
Convergence Analysis of Proximal Gradient with Momentum for Nonconvex Optimization

Qunwei Li¹ Yi Zhou¹ Yingbin Liang¹ Pramod K. Varshney¹

Abstract

In this work, we investigate the accelerated proximal gradient method for nonconvex programming (APGnc). The method compares between a usual proximal gradient step and a linear extrapolation step, and accepts the one that has a lower function value to achieve a monotonic decrease. In specific, under a general nonsmooth and nonconvex setting, we provide a rigorous argument to show that the limit points of the sequence generated by APGnc are critical points of the objective function. Then, by exploiting the Kurdyka-Łojasiewicz (KŁ) property for a broad class of functions, we establish the linear and sub-linear convergence rates of the function value sequence generated by APGnc. We further propose a stochastic variance reduced APGnc (SVRG-APGnc), and establish its linear convergence under a special case of the KŁ property. We also extend the analysis to the inexact version of these methods and develop an adaptive momentum strategy that improves the numerical performance.

1. Introduction

Many problems in machine learning, data mining, and signal processing can be formulated as the following composite minimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x}). \quad (\text{P})$$

Typically, $f : \mathbb{R}^d \rightarrow \mathbb{R}$ captures the loss of data fitting and can be written as $f = \frac{1}{n} \sum_{l=1}^n f_l$ with each f_l corresponding to the loss of one sample. The second term $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is the regularizer that promotes desired structures on the solution based on prior knowledge of the problem.

¹Syracuse University, NY, USA. Correspondence to: Qunwei Li <qli33@syr.edu>.

Algorithm 1 APG

Input: $\mathbf{y}_1 = \mathbf{x}_1 = \mathbf{x}_0, t_1 = 1, t_0 = 0, \eta < \frac{1}{L}$.
for $k = 1, 2, \dots$ **do**
 $\mathbf{y}_k = \mathbf{x}_k + \frac{t_{k-1}-1}{t_k}(\mathbf{x}_k - \mathbf{x}_{k-1})$.
 $\mathbf{x}_{k+1} = \text{prox}_{\eta g}(\mathbf{y}_k - \eta \nabla f(\mathbf{y}_k))$.
 $t_{k+1} = \frac{\sqrt{4t_k^2+1}+1}{2}$.
end for

In practice, many problems of (P) are formulated, either naturally or intentionally, into a convex model to guarantee the tractability of algorithms. In particular, such convex problems can be efficiently minimized by many first-order algorithms, among which the accelerated proximal gradient (APG) method (also referred to as FISTA (Beck & Teboulle, 2009b)) is proven to be the best for minimizing such class of convex functions. We present one of its basic forms in Algorithm 1. Compared to the usual proximal gradient step, the APG algorithm takes an extra linear extrapolation step for acceleration. It has been shown (Beck & Teboulle, 2009b) that the APG method reduces the function value gap at a rate of $\mathcal{O}(1/k^2)$ where k denotes the number of iterations. This convergence rate meets the theoretical lower bound of first-order gradient methods for minimizing smooth convex functions. The reader can refer to (Tseng, 2010) for other variants of APG.

Algorithm 2 Monotone APG (mAPG)

Input: $\mathbf{y}_1 = \mathbf{x}_1 = \mathbf{x}_0, t_1 = 1, t_0 = 0, \eta < \frac{1}{L}$.
for $k = 1, 2, \dots$ **do**
 $\mathbf{y}_k = \mathbf{x}_k + \frac{t_{k-1}}{t_k}(\mathbf{z}_k - \mathbf{x}_k) + \frac{t_{k-1}-1}{t_k}(\mathbf{x}_k - \mathbf{x}_{k-1})$.
 $\mathbf{z}_{k+1} = \text{prox}_{\eta g}(\mathbf{y}_k - \eta \nabla f(\mathbf{y}_k))$.
 $\mathbf{v}_{k+1} = \text{prox}_{\eta g}(\mathbf{x}_k - \eta \nabla f(\mathbf{x}_k))$.
 $t_{k+1} = \frac{\sqrt{4t_k^2+1}+1}{2}$.
 if $F(\mathbf{z}_{k+1}) \leq F(\mathbf{v}_{k+1})$ **then**
 $\mathbf{x}_{k+1} = \mathbf{z}_{k+1}$,
 else if $F(\mathbf{v}_{k+1}) \leq F(\mathbf{z}_{k+1})$ **then**
 $\mathbf{x}_{k+1} = \mathbf{v}_{k+1}$.
 end if
end for

