pi}_{U_k}^{(L_\mu+M)/\sigma_1}(x_0)$ and $x_{k+1} = \widetilde{\Pi}_{u_k}^{L_\mu+M}(y_k)$
8: **end for**
9: **Output**: x_t

mapping is easy to compute. In particular, we will consider g as a smooth component in $f_\mu + g$ and use the gradient of both f_μ and g in the updating. The detailed updates are presented in Algorithm 3, where

$$\widetilde{\Pi}_u^c(x) = \arg\min_{z\in\Omega_1} \langle u, z\rangle + \frac{c}{2}\|z - x\|^2 \tag{16}$$

To present the convergence guarantee, we assume that the function g is M-smooth w.r.t $\|x\|$, then the smoothness parameter of objective function $F_\mu(x) = f_\mu(x) + g(x)$ is

$$L = L_\mu + M = \frac{\|A\|^2}{\mu} + M \tag{17}$$

Then, we state the convergence result of Algorithm 3 in the following corollary.

Corollary 5 *Let* $\theta_k = \frac{2}{k+2}$, $\alpha_k = \frac{2}{k+1}, k \geq 0$ *or* $\alpha_{k+1} = \theta_{k+1} = \frac{\sqrt{\theta_k^4 + 4\theta_k^2} - \theta_k^2}{2}, k \geq 0$. *For any* $x \in \Omega_1$, *we have*

$$F(x_t) - F(x) \leq \frac{\mu D^2}{2} + \frac{2\|A\|^2\|x - x_0\|^2}{\mu t^2} + \frac{2M\|x - x_0\|^2}{t^2} \tag{18}$$

Remark: In order to have $F(x_t) \leq F(x_*) + \epsilon$, we can consider $x = x_*$ in Corollary 5, i.e.

$$F(x_t) - F(x_*) \leq \frac{\mu D^2}{2} + \frac{2\|A\|^2\|x_* - x_0\|^2}{\mu t^2} + \frac{2M\|x_* - x_0\|^2}{t^2} \tag{19}$$

In particular, we set

$$\mu = \frac{2\epsilon}{3D^2}$$

and

$$t \geq \max\left\{\frac{3D\|A\|\|x_* - x_0\|}{\epsilon}, \frac{\sqrt{6M}\|x_* - x_0\|}{\sqrt{\epsilon}}\right\}$$

Algorithm 3 also achieves the iteration complecity of $O(1/\epsilon)$.

Similarly, we can develop the HOPS algorithm and present it in Algorithm 4. The iteration complexity of HOPS is established in Theorem 6.

Algorithm 4 Homotopy Smoothing (HOPS) for solving (1) (g is smooth)

1: **Input**: the number of stages m and the number of iterations t per-stage, and the initial solution $x_0 \in \Omega_1$ and a parameter $b > 1$.
2: Let $\mu_1 = \frac{2\epsilon_0}{3bD^2}$
3: **for** $s = 1, \ldots, m$ **do**
4: Let $x_s = \text{APG}(x_{s-1}, t, L_{\mu_s})$
5: Update $\mu_{s+1} = \mu_s / b$
6: **end for**
7: **Output**: x_m

Theorem 6 *Suppose Assumption 1 holds and $F(x)$ obeys the local error bound condition. Let HOPS run with $t = O(1/\epsilon^{1-\theta}) \geq \max\left\{\frac{3D\|A\|bc}{\epsilon^{1-\theta}}, \frac{\sqrt{6M\epsilon_s}bc}{\epsilon^{1-\theta}}\right\}$ iterations for each stage, and $m = \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$. Then*

$$F(x_m) - F_* \leq 2\epsilon.$$

Hence, the iteration complexity for achieving an 2ϵ-optimal solution is $\widetilde{O}(1/\epsilon^{1-\theta})$.

Proof Let $x^\dagger_{s,\epsilon}$ denote the closest point to x_s in the ϵ sublevel set and define $\epsilon_s \triangleq \frac{\epsilon_0}{b^s}$. We will show by induction that $F(x_s) - F_* \leq \epsilon_s + \epsilon$ for $s = 0, 1, \ldots$ which leads to our conclusion when $s = m$. The inequality holds obviously for $s = 0$. Assuming $F(x_{s-1}) - F_* \leq \epsilon_{s-1} + \epsilon$, we need to show that $F(x_s) - F_* \leq \epsilon_s + \epsilon$. We apply Corollary 5 to the s-th epoch of Algorithm 4 and get

$$F(x_s) - F(x^\dagger_{s-1,\epsilon}) \leq \frac{\mu_s D^2}{2} + \frac{2\|A\|^2\|x^\dagger_{s-1,\epsilon} - x_{s-1}\|^2}{\mu_s t^2} + \frac{2M\|x^\dagger_{s-1,\epsilon} - x_{s-1}\|^2}{t^2} \tag{20}$$

First, we assume $F(x_{s-1}) - F_* \leq \epsilon$, i.e. $x_{s-1} \in \mathcal{S}_\epsilon$. Then we have $x^\dagger_{s-1,\epsilon} = x_{s-1}$ and

$$F(x_s) - F(x^\dagger_{s-1,\epsilon}) \leq \frac{D^2\mu_s}{2} \leq \frac{\epsilon_s}{3}$$

As a result,

$$F(x_s) - F_* \leq F(x^\dagger_{s-1,\epsilon}) - F_* + \frac{\epsilon_s}{3} \leq \epsilon + \epsilon_s$$

Next, we consider $F(x_{s-1}) - F_* > \epsilon$, i.e. $x_{s-1} \notin \S_\epsilon$. Then we have $F(x^\dagger_{s-1,\epsilon}) - F_* = \epsilon$. Recall that

$$\|x_{s-1} - x^\dagger_{s-1,\epsilon}\| \leq \frac{c\epsilon_{s-1}}{\epsilon^{1-\theta}} \tag{21}$$