Although convex problems are tractable and can be glob-

Algorithm 3 APG non-convex problem (APGnc)

Input: $\mathbf{y}_1 = \mathbf{x}_0, \beta_k = \frac{k}{k+3}, \eta < \frac{1}{L}$.
for $k = 1, 2, \dots$ **do**
 $\mathbf{x}_k = \text{prox}_{\eta g}(\mathbf{y}_k - \eta \nabla f(\mathbf{y}_k))$.
 $\mathbf{v}_k = \mathbf{x}_k + \beta_k(\mathbf{x}_k - \mathbf{x}_{k-1})$.
 if $F(\mathbf{x}_k) \leq F(\mathbf{v}_k)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{x}_k$,
 else if $F(\mathbf{v}_k) \leq F(\mathbf{x}_k)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{v}_k$.
 end if
end for

ally minimized, many applications naturally require to solve nonconvex optimization problems of (\mathbf{P}) . Recently, several variants of the APG method have been proposed for nonconvex problems, and two major ones are presented in Algorithm 2 and Algorithm 3, respectively. The major difference to the original APG is that the modified methods only accept the new iterate when the corresponding function value is sufficiently decreased, which leads to a more stable convergence behavior. In particular, (Li & Lin, 2015) analyzed mAPG (Algorithm 2) by exploiting the Kurdyka-Łojasiewicz (KL) property, which is a local geometrical structure very generally held by a large class of nonconvex objective functions, and has been successfully exploited to characterize the asymptotic convergence behavior of many first order methods. It was shown in (Li & Lin, 2015) that mAPG achieves the $\mathcal{O}(1/k^2)$ convergence rate for convex problems of (\mathbf{P}) , and converges to a critical point at sublinear and linear rates under different cases of the KL property for nonconvex problems. Despite the desirable convergence rate, mAPG requires two proximal steps, which doubles the computational complexity of the original APG. In comparison, the APGnc (Algorithm 3) requires only one proximal step, and hence computes faster than mAPG in each iteration. However, the analysis of APGnc in (Yao & Kwok, 2016) does not exploit the KL property and no convergence rate of the function value is established. Hence, there is still no formal theoretical comparison of the overall performance (which depends on both computation per iteration and convergence rate) between mAPG and APGnc. It is unclear whether the computational saving per iteration in APGnc is at the cost of lower convergence rate.

The goal of this paper is to provide a comprehensive analysis of the APGnc algorithm under the KL framework, thus establishing a rigorous comparison between mAPG and APGnc and formally justifying the overall advantage of APGnc.

1.1. Main Contributions

This paper provides the convergence analysis of APGnc type algorithms for *nonconvex* problems of (\mathbf{P}) under the KL framework as well as the inexact situation. We also study the stochastic variance reduced APGnc algorithm and its inexact situation. Our analysis requires novel technical treatments to exploit the KL property due to the joint appearance of the following ingredients in the algorithms including momentum terms, inexact errors, and stochastic variance reduced gradients. Our contributions are summarized as follows.

For APGnc applied to nonconvex problems of (\mathbf{P}) , we show that the limit points of the sequences generated by APGnc are critical points of the objective function. Then, by exploiting different cases of the Kurdyka-Łojasiewicz property of the objective function, we establish the linear and sub-linear convergence rates of the function value sequence generated by APGnc. Our results formally show that APGnc (with one proximal map per iteration) achieves the same convergence properties as well as the convergence rates as mAPG (with two proximal maps per iteration) for nonconvex problems, thus establishing its overall computational advantage.

We further propose an APGnc⁺ algorithm, which is an improved version of APGnc by adapting the momentum step-size (see Algorithm 4), and shares the same theoretical convergence rate as APGnc but numerically performs better than APGnc.

Furthermore, we study the inexact APGnc in which the computation of the gradient and the proximal mapping may have errors. We show that the algorithm still achieves the convergence rate at the same order as the exact case as long as the inexactness is properly controlled. We also explicitly characterize the impact of errors on the constant factors that affect the convergence rate.

To facilitate the solution to large-scale optimization problems, we study the stochastic variance reduced APGnc (SVRG-APGnc), and show that such an algorithm achieves linear convergence rate under a certain case of the KL property. We further analyze the inexact SVRG-APGnc and show that it also achieves the linear convergence under the same KL property as long as the error in the proximal mapping is bounded properly. This is the first analysis of the SVRG proximal algorithm with *momentum* that exploits the KL structure to establish linear convergence rate for nonconvex programming.

Our numerical results further corroborate the theoretic analysis. We demonstrate that APGnc/APGnc⁺ outperforms APG and mAPG for nonconvex problems in both exact and inexact cases, and in both deterministic and stochastic variants of the algorithms. Furthermore,

APGnc⁺ outperforms APGnc due to properly chosen momentum stepsize.

1.2. Comparison to Related Work

APG algorithms: The original accelerated gradient method for minimizing a single smooth convex function dates back to (Nesterov, 1983), and is further extended as APG in the composite minimization framework in (Beck & Teboulle, 2009b; Tseng, 2010). While these APG variants generate a sequence of function values that may oscillate, (Beck & Teboulle, 2009a) proposed another variant of APG that generates a non-increasing sequence of function values. Then, (Li & Lin, 2015) further proposed an mAPG that generates a sufficiently decreasing sequence of function values, and established the asymptotic convergence rates under the KL property. Recently, (Yao & Kwok, 2016) proposed APGnc, which is a more efficient version of APG for nonconvex problems, but the analysis only characterizes fixed points and did not exploit the KL property to characterize the convergence rate. A unified treatment of accelerated gradient method for nonconvex stochastic optimization is presented in (Ghadimi & Lan, 2016). But the discussion does not exploit the KL property, and requires function g to be convex. Our study establishes the convergence rate analysis of APGnc under the KL property.

Nonconvex optimization under KL: The KL property (Bolte et al., 2007) is an extension of the Łojasiewicz gradient inequality (Łojasiewicz, 1965) to the nonsmooth case. Many first-order descent methods, under the KL property, can be shown to converge to a critical point (Attouch & Bolte, 2009; Attouch et al., 2010; Bolte et al., 2014) with different types of asymptotic convergence rates. (Li & Lin, 2015) and our paper focuses on the first-order algorithms with *momentum*, and respectively analyze mAPG and APGnc by exploiting the KL property.

Inexact algorithms under KL: (Attouch et al., 2013; Frankel et al., 2015) studied the inexact proximal algorithm under the KL property. This paper studies the inexact proximal algorithm with momentum (i.e., APGnc) under the KL property. While (Yao & Kwok, 2016) also studied the inexact APGnc, the analysis did not exploit the KL property to characterize the convergence rate.

Nonconvex SVRG: SVRG was first proposed in (Johnson & Zhang, 2013), to accelerate the stochastic gradient method for strongly convex objective functions, and was studied for the convex case in (Zhu & Yuan, 2016). Recently, SVRG was further studied for smooth nonconvex optimization in Reddi et al. (2016a). Then in (Reddi et al., 2016b), the proximal SVRG was proposed and studied for nonsmooth and nonconvex optimization. Our paper further incorporates SVRG for the proximal gradient with *momentum* in the nonconvex case. Furthermore, we exploit a cer-

tain KL property in our analysis that is very different from the PL property exploited in (Reddi et al., 2016a), and requires special technical treatment in convergence analysis.

2. Preliminaries and Assumptions

In this section, we first introduce some technical definitions that are useful later on, and then describe the assumptions on the problem **(P)** that we take in this paper.

Throughout this section, $h : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is an extended real-valued function that is proper, i.e., its domain $\text{dom } h := \{\mathbf{x} \in \mathbb{R}^d : h(\mathbf{x}) < \infty\}$ is nonempty, and is closed, i.e., its sublevel sets $\{\mathbf{x} \in \mathbb{R}^d : h(\mathbf{x}) \leq \alpha\}$ are closed for all $\alpha \in \mathbb{R}$. Note that a proper and closed function h can be nonsmooth and nonconvex, hence we consider the following generalized notion of derivative.