Combining (20) and (21) and using the fact that $\mu_s = \frac{2\epsilon_s}{3D^2}$ and $t \geq \max\left\{\frac{3D\|A\|bc}{\epsilon^{1-\theta}}, \frac{\sqrt{6M\epsilon_s}bc}{\epsilon^{1-\theta}}\right\}$, we get

$$\begin{aligned} F(x_s) - F(x^\dagger_{s-1,\epsilon}) &\leq \frac{\epsilon_s}{3} + \frac{3D^2\|A\|^2c^2\epsilon^2_{s-1}}{\epsilon_s\epsilon^{2(1-\theta)}t^2} + \frac{2Mc^2\epsilon^2_{s-1}}{\epsilon^{2(1-\theta)}t^2} \\ &\leq \frac{\epsilon_s}{3} + \frac{\epsilon^2_{s-1}}{3\epsilon_s b^2} + \frac{\epsilon^2_{s-1}}{3\epsilon_s b^2} = \epsilon_s \end{aligned}$$

which together with the fact that $F(x^\dagger_{s-1,\epsilon}) = F_* + \epsilon$ implies

$$F(x_s) - F_* \leq \epsilon + \epsilon_s$$

Therefore by induction, we have

$$F(x_m) - F_* \leq \epsilon_m + \epsilon = \frac{\epsilon_0}{b^m} + \epsilon \leq 2\epsilon$$

where the last inequality is due to the value of $m = \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$.

In fact, the number of iteration in each stage depends on s, then the iteration complexity for achieving an 2ϵ-optimal solution is

$$\begin{aligned}
\sum_{s=1}^{m} \max\left\{\frac{3D\|A\|bc}{\epsilon^{1-\theta}}, \frac{\sqrt{6M\epsilon_s}bc}{\epsilon^{1-\theta}}\right\} &\leq \sum_{s=1}^{m} \frac{3D\|A\|bc + \sqrt{6M\epsilon_s}bc}{\epsilon^{1-\theta}} \\
&= \frac{3D\|A\|bc}{\epsilon^{1-\theta}} \left\lceil \log_b\left(\frac{\epsilon_0}{\epsilon}\right)\right\rceil + \sum_{s=1}^{m} \frac{\sqrt{6M\epsilon_0}bc}{\sqrt{b^s}\epsilon^{1-\theta}} \\
&\leq \frac{3D\|A\|bc}{\epsilon^{1-\theta}} \left\lceil \log_b\left(\frac{\epsilon_0}{\epsilon}\right)\right\rceil + \frac{\sqrt{6M\epsilon_0}bc}{(\sqrt{b}-1)\epsilon^{1-\theta}}
\end{aligned}$$

■

4.5. HOPS with a p-norm

As we mentioned in the paper, we can generalize the results to a smoothness definition with respect to a p-norm $\|x\|_p$ with $p \in (1, 2]$, which makes $\frac{1}{2}\|x\|_p^2$ a $(p-1)$-strongly convex function w.r.t $\|\cdot\|_p$. Algorithm 1 and Algorithm 3 remain the same with $\sigma_1 = p - 1$ except that the norm $\|\cdot\|$ is replaced with $\|\cdot\|_p$ in the updates of x_k and z_k. In order to have efficient updates using a p-norm, we assume $\Omega_1 = \mathbb{R}^d$. Similarly as before, we introduce several notations. Let $dist_p(x, \Omega_*) = \min_{z\in\Omega_*} \|x - z\|_p$. Let $x^\dagger_\epsilon$ denote the closest point in the ϵ-sublevel set to x measured in p-norm, i.e,

$$x^\dagger_\epsilon = \arg\min_{z\in\mathbb{R}^d} \|z - x\|_p^2, \quad s.t. \quad F(z) \leq F_* + \epsilon$$

The following lemma is a generalization of Lemma 1 to a p-norm.

Lemma 4 ((Yang and Lin, 2016)) *For any $x \in \mathbb{R}^d$ and $\epsilon > 0$, we have*

$$\|x - x^\dagger_\epsilon\|_p \leq \frac{dist_p(x^\dagger_\epsilon, \Omega_*)}{\epsilon}(F(x) - F(x^\dagger_\epsilon))$$

where $x^\dagger_\epsilon \in \mathcal{S}_\epsilon$ is the closest point in the ϵ-sublevel set to x.

The proof of the above lemma can be also found in (Yang and Lin, 2016). For completeness, we give the proof in Appendix.

To establish the improved convergence, we assume the following local error bound condition using the p-norm.

Definition 5 (Local error bound) *A function $F(x)$ is said to satisfy a local error bound condition w.r.t a p-norm if there exist $\theta \in (0, 1]$ and $c > 0$ such that for any $x \in \mathcal{S}_\epsilon$*

$$dist_p(x, \Omega_*) \leq c(F(x) - F_*)^\theta \tag{22}$$

The convergence of APG with a p-norm is similar to Corollary 3 in the paper.

Corollary 7 *For any $x \in \mathbb{R}^d$, by running APG with a p-norm, we have*

$$F(x_t) - F(x) \leq \mu D^2/2 + \frac{2L_\mu \|x - x_0\|_p^2}{t^2} \tag{23}$$

Finally, we have the similar convergence as Theorem 4 for HOPS except that $\|A\|$ is defined using the p-norm of x.

5. Primal-Dual Homotopy Smoothing (PD-HOPS)

We note that the required number of iterations per-stage t for finding an ϵ accurate solution depends on unknown constant c and sometimes θ. Thus, an inappropriate setting of t may lead to a less accurate solution. To address this issue, we present a primal-dual homotopy smoothing. Basically, we also apply the homotopy smoothing to the dual problem:

$$\max_{u \in \Omega_2} \Phi(u) \triangleq -\phi(u) + \underbrace{\min_{x \in \Omega_1} \langle A^\top u, x\rangle + g(x)}_{\psi(u)} \tag{24}$$

Denote by Φ_* the optimal value of the above problem. Under some mild conditions, it is easy to see that $\Phi_* = F_*$. By extending the analysis and result to the dual problem, we can obtain that $F(x_m) - \Phi(u_m) \leq 4\epsilon$. Thus, we can use the duality gap $F(x_s) - \Phi(u_s)$ as a certificate to monitor the progress of optimization. In this section, we present the details of primal-dual HOPS.