Definition 1 (Subdifferential, (Rockafellar & Wets, 1997)). *The Fréchet subdifferential $\hat{\partial}h$ of h at $\mathbf{x} \in \text{dom } h$ is the set of $\mathbf{u} \in \mathbb{R}^d$ such that*

$$\liminf_{\mathbf{z} \neq \mathbf{x}, \mathbf{z} \rightarrow \mathbf{x}} \frac{h(\mathbf{z}) - h(\mathbf{x}) - \mathbf{u}^\top (\mathbf{z} - \mathbf{x})}{\|\mathbf{z} - \mathbf{x}\|} \geq 0,$$

while the (limiting) subdifferential ∂h at $\mathbf{x} \in \text{dom } h$ is the graphical closure of $\hat{\partial}h$:

$$\{\mathbf{u} : \exists (\mathbf{x}^k, h(\mathbf{x}^k)) \rightarrow (\mathbf{x}, h(\mathbf{x})), \hat{\partial}h(\mathbf{x}^k) \ni \mathbf{u}^k \rightarrow \mathbf{u}\}.$$

In particular, this generalized derivative reduces to ∇h when h is continuously differentiable, and reduces to the usual subdifferential when h is convex.

Definition 2 (Critical point). *A point $\mathbf{x} \in \mathbb{R}^d$ is a critical point of h iff $\mathbf{0} \in \partial h(\mathbf{x})$.*

Definition 3 (Distance). *The distance of a point $\mathbf{x} \in \mathbb{R}^d$ to a closed set $\Omega \subseteq \mathbb{R}^d$ is defined as:*

$$\text{dist}_\Omega(\mathbf{x}) := \min_{\mathbf{y} \in \Omega} \|\mathbf{y} - \mathbf{x}\|. \quad (1)$$

Definition 4 (Proximal map, e.g. (Rockafellar & Wets, 1997)). *The proximal map of a point $\mathbf{x} \in \mathbb{R}^d$ under a proper and closed function h with parameter $\eta > 0$ is defined as:*

$$\text{prox}_{\eta h}(\mathbf{x}) := \text{argmin}_{\mathbf{z}} h(\mathbf{z}) + \frac{1}{2\eta} \|\mathbf{z} - \mathbf{x}\|^2, \quad (2)$$

where $\|\cdot\|$ is the Euclidean l_2 norm.

We note that when h is convex, the corresponding proximal map is the minimizer of a strongly convex function, i.e., a singleton. But for nonconvex h , the proximal map can be set-valued, in which case $\text{prox}_{\eta h}(\mathbf{x})$ stands for an arbitrary element from that set. The proximal map is a popular tool to handle the nonsmooth part of the objective function, and is the key component of proximal-like algorithms (Beck & Teboulle, 2009b; Bolte et al., 2014).

Definition 5 (Uniformized KL property, (Bolte et al., 2014)). *Function h is said to satisfy the uniformized KL property if for every compact set $\Omega \subset \text{dom}h$ on which h is constant, there exist $\varepsilon, \lambda > 0$ such that for all $\bar{\mathbf{x}} \in \Omega$ and all $\mathbf{x} \in \{\mathbf{x} \in \mathbb{R}^d : \text{dist}_\Omega(\mathbf{x}) < \varepsilon\} \cap [\mathbf{x} : h(\bar{\mathbf{x}}) < h(\mathbf{x}) < h(\bar{\mathbf{x}}) + \lambda]$, one has*

$$\varphi'(h(\mathbf{x}) - h(\bar{\mathbf{x}})) \cdot \text{dist}_{\partial h(\mathbf{x})}(\mathbf{0}) \geq 1, \quad (3)$$

where the function $\varphi : [0, \lambda) \rightarrow \mathbb{R}_+$ takes the form $\varphi(t) = \frac{c}{\theta} t^\theta$ for some constants $c > 0, \theta \in (0, 1]$.

The above definition is a modified version of the original KL property (Bolte et al., 2010; Kurdyka, 1998), and is more convenient for our analysis later. The KL property is a generalization of the Łojasiewicz gradient inequality to nonsmooth functions (Bolte et al., 2007), and it is a powerful tool to analyze a class of first-order descent algorithms (Attouch & Bolte, 2009; Attouch et al., 2010; Bolte et al., 2014). In particular, the class of semi-algebraic functions satisfy the above KL property. This function class covers most objective functions in real applications, for instance, all l_p where $p \geq 0$ and is rational, real polynomials, rank, etc. For a more detailed discussion and a list of examples of KL functions, see (Bolte et al., 2014) and (Attouch et al., 2010).

We adopt the following assumptions on the problem (P) in this paper.

Assumption 1. *Regarding the functions f, g (and $F = f + g$) in (P)*

1. *They are proper and lower semicontinuous; $\inf_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) > -\infty$; the sublevel set $\{\mathbf{x} \in \mathbb{R}^d : F(\mathbf{x}) \leq \alpha\}$ is bounded for all $\alpha \in \mathbb{R}$;*
2. *They satisfy the uniformized KL property;*
3. *Function f is continuously differentiable and the gradient ∇f is L -Lipschitz continuous.*

Note that the sublevel set of F is bounded when either f or g has bounded sublevel set, i.e., $f(\mathbf{x})$ or $g(\mathbf{x}) \rightarrow +\infty$ as $\|\mathbf{x}\| \rightarrow +\infty$. Of course, we do *not* assume convexity on either f or g , and the KL property serves as an alternative in this general setting.

3. Main Results

In this section, we provide our main results on the convergence analysis of APGnc and SVRG-APGnc as well as inexact variants of these algorithms. All proofs of the theorems are provided in supplemental materials.

3.1. Convergence Analysis

In this subsection, we characterize the convergence of APGnc. Our first result characterizes the behavior of the limit points of the sequence generated by APGnc.

Theorem 1. *Let Assumption 1. {1,3} hold for the problem (P). Then with stepsize $\eta < \frac{1}{L}$, the sequence $\{\mathbf{x}_k\}$ generated by APGnc satisfies*

1. *$\{\mathbf{x}_k\}$ is a bounded sequence;*
2. *The set of limit points Ω of $\{\mathbf{x}_k\}$ forms a compact set, on which the objective function F is constant;*
3. *All elements of Ω are critical points of F .*

Theorem 1 states that the sequence $\{\mathbf{x}_k\}$ generated by APGnc eventually approaches a compact set (i.e., a closed and bounded set in \mathbb{R}^d) of critical points, and the objective function remains constant on it. Here, approaching critical points establishes the first step for solving general nonconvex problems. Moreover, the compact set Ω meets the requirements of the uniform KL property, and hence provides a seed to exploit the KL property around it. Next, we further utilize the KL property to establish the asymptotic convergence rate for APGnc. In the following theorem, θ is the parameter in the uniformized KL property via the function φ that takes the form $\varphi(t) = \frac{c}{\theta} t^\theta$ for some $c > 0, \theta \in (0, 1]$.

Theorem 2. *Let Assumption 1. {1,2,3} hold for the problem (P). Let $F(\mathbf{x}) \equiv F^*$ for all $\mathbf{x} \in \Omega$ (the set of limit points), and denote $r_k := F(\mathbf{x}_k) - F^*$. Then with stepsize $\eta < \frac{1}{L}$, the sequence $\{r_k\}$ satisfies for k_0 large enough*

1. *If $\theta = 1$, then r_k reduces to zero in finite steps;*
2. *If $\theta \in [\frac{1}{2}, 1)$, then $r_k \leq (\frac{c^2 d_1}{1+c^2 d_1})^{k-k_0} r_{k_0}$;*
3. *If $\theta \in (0, \frac{1}{2})$, then $r_k \leq (\frac{c}{(k-k_0)d_2(1-2\theta)})^{\frac{1}{1-2\theta}}$,*

where $d_1 = (\frac{1}{\eta} + L) / (\frac{1}{2\eta} - \frac{L}{2})$ and $d_2 = \min\{\frac{1}{2cd_1}, \frac{c}{1-2\theta}(2^{\frac{2\theta-1}{2\theta-2}} - 1)r_{k_0}^{2\theta-1}\}$.

Theorem 2 characterizes three types of convergence behaviors of APGnc, depending on θ that parameterizes the KL property that the objective function satisfies. An illustrative example for the first kind ($\theta = 1$) can take a form similar to $F(x) = |x|$ for $x \in \mathbb{R}$ around the critical points. The function is ‘sharp’ around its critical point $x = 0$ and thus the iterates slide down quickly onto it within finite steps. For the second kind ($\theta \in [\frac{1}{2}, 1)$), example functions can take a form similar to $F(x) = x^2$ around the critical points. That is, the function is strongly convex-like and hence the convergence rate is typically linear. Lastly, functions of the third kind are ‘flat’ around its critical points and thus the convergence is slowed down to sub-linear rate. We note that characterizing the value of θ for a given function is a highly non-trivial problem that takes much independent effort (Kurdyka & Spodzieja, 2011; Li & Kei, 2016). Nevertheless, KL property provides a general picture of the asymptotic convergence behaviors of APGnc.