5.1. Nesterov's Smoothing on the Dual Problem

We construct a smooth function from $\psi_\eta(u)$ that well approximates $\psi(u)$:

$$\psi_\eta(u) = \min_{x \in \Omega_1} \langle A^\top u, x\rangle + g(x) + \eta\omega(x)$$

where $\omega(x)$ is a 1-strongly convex function w.r.t. x in terms of a norm $\|\cdot\|$ [5]. Similarly, we know that $\psi_\eta(u)$ is a smooth function of u with respect to an Euclidean norm $\|u\|$ with smoothness parameter $L_\eta = \frac{1}{\eta}\|A\|_+^2$, where $\|A\|_+$ is defined by $\|A\|_+ = \max_{\|x\| \leq 1} \max_{\|u\|_+ \leq 1} \langle A^\top u, x\rangle$. Denote by

$$x_\eta(u) = \arg\min_{x \in \Omega_1} \langle A^\top u, x\rangle + g(x) + \eta\omega(x)$$

The gradient of $\psi_\eta(u)$ is computed by $\nabla\psi_\eta(u) = Ax_\eta(u)$. We can see that when η is very small, $\psi_\eta(u)$ gives a good approximation of $\psi(u)$. This motivates us to solve the following composite optimization problem

$$\max_{u \in \Omega_2} \Phi_\eta(u) \triangleq -\phi(u) + \psi_\eta(u)$$

5. This could be a general norm.

Algorithm 5 An Accelerated Proximal Gradient Method for solving dual problem (24): DAPG(u_0, t, L_η)

Input: the number of iterations t, the initial solution u_0, and the smoothness constant L_η
Let $\theta_0 = 1$, $V_{-1} = 0$, $\Gamma_{-1} = 0$, $r_0 = u_0$
Let α_k and θ_k be two sequences given in Theorem 8.
for $k = 0, \ldots, t-1$ **do**
  Compute $w_k = (1-\theta_k)u_k + \theta_k r_k$
  Compute $v_k = \nabla\psi_\eta(w_k)$, $V_k = V_{k-1} - \frac{v_k}{\alpha_k}$, and $\Gamma_k = \Gamma_{k-1} + \frac{1}{\alpha_k}$
  Compute $r_{k+1} = \Pi^{L_\eta/\sigma_2}_{V_k, \Gamma_k\phi}(u_0)$ and $u_{k+1} = \Pi^{L_\eta}_{-v_k,\phi}(w_k)$
end for
Output: u_t

Algorithm 6 Homotopy Smoothing (HOPS) for solving dual problem (24)

Input: the number of stages m and the number of iterations t per-stage, and the initial solution $u_0 \in \Omega_2$ and a parameter $b > 1$.
Let $\eta_1 = \epsilon_0/(b\widetilde{D}^2)$
for $s = 1, \ldots, m$ **do**
  Let $u_s = \text{DAPG}(u_{s-1}, t, L_{\eta_s})$
  Update $\eta_{s+1} = \eta_s/b$
end for
Output: u_m

Similar to solving the primal problem, an accelerated proximal gradient method for dual problem can be employed to solve the above problem. We present the details in Algorithm 5. We present the convergence results for Algorithm 5 in the following theorem:

Theorem 8 *(Nesterov, 2005a; Tseng, 2008) Let* $\theta_k = \frac{2}{k+2}$, $\alpha_k = \frac{2}{k+1}, k \geq 0$ *or* $\alpha_{k+1} = \theta_{k+1} = \frac{\sqrt{\theta_k^4 + 4\theta_k^2} - \theta_k^2}{2}, k \geq 0$. *For any* $u \in \Omega_2$, *we have*

$$\Phi_\eta(u) - \Phi_\eta(u_t) \leq \frac{2L_\eta \|u - u_0\|^2}{t^2} \tag{25}$$

5.2. HOPS for the Dual Problem

Similar to primal problem, we can also develop the HOPS for dual problem, which is presented in Algorithm 6. A convergence can be established similarly by exploring a local error bound condition on $\Phi(u)$. To present the convergence result, we make the following assumptions, which are similar as the primal problem.

Assumption 9 *For a concave maximization problem (24), we assume (i) there exist* $u_0 \in \Omega_2$ *and* $\epsilon_0 \geq 0$ *such that* $\max_{u\in\Omega_2} \Phi(u) - \Phi(u_0) \leq \epsilon_0$; *(ii) let* $\psi(u) = \min_{x\in\Omega_1} \langle A^\top u, x\rangle + g(x)$, *where* $g(x)$ *is a convex function; (iii) There exists a constant* $\widetilde{D}$ *such that* $\max_{x\in\Omega_1} \omega(x) \leq \widetilde{D}^2/2$.

Algorithm 7 Primal-Dual Homotopy Smoothing (PD-HOPS) for solving (1)

Input: the number of stages m, initial solutions $x_0 \in \Omega_1, u_0 \in \Omega_2$ and a parameter $b > 1$
Let $\epsilon_1 = \frac{\epsilon_0}{b}$, $\mu_1 = \frac{\epsilon_1}{\widetilde{D}^2}$, $\eta_1 = \frac{\epsilon_1}{\widetilde{D}^2}$
for $s = 1, \ldots, m$ **do**
  for $k = 0, 1, \ldots,$ **do**
    Update the sequence of x_{k+1} as in Algorithm 1 starting from x_{s-1}
    Update the sequence of u_{k+1} as in Algorithm 5 starting from u_{s-1}
    Check occasionally if $F(x_{k+1}) - \Phi(u_{k+1}) \leq 2(\epsilon_s + \epsilon)$; break the loop if it is true
  end for
  Update $x_s = x_{k+1}$ and $u_s = u_{k+1}$
  Update $\epsilon_{s+1} = \epsilon_s/b$, $\mu_{s+1} = \mu_s/b$ and $\eta_{s+1} = \eta_s/b$
end for
Output: (x_m, u_m)

Let $\widetilde{\Omega}_*$ denote the optimal solution set of (24). For any $u \in \Omega_2$, let u^* denote the closest optimal solution in $\widetilde{\Omega}_*$ to u, i.e., $u^* = \arg\min_{v \in \widetilde{\Omega}_*} \|v - u\|^2$. We denote by $\widetilde{\mathcal{L}}_\epsilon$ the ϵ-level set of $\Phi(u)$ and by $\widetilde{\mathcal{S}}_\epsilon$ the ϵ-sublevel set of $\Phi(u)$, respectively, i.e.,

$$\widetilde{\mathcal{L}}_\epsilon = \{u \in \Omega_2 : \Phi(u) = \Phi_* - \epsilon\}, \quad \widetilde{\mathcal{S}}_\epsilon = \{u \in \Omega_2 : \Phi(u) \geq \Phi_* - \epsilon\} \tag{26}$$

A local error bound condition is also imposed.

Definition 6 (Local error bound (LEB)) *A function $\Phi(u)$ is said to satisfy a local error bound condition if there exist $\tilde{\theta} \in (0, 1]$ and $\tilde{c} > 0$ such that for any $u \in \widetilde{\mathcal{S}}_\epsilon$*

$$dist(u, \widetilde{\Omega}_*) \leq \tilde{c}(\Phi_* - \Phi(u))^{\tilde{\theta}} \tag{27}$$

Theorem 10 *Suppose Assumption 9 holds and $\Phi(u)$ obeys the local error bound condition. Let HOPS for dual problem run with $t = O\left(\frac{2b\tilde{c}\widetilde{D}\|A\|_+}{\epsilon^{1-\tilde{\theta}}}\right) \geq \frac{2b\tilde{c}\widetilde{D}\|A\|_+}{\epsilon^{1-\tilde{\theta}}}$ iterations for each stage, and $m = \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$. Then*

$$\Phi_* - \Phi(u_m) \leq 2\epsilon.$$

Hence, the iteration complexity for achieving an 2ϵ-optimal solution is $\frac{2b\tilde{c}\widetilde{D}\|A\|_+}{\epsilon^{1-\tilde{\theta}}} \lceil \log_b(\frac{\epsilon_0}{\epsilon}) \rceil$ in the worst-case.

The above theorem can be proved similarly as Theorem 4.

5.3. Primal-Dual HOPS

As mentioned before, we can use the duality gap $F(x_s) - \Phi(u_s)$ as a certificate to monitor the progress of optimization to address the problem of detecting the number of iterations per-stage t. We describe the details in Algorithm 7. Suppose Assumptions 1 and 9 hold, following the analysis as in the proof of Theorem 4, when the number of iterations in the s-th

Table 1: Comparison of different optimization algorithms by the number of iterations and running time in second (mean ± standard deviation) for achieving a solution that satisfies $F(x) - F_* \leq \epsilon$.

	Linear Classification		Image Denoising		Matrix Decomposition	
	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$	$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$	$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$
PD	9861 (1.58±0.02)	27215 (4.33±0.06)	8078 (22.01±0.51)	34292 (94.26±2.67)	2523 (4.02±0.10)	3441 (5.65±0.20)
APG-D	4918 (2.44±0.22)	28600 (11.19±0.26)	179204 (924.37±59.67)	1726043 (9032.69±539.01)	1967 (6.85±0.08)	8622 (30.36±0.11)
APG-F	3277 (1.33±0.01)	19444 (7.69±0.07)	14150 (40.90±2.28)	91380 (272.45±14.56)	1115 (3.76±0.06)	4151 (9.16±0.10)
HOPS-D	1012 (0.44±0.02)	4101 (1.67±0.01)	3542 (13.77±0.13)	4501 (17.38±0.10)	224 (1.36±0.02)	313 (1.51±0.03)
HOPS-F	1009 (0.46±0.02)	4102 (1.69±0.04)	2206 (6.99±0.15)	3905 (16.52±0.08)	230 (0.91±0.01)	312 (1.23±0.01)
PD-HOPS	846 (0.36±0.01)	3370 (1.27±0.02)	2538 (7.97±0.13)	3605 (11.39±0.10)	124 (0.45±0.01)	162 (0.64±0.01)

epoch denoted by t_s satisfies $t_s \geq \max\{\frac{2bcD\|A\|}{\epsilon^{1-\theta}}, \frac{2b\tilde{c}\widetilde{D}\|A\|_+}{\epsilon^{1-\tilde{\theta}}}\}$, we can have $F(x_s) - F_* \leq \epsilon + \epsilon_s$ and $\Phi_* - \Phi(u_s) \leq \epsilon + \epsilon_s$, so that

$$F(x_s) - \Phi(u_s) \leq 2(\epsilon + \epsilon_s) \tag{28}$$

Hence, as long as the above inequality holds, we restart the next stage. Then with at most $m = \lceil \log_b(\epsilon_0/\epsilon) \rceil$ epochs we have

$$F(x_m) - \Phi(u_m) \leq 2(\epsilon + \epsilon_m) \leq 4\epsilon. \tag{29}$$

Similarly, we can show that PD-HOPS enjoys an $\widetilde{O}(\max\{1/\epsilon^{1-\theta}, 1/\epsilon^{1-\tilde{\theta}}\})$ iteration complexity, where $\tilde{\theta}$ is the exponent constant in the local error bound of the objective function for dual problem. For example, for linear classification problems with a piecewise linear loss and ℓ_1 norm regularizer we can have $\theta = 1$ and $\tilde{\theta} = 1$, and PD-HOPS enjoys a linear convergence.