Algorithm 4 APGnc with adaptive momentum (APGnc⁺)

Input: $\mathbf{y}_1 = \mathbf{x}_0, \beta, t \in (0, 1), \eta < \frac{1}{L}$.
for $k = 1, 2, \dots$ **do**
 $\mathbf{x}_k = \text{prox}_{\eta g}(\mathbf{y}_k - \eta \nabla f(\mathbf{y}_k))$.
 $\mathbf{v}_k = \mathbf{x}_k + \beta(\mathbf{x}_k - \mathbf{x}_{k-1})$.
 if $F(\mathbf{x}_k) \leq F(\mathbf{v}_k)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{x}_k, \beta \leftarrow t\beta$.
 else if $F(\mathbf{v}_k) \leq F(\mathbf{x}_k)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{v}_k, \beta \leftarrow \min\{\frac{\beta}{t}, 1\}$.
 end if
end for

3.2. APGnc with Adaptive Momentum

The original APGnc sets the momentum parameter $\beta_k = \frac{k}{k+3}$, which can be theoretically justified only for convex problems. We here propose an alternative choice of the momentum stepsize that is more intuitive for nonconvex problems, and refer to the resulting algorithm as APGnc⁺ (See Algorithm 4). The idea is to enlarge the momentum β to further exploit the opportunity of acceleration when the extrapolation step \mathbf{v}_k achieves a lower function value. Since the proofs of Theorem 1 and Theorem 2 do not depend on the exact value of the momentum stepsize, APGnc and APGnc⁺ have the same order-level convergence rate. However, we show in Section 4 that APGnc⁺ improves upon APGnc numerically.

3.3. Inexact APGnc

We further consider inexact APGnc, in which computation of the proximal gradient step may be inexact, i.e.,

$$\mathbf{x}_k = \text{prox}_{\eta g}^{\epsilon_k}(\mathbf{y}_k - \eta(\nabla f(\mathbf{y}_k) + \mathbf{e}_k)),$$

where \mathbf{e}_k captures the inexactness of computation of $\nabla f(\mathbf{y}_k)$, and ϵ_k captures the inexactness of evaluation of the proximal map as given by

$$\begin{aligned} \mathbf{x} &= \text{prox}_{\eta g}^{\epsilon}(\mathbf{y}) \\ &= \{\mathbf{u} \mid g(\mathbf{u}) + \frac{1}{2\eta}\|\mathbf{u} - \mathbf{y}\|^2 \\ &\leq \epsilon + g(\mathbf{v}) + \frac{1}{2\eta}\|\mathbf{v} - \mathbf{y}\|^2, \quad \forall \mathbf{v} \in \mathbb{R}^d\}. \end{aligned} \quad (4)$$

The inexact proximal algorithm has been studied in (Attouch et al., 2013) for nonconvex functions under the KŁ property. Our study here is the first treatment of inexact proximal algorithms with momentum (i.e., APG-like algorithms). Furthermore, previous studies addressed only the inexactness of gradient computation for nonconvex problems, but our study here also includes the inexactness of the proximal map for nonconvex problems requiring only g to be convex as the second case we specify below.

We study the following two cases.

1. g is convex;
2. g is nonconvex, and $\epsilon = 0$.

In the first case, $\partial g(\mathbf{x})$ reduces to the usual subdifferential of convex functions, and the inexactness ϵ naturally induces the following ϵ -subdifferential

$$\partial_{\epsilon} g(\mathbf{x}) = \{\mathbf{u} \mid g(\mathbf{y}) \geq g(\mathbf{x}) + \langle \mathbf{y} - \mathbf{x}, \mathbf{u} \rangle - \epsilon, \forall \mathbf{y} \in \mathbb{R}^d\}.$$

Moreover, since the KŁ property utilizes the information of ∂F , we then need to characterize the perturbation of ∂g under the inexactness ϵ . This leads to the following definition.