6. Experiments

In this section, we present some experimental results to demonstrate the effectiveness of HOPS and PD-HOPS by comparing with two state-of-the-art algorithms, the first-order Primal-Dual (PD) method (Chambolle and Pock, 2011) and Accelerated Proximal Gradient (APG) methods. For APG, we implement two variants, where APG-D refers to the variant with the dual averaging style of update on one sequence of points (i.e., Algorithm 1) and APG-F refers to the variant of the FISTA style (Beck and Teboulle, 2009). Similarly, we also implement the two variants for HOPS. We conduct experiments for solving three problems: (1) an ℓ_1-norm regularized hinge loss for linear classification on the w1a dataset [6]; (2) a total variation based ROF model (Rudin et al., 1992) for image denoising on the Cameraman picture [7]; (3) a nuclear norm regularized absolute error minimization for low-rank and sparse matrix decomposition on a synthetic data. The three problems are discussed in details below.

6. https://www.csie.ntu.edu.tw/∼cjlin/libsvmtools/datasets/
7. http://pages.cs.wisc.edu/∼swright/TVdenoising/

- **Linear Classification:** In linear classification problems, the goal is to solve the following optimization problem:

$$\min_{x\in\mathbb{R}^d} \quad \frac{1}{n}\sum_{i=1}^{n} \ell(x^\top a_i, y_i) + \lambda r(x)$$

where $(a_i, y_i), i = 1, 2, \ldots, n$ denote pairs of and label of training data, $\ell(x^\top a_i, y_i)$ is loss function, $r(x)$ is regularizer, and λ is regularization parameter. In our experiment, we use the hinge loss (a non-smooth function) $\ell(zy) = \max(0, 1 - zy) = \max_{\alpha\in[0,1]} \alpha(1 - zy)$ for loss function and the ℓ_1-norm for regularizer:

$$\min_{x\in\mathbb{R}^d} F(x) \triangleq \frac{1}{n}\sum_{i=1}^{n} \max_{u_i\in[0,1]} u_i(1 - y_i a_i^\top x) + \lambda\|x\|_1 \tag{30}$$

We first write (30) into the following equivalent minimax formulation

$$\min_{x\in\mathbb{R}^d} \max_{u\in[0,1]^n} u^\top Ax + \frac{u^\top \mathbf{1}}{n} + \lambda\|x\|_1 \tag{31}$$

where matrix $A = -\frac{1}{n}(y_1 a_1, y_2 a_2, \ldots, y_n a_n)^\top$ and $\mathbf{1}$ is a vector of all ones. Thus, $f(x) = \max_{u\in[0,1]^n} u^\top Ax + \frac{u^\top \mathbf{1}}{n}$ and $g(x) = \lambda\|x\|_1$. To apply Nesterov's smoothing technique, we construct the following smoothed function

$$f_\mu(x) = \max_{u\in[0,1]^n} u^\top Ax + \frac{u^\top \mathbf{1}}{n} - \frac{\mu}{2}\|u\|_2^2 \tag{32}$$

We construct the experiment on the w1a dataset, which contains $2,477$ training examples and 300 features. We fix the regularization parameter $\lambda = n^{-1}$.

- **Image Denoising:** For total variation (TV) based image denoising problem, we consider the following ROF model:

$$\min_x \int_\Omega |\nabla x| + \frac{\lambda}{2}\|x - h\|_2^2, \tag{33}$$

where h is the observed noisy image, $\Omega \subset \mathbb{R}^{m\times n}$ is the image domain, $\int_\Omega |\nabla x|$ is the TV regularization term, and λ is the trade-off parameter between regularization and fidelity. Following the ROF setting in (Chambolle and Pock, 2011), we obtain the following discrete version:

$$\min_{x\in X} F(x) \triangleq \|\nabla x\|_1 + \frac{\lambda}{2}\|x - h\|_2^2. \tag{34}$$

where $X = \mathbb{R}^{mn}$ is a finite dimensional vector space, $\nabla x \in Y$ and $Y = X \times X$. The discrete gradient operator ∇x is defined as following that has two components:

$$(\nabla x)^1_{i,j} = \begin{cases} x_{i+1,j} - x_{i,j} & \text{if } i < m \\ 0 & \text{if } i = m \end{cases}$$

$$(\nabla x)^2_{i,j} = \begin{cases} x_{i,j+1} - x_{i,j} & \text{if } j < n \\ 0 & \text{if } j = n, \end{cases}$$

and $\|\nabla x\|_1$ is defined as

$$\|\nabla x\|_1 = \sum_{i,j} |(\nabla x)_{i,j}| = \sum_{i,j} \sqrt{((\nabla x)^1_{i,j})^2 + ((\nabla x)^2_{i,j})^2}.$$

According to (Chambolle and Pock, 2011), we have the minimax formulation of ROF model as

$$\min_{x \in X} \max_{u \in \Omega_2} -\langle x, \text{div} u \rangle + \frac{\lambda}{2}\|x - h\|_2^2 \tag{35}$$

where $\Omega_2 = \{u : u \in Y, \|u\|_\infty \leq 1\}$, $\|u\|_\infty = \max_{i,j} \sqrt{(u^1_{i,j})^2 + (u^2_{i,j})^2}$, and div$u$ is the discrete divergence operator (Chambolle and Pock, 2011). Thus, $f(x) = \max_{u \in \Omega_2} -\langle x, \text{div} u \rangle$ and $g(x) = \frac{\lambda}{2}\|x - h\|_2^2$. By using Nesterov's smoothing technique, we have the following smoothed function

$$\max_{u \in \Omega_2} -\langle x, \text{div} u \rangle - \frac{\mu}{2}\|u\|_2^2. \tag{36}$$

In our experiment, we use Cameraman picture of size 256×256 with additive zero mean Gaussian noise with standard deviation $\sigma = 0.05$ and we set $\lambda = 20$.

- **Matrix Decomposition:** In low-rank and sparse matrix decomposition problem, suppose given a data matrix $O \in \mathbb{R}^{m \times n}$, we aim to decompose it as

$$O = X + E$$

where $X \in \mathbb{R}^{m \times n}$ is a low-rank matrix, and $E \in \mathbb{R}^{m \times n}$ represents errors and it is sparse. We use nuclear norm regularized absolute error minimization:

$$\min_{X \in \mathbb{R}^{m \times n}} F(X) = \|X\|_* + \lambda \|E\|_1$$
$$\text{s.t. } O = X + E$$

where $\|X\|_* = \sum_i \sigma_i(X)$ denotes the nuclear norm of matrix X, i.e., the summation of singular values of matrix X, and $\|E\|_1 = \sum_{ij} |E_{ij}|$ denotes the ℓ_1-norm of E. The above formulation is equavilent to

$$\min_{X \in \mathbb{R}^{m \times n}} F(X) = \|X\|_* + \lambda \|O - X\|_1 \tag{37}$$

We first write (37) into the following equivalent minimax formulation

$$\min_{X \in \mathbb{R}^{m \times n}} \max_{\|U\|_\infty \leq 1} -\lambda \langle X, U \rangle + \lambda \langle O, U \rangle + \|X\|_* \tag{38}$$

where $U \in \mathbb{R}^{m\times n}$ and $\|U\|_\infty = \max_{ij}|U_{ij}|$. Thus, $f(X) = \max_{\|U\|_\infty\leq 1} -\lambda\langle X, U\rangle + \lambda\langle O, U\rangle$ and $g(X) = \|X\|_*$. To apply Nesterov's smoothing technique, we consider the following smoothed function

$$f_\mu(X) = \max_{\|U\|_\infty\leq 1} -\lambda\langle X, U\rangle + \lambda\langle M, U\rangle - \frac{\mu}{2}\|U\|_F^2 \tag{39}$$

We set the regularization parameter $\lambda = (\max\{m, n\})^{-0.5}$. We conduct experiment on a synthetic data with $m = n = 100$. To generate the corrupted matrix $O \in \mathbb{R}^{m\times n}$, we first obtain two orthogonal matrices $S_1 \in \mathbb{R}^{m\times k}$ and $S_2 \in \mathbb{R}^{n\times k}$ ($k = 10$) by Gaussian distribution. The low rank matrix X can be calculated by $X = S_1 S_2^\top$. Then we randomly add Gaussian noise to 10% elements of X and obtain the corrupted matrix O.

To make fair comparison, we stop each algorithm when the optimality gap is less than a given ϵ and count the number of iterations and the running time that each algorithm requires. The optimal value is obtained by running PD with a sufficiently large number of iterations such that the duality gap is very small. We repeat each algorithm 10 times for solving a particular problem and then report the averaged running time in second and the corresponding standard deviations. The running time of PD-HOPS only accounts the time for updating the primal variable since the updates for the dual variable are fully decoupled from the primal updates and can be carried out in parallel. For APG, we use the backtracking trick to tune L_μ. For HOPS, we tune the number of iterations t in each epoch among several values in the range of $\{10, 50, 100, 150, 200, 250, 300, 350, 400, 500, 1000\}$ and the parameter b among $\{1.2, 2, 2.5, 3, 3.5, 4, 5, 10, 25\}$, and report the best results. We also tune the values of parameters σ and τ and report the best results for PD. We present the comparison of different algorithms on different tasks in Table 1, where for PD-HOPS we only report the results of using the faster variant of APG, i.e., APG-F. From the results, we can see that (i) HOPS converges consistently faster than their APG variants especially when ϵ is small; (ii) PD-HOPS allows for choosing the number of iterations at each epoch automatically, yielding faster convergence speed than HOPS with manual tuning; (iii) both HOPS and PD-HOPS are significantly faster than PD.

7. Conclusions

In this paper, we have developed a homotopy smoothing (HOPS) algorithm for solving a family of structured non-smooth optimization problems with formal guarantee on the iteration complexities. We show that the proposed HOPS can achieve a lower iteration complexity of $\widetilde{O}(1/\epsilon^{1-\theta})$ with $\theta \in (0, 1]$ for obtaining an ϵ-optimal solution under a mild local error bound condition. The experimental results on three different tasks demonstrate the effectiveness of HOPS.

Acknowlegements

Y. Xu and T. Yang are partially supported by National Science Foundation (IIS-1463988, IIS-1545995).

References

Zeyuan Allen-Zhu and Elad Hazan. Optimal black-box reductions between optimization objectives. *CoRR*, abs/1603.05642, 2016.

Hédy Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-lojasiewicz inequality. *Math. Oper. Res.*, 35:438–457, 2010.

Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2:183–202, 2009.

Stephen Becker, Jérôme Bobin, and Emmanuel J. Candès. Nesta: A fast and accurate first-order method for sparse recovery. *SIAM J. Img. Sci.*, 4:1–39, 2011. ISSN 1936-4954.

Jérôme Bolte, Aris Daniilidis, and Adrian Lewis. The łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM J. Optim.*, 17:1205–1223, 2006.

James V. Burke and Paul Tseng. A unified analysis of hoffman's bound via fenchel duality. *SIAM J. Optim.*, 6(2):265–282, 1996.

Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Img. Vis.*, 40:120–145, 2011. ISSN 0924-9907.

Xi Chen, Qihang Lin, Seyoung Kim, Jaime G. Carbonell, and Eric P. Xing. Smoothing proximal gradient method for general structured sparse regression. *Ann. Appl. Stat.*, 6 (2):719–752, 06 2012.

Robert M Freund and Haihao Lu. New computational guarantees for solving convex optimization problems with first order methods, via a function growth condition measure. *CoRR*, abs/1511.02974, 2015.

Andrew Gilpin, Javier Peña, and Tuomas Sandholm. 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.

Rafal Goebel and Ralph Tyrell Rockafellar. Local strong convexity and local lipschitz continuity of the gradient of convex functions. *J. Convex Anal.*, 15(2):263–270, 2008.