Definition 6. For any $\mathbf{x} \in \mathbb{R}^d$, let $\mathbf{u}' \in \partial_{\epsilon} g(\mathbf{x})$ such that $\nabla f(\mathbf{x}) + \mathbf{u}'$ has the minimal norm. Then the perturbation between ∂g and $\partial_{\epsilon} g$ is defined as $\xi := \text{dist}_{\partial g(\mathbf{x})}(\mathbf{u}')$.

The following theorem states that for nonconvex functions, as long as the inexactness parameters \mathbf{e}_k, ϵ_k and ξ_k are properly controlled, then the inexact APGnc converges at the same order-level rate as the corresponding exact algorithm.

Theorem 3. Consider the above two cases for inexact APGnc under Assumption 1. $\{1, 2, 3\}$. If for all $k \in \mathbb{N}$

$$\|\mathbf{e}_k\| \leq \gamma \|\mathbf{x}_k - \mathbf{y}_k\|, \quad (5)$$

$$\epsilon_k \leq \delta \|\mathbf{x}_k - \mathbf{y}_k\|^2, \quad (6)$$

$$\xi_k \leq \lambda \|\mathbf{x}_k - \mathbf{y}_k\|, \quad (7)$$

then all the statements in Theorem 1 remain true and the convergence rates in Theorem 2 remain at the same order with the constants $d_1 = (\frac{1}{\eta} + L + C)^2 / (\frac{1}{2\eta} - \frac{L}{2} - C)$, where $C > 0$ depends on γ, δ and λ , and $d_2 = \min\{\frac{1}{2cd_1}, \frac{c}{1-2\theta}(2^{\frac{2\theta-1}{2\theta-2}} - 1)r_{k_0}^{2\theta-1}\}$. Correspondingly, a smaller stepsize $\eta < \frac{1}{2C+L}$ should be used.

It can be seen that, due to the inexactness, the constant factor d_1 in Theorem 2 is enlarged, which further leads to a smaller d_2 in Theorem 2. Hence, the corresponding convergence rates are slower compared to the exact case, but remain at the same order.

3.4. Stochastic Variance Reduced APGnc

In this subsection, we study the stochastic variance reduced APGnc algorithm, referred to as SVRG-APGnc. The main steps are presented in Algorithm 5. The major difference from APGnc is that the single proximal gradient step is replaced by a loop of stochastic proximal gradient steps using variance reduced gradients.

Due to the stochastic nature of the algorithm, the iterate sequence may not stably stay in the local KŁ region, and hence the standard KŁ approach fails. We then focus on the analysis of the special but important case of the global KŁ

Algorithm 5 SVRG-APGnc

Input: $\mathbf{y}_0 = \mathbf{x}_0^0, \beta_k = \frac{k}{k+3}, m, \eta < \frac{1}{2mL}$.
for $k = 0, 1, 2, \dots$ **do**
 $\mathbf{x}_k^0 = \mathbf{y}_k, \mathbf{g}_k = \nabla f(\mathbf{y}_k)$.
for $t = 0, 1, \dots, m-1$ **do**
 sample ξ from $\{1, 2, \dots, n\}$.
 $\mathbf{v}_k^t = \nabla f_\xi(\mathbf{x}_k^t) - \nabla f_\xi(\mathbf{y}_k) + \mathbf{g}_k$.
 $\mathbf{x}_k^{t+1} = \text{prox}_{\eta g}(\mathbf{x}_k^t - \eta \mathbf{v}_k^t)$.
end for
 $\mathbf{z}_k = \mathbf{x}_k^m + \beta_k(\mathbf{x}_k^m - \mathbf{x}_{k-1}^m)$.
if $F(\mathbf{x}_k^m) \leq F(\mathbf{z}_k)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{x}_k^m$,
else if $F(\mathbf{z}_k) \leq F(\mathbf{x}_k^m)$ **then**
 $\mathbf{y}_{k+1} = \mathbf{z}_k$.
end if
end for

property with $\theta = \frac{1}{2}$. In fact, if $g = 0$, the KŁ property in such a case reduces to the well known Polyak-Łojasiewicz (PL) inequality studied in (Karimi et al., 2016). Various nonconvex problems have been shown to satisfy this property such as quadratic phase retrieval loss function (Zhou et al., 2016) and neural network loss function (Hardt & Ma, 2016). The following theorem characterizes the convergence rate of SVRG-APGnc under the KŁ property with $\theta = \frac{1}{2}$.