Ke Hou, Zirui Zhou, Anthony Man-Cho So, and Zhi-Quan Luo. On the linear convergence of the proximal gradient method for trace norm regularization. In *NIPS*, pages 710–718, 2013.

Juan Peypouquet Bruce Suter Jerome Bolte, Trong Phong Nguyen. From error bounds to the complexity of first-order descent methods for convex functions. *CoRR*, abs/1510.08234, 2015.

Guanghui Lan, Zhaosong Lu, and Renato D. C. Monteiro. Primal-dual first-order methods with $O(1/\epsilon)$ iteration-complexity for cone programming. *Math. Program.*, 126(1):1–29, 2011.

Qihang Lin, Xi Chen, and Javier Peña. A smoothing stochastic gradient method for composite optimization. *Optimization Methods and Software*, 29(6):1281–1301, 2014.

S. Łojasiewicz. *Ensembles semi-analytiques*. Institut des Hautes Etudes Scientifiques, 1965.

Zhi-Quan Luo and Paul Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.*, 46(1):157–178, 1993.

I Necoara, Yu Nesterov, and F Glineur. Linear convergence of first order methods for non-strongly convex optimization. *CoRR, abs/1504.06298*, 2015.

Arkadi Nemirovski. Prox-method with rate of convergence o(1/t) for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. Optim.*, 15(1):229–251, January 2005.

Yu. Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1): 127–152, 2005a.

Yu. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.

Yurii Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal on Optimization*, 16(1):235–249, 2005b.

Yurii Nesterov. Smoothing technique and its applications in semidefinite optimization. *Math. Program.*, 110(2):245–259, 2007.

Francesco Orabona, Andreas Argyriou, and Nathan Srebro. PRISMA: PRoximal Iterative SMoothing Algorithm. *arXiv.org*, math.OC, 2012.

Hua Ouyang and Alexander G. Gray. Stochastic smoothing for nonsmooth minimizations: Accelerating SGD by exploiting structure. In *Proceedings of the 29th International Conference on Machine Learning, (ICML)*, 2012.

Jong-Shi Pang. Error bounds in mathematical programming. *Math. Program.*, 79:299–332, 1997.

R.T. Rockafellar. *Convex Analysis*. Princeton mathematical series. Princeton University Press, 1970. URL https://books.google.com/books?id=wDgAoAEACAAJ.

Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.

Quoc Tran-Dinh. Adaptive smoothing algorithms for nonsmooth composite convex minimization. *CoRR*, abs/1509.00106, 2015.

P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. *submitted to SIAM J. Optim.*, 2008.

Lin Xiao and Tong Zhang. A proximal-gradient homotopy method for the sparse least-squares problem. *SIAM J. Optim.*, 23(2):1062–1091, 2013.

Tianbao Yang and Qihang Lin. Rsg: Beating subgradient method without smoothness and strong convexity. *CoRR*, abs/1512.03107, 2016.

Tianbao Yang, Mehrdad Mahdavi, Rong Jin, and Shenghuo Zhu. An efficient primal-dual prox method for non-smooth optimization. *Machine Learning*, 2014.

Hui Zhang. New analysis of linear convergence of gradient-type methods via unifying error bound conditions. *arXiv:1606.00269*, 2016.

Zirui Zhou and Anthony Man-Cho So. A unified approach to error bounds for structured convex optimization problems. *CoRR, abs/1512.03518*, 2015.

Appendix A. Proof of Lemma 1

The lemma is an immediate result from (Yang and Lin, 2016). For completeness, we give the proof here.

Proof Consider $\|x\|$ to be an Euclidean norm. We first recall the definition of $x_\epsilon^\dagger$:

$$x_\epsilon^\dagger = \arg\min_{z\in\mathcal{S}_\epsilon} \|z - x\|^2 \tag{40}$$

where $\mathcal{S}_\epsilon = \{x \in \Omega_1 : F(x) \leq F_* + \epsilon\}$ is the sublevel set. We assume $x \notin \mathcal{S}_\epsilon$, otherwise the conclusion holds trivially. Thus $F(x_\epsilon^\dagger) = F_* + \epsilon$. By the first-order optimality conditions of (40), we have for any $z \in \Omega_1$, there exists $\zeta \geq 0$ (the Lagrangian multiplier of problem (40))

$$(x_\epsilon^\dagger - x + \zeta\partial F(x_\epsilon^\dagger))^\top (z - x_\epsilon^\dagger) \geq 0 \tag{41}$$

Let $z = x$ we have

$$\zeta\partial F(x_\epsilon^\dagger)^\top (x - x_\epsilon^\dagger) \geq \|x - x_\epsilon^\dagger\|^2$$

We argue that $\zeta > 0$, otherwise $x = x_\epsilon^\dagger$ contradicting to the assumption $x \notin \mathcal{S}_\epsilon$. Therefore

$$F(x) - F(x_\epsilon^\dagger) \geq \partial F(x_\epsilon^\dagger)^\top (x - x_\epsilon^\dagger) \geq \frac{\|x - x_\epsilon^\dagger\|^2}{\zeta} = \frac{\|x - x_\epsilon^\dagger\|}{\zeta}\|x - x_\epsilon^\dagger\| \tag{42}$$

Next we prove that ζ is upper bounded. Since

$$-\epsilon = F(x_\epsilon^*) - F(x_\epsilon^\dagger) \geq (x_\epsilon^* - x_\epsilon^\dagger)^\top \partial F(x_\epsilon^\dagger)$$

where x_ϵ^* is the closest point to $x_\epsilon^\dagger$ in the optimal set. Let $z = x_\epsilon^*$ in the inequality of (41), we have

$$(x_\epsilon^\dagger - x)^\top (x_\epsilon^* - x_\epsilon^\dagger) \geq \zeta(x_\epsilon^\dagger - x_\epsilon^*)^\top \partial F(x_\epsilon^\dagger) \geq \zeta\epsilon$$

Thus

$$\zeta \leq \frac{(x_\epsilon^\dagger - x)^\top (x_\epsilon^* - x_\epsilon^\dagger)}{\epsilon} \leq \frac{dist(x_\epsilon^\dagger, \Omega_*)\|x_\epsilon^\dagger - x\|}{\epsilon}$$

Therefore

$$\frac{\|x - x_\epsilon^\dagger\|}{\zeta} \geq \frac{\epsilon}{dist(x_\epsilon^\dagger, \Omega_*)}$$

Combining the above inequality with (42) we have

$$\|x - x_\epsilon^\dagger\| \leq \frac{dist(x_\epsilon^\dagger, \Omega_*)}{\epsilon}(F(x) - F(x_\epsilon^\dagger))$$

which completes the proof. ■

Appendix B. Proof of Lemma 4

The proof of Lemma 4 can be also found in (Yang and Lin, 2016). For completeness, we give the proof here.

Proof We assume that $x \notin \mathcal{S}_\epsilon$, otherwise $x = x_\epsilon^\dagger$ and the lemma holds trivially. Thus $x_\epsilon^\dagger \in \mathcal{L}_\epsilon$, i.e., $F(x_\epsilon^\dagger) = F_* + \epsilon$. Note that

$$\frac{\partial \|x\|_p}{\partial x_i} = \frac{|x_i|^{p-1} sign(x_i)}{\|x\|_p^{p-1}}, \quad (\nabla \frac{1}{2}\|x\|_p^2)_i = \|x\|_p^{2-p}|x_i|^{p-1} sign(x_i)$$

By the definition of $x_\epsilon^\dagger$ and the Lagrangian theory, there exists a Lagrangian multiplier $\zeta \geq 0$ and a subgradient $v_\epsilon^\dagger \in \partial F(x_\epsilon^\dagger)$ such that

$$\|x_\epsilon^\dagger - x\|_p^{2-p}|[x_\epsilon^\dagger - x]_i|^{p-1}\text{sign}([x_\epsilon^\dagger - x]_i) + \zeta[v_\epsilon^\dagger]_i = 0, \forall i.$$

It is clear that $\zeta > 0$, otherwise $x = x_\epsilon^\dagger$ that contradicts to $x \notin \mathcal{S}_\epsilon$. By the convexity of $F(\cdot)$ we have

$$\begin{aligned} F(x) - F(x_\epsilon^\dagger) \geq (x - x_\epsilon^\dagger)^\top v_\epsilon^\dagger &= \frac{1}{\zeta}\|x_\epsilon^\dagger - x\|_p^{2-p}\sum_{i=1}^d |[x_\epsilon^\dagger - x]_i|^{p-1}\text{sign}([x_\epsilon^\dagger - x]_i)[x_\epsilon^\dagger - x]_i \\ &= \frac{1}{\zeta}\|x_\epsilon^\dagger - x\|_p^{2-p}\|x_\epsilon^\dagger - x\|_p^p = \frac{\|x_\epsilon^\dagger - x\|_p^2}{\zeta} \end{aligned}$$

Next, we bound ζ. Let $1/p + 1/q = 1$. Since

$$\begin{aligned} \zeta^q \sum_i |[v_\epsilon^\dagger]_i|^q &= \|x_\epsilon^\dagger - x\|_p^{q(2-p)} \sum_{i=1}^d |[x_\epsilon^\dagger - x]_i|^{q(p-1)} = \|x_\epsilon^\dagger - x\|_p^{q(2-p)} \sum_{i=1}^d |[x_\epsilon^\dagger - x]_i|^p \\ &= \|x_\epsilon^\dagger - x\|_p^{q(2-p)+p} = \|x_\epsilon^\dagger - x\|_p^{p/(p-1)} \end{aligned} \tag{43}$$

Thus

$$\frac{1}{\zeta} \geq \frac{\|v_\epsilon^\dagger\|_q}{\|x_\epsilon^\dagger - x\|_p^{p/(q(p-1))}}$$

To lower bound $\|v_\epsilon^\dagger\|_q$, we explore the convexity of $F(x)$. By the convexity of $F(\cdot)$,

$$F(x_\epsilon^*) - F(x_\epsilon^\dagger) \geq (x_\epsilon^* - x_\epsilon^\dagger)^\top v_\epsilon^\dagger$$

where x_ϵ^* is the closest point in Ω_* to $x_\epsilon^\dagger$. Then we have

$$\|x_\epsilon^* - x_\epsilon^\dagger\|_p \|v_\epsilon^\dagger\|_q \geq -(x_\epsilon^* - x_\epsilon^\dagger)^\top v_\epsilon^\dagger \geq F(x_\epsilon^\dagger) - F(x_\epsilon^*) = F(x_\epsilon^\dagger) - F_* = \epsilon$$

Then

$$\frac{1}{\zeta} \geq \frac{\epsilon}{\|x_\epsilon^\dagger - x\|_p^{p/(q(p-1))} \|x_\epsilon^* - x_\epsilon^\dagger\|_p} = \frac{\epsilon}{\|x_\epsilon^\dagger - x\|_p \|x_\epsilon^* - x_\epsilon^\dagger\|_p}$$

Therefore

$$F(x) - F(x_\epsilon^\dagger) \geq \frac{\|x_\epsilon^\dagger - x\|_p^2}{\zeta} \geq \frac{\epsilon \|x_\epsilon^\dagger - x\|_p^2}{\|x_\epsilon^\dagger - x\|_p \|x_\epsilon^* - x_\epsilon^\dagger\|_p} = \frac{\epsilon \|x_\epsilon^\dagger - x\|_p}{\|x_\epsilon^* - x_\epsilon^\dagger\|_p}$$

i.e.,

$$\|x_\epsilon^\dagger - x\|_p \leq \frac{\|x_\epsilon^* - x_\epsilon^\dagger\|_p}{\epsilon}(F(x) - F(x_\epsilon^\dagger)) = \frac{dist_p(x_\epsilon^\dagger, \Omega_*)}{\epsilon}(F(x) - F(x_\epsilon^\dagger))$$

■