Theorem 4. *Let $\eta = \rho/L$, where $\rho < 1/2$ and satisfies $4\rho^2 m^2 + \rho \leq 1$. If the problem (P) satisfies the KŁ property globally with $\theta = 1/2$, then the sequence $\{\mathbf{y}_k\}$ generated by Algorithm 5 satisfies*

$$\mathbb{E}[F(\mathbf{y}_k) - F^*] \leq \left(\frac{d}{d+1}\right)^k (F(\mathbf{y}_0) - F^*), \quad (8)$$

where $d = \frac{c^2(L + \frac{1}{\eta})^2 + \eta L^2 m}{\frac{1}{2\eta} - L}$, and F^* is the optimal function value.

Hence, SVRG-APGnc also achieves the linear convergence rate under the KŁ property with $\theta = \frac{1}{2}$. We note that Theorem 4 differs from the linear convergence result established in (Reddi et al., 2016b) for the SVRG proximal gradient in two folds: (1) we analyze proximal gradient with momentum but (Reddi et al., 2016b) studied proximal gradient algorithm; (2) the KŁ property with $\theta = \frac{1}{2}$ here is different from the generalized PL inequality for composite functions adopted by (Karimi et al., 2016). In order to exploit the KŁ property, our analysis of the convergence rate requires novel treatments of bounds, which can be seen in the proof of Theorem 4 in ??.

3.5. Inexact SVRG-APGnc

We further study the inexact SVRG-APGnc algorithm, and the setting of inexactness is the same as that in Section 3.3.

Here, we focus on the case where g is convex and $\mathbf{e}_k = 0$. The following theorem characterizes the convergence rate under such an inexact case.

Theorem 5. *Let g be convex and consider only the inexactness ϵ in the proximal map. Assume the KŁ property is globally satisfied with $\theta = 1/2$. Set $\eta = \rho/L$ where $\rho < 1/2$ and satisfies $8\rho^2 m^2 + \rho \leq 1$. Assume that $\sum_{t=0}^{m-1} 3\mathbb{E}[\epsilon_k^t] \leq \alpha \sum_{t=0}^{m-1} \mathbb{E}[\|\bar{\mathbf{x}}_k^{t+1} - \mathbf{x}_k^t\|^2]$ for some $\alpha > 0$, and define $\bar{\mathbf{x}}_k^{t+1} = \text{prox}_{\eta g}(\mathbf{x}_k^t - \eta \nabla f(\mathbf{x}_k^t))$. Then the sequence $\{\mathbf{y}_k\}$ satisfies*

$$\mathbb{E}[F(\mathbf{y}_k) - F^*] \leq \left(\frac{d}{d+1}\right)^k (F(\mathbf{y}_0) - F^*), \quad (9)$$

where $d = \frac{c^2(L + \frac{1}{\eta})^2 + 2\eta L^2 m + \frac{1}{2\eta}}{\frac{1}{2\eta} - L - \alpha}$, and F^* is the optimal function value.

The convergence analysis for stochastic methods in inexact case has never been addressed before. To incorporate the KŁ property in deriving the convergence rate, we use a reference sequence generated by exact proximal mapping. Even though this sequence is not actually generated by the algorithm, we can reach to the convergence rate by analyzing the relation between the reference sequence and the actual sequence generated by the algorithm.

Compared to the exact case, the convergence rate remains at the same order, i.e., the linear convergence, but the convergence is slower due to the larger parameter d caused by the error parameter α .

4. Experiments

In this section, we compare the efficiency of APGnc and SVRG-APGnc with other competitive methods via numerical experiments. In particular, we focus on the non-negative principle component analysis (NN-PCA) problem, which can be formulated as

$$\min_{\mathbf{x} \geq 0} -\frac{1}{2} \mathbf{x}^T \left(\sum_{i=1}^n \mathbf{z}_i \mathbf{z}_i^T \right) \mathbf{x} + \gamma \|\mathbf{x}\|^2. \quad (10)$$

It can be equivalently written as

$$\min_{\mathbf{x}} -\frac{1}{2} \mathbf{x}^T \left(\sum_{i=1}^n \mathbf{z}_i \mathbf{z}_i^T \right) \mathbf{x} + \gamma \|\mathbf{x}\|^2 + \mathbf{1}_{\{\mathbf{x} \geq 0\}}. \quad (11)$$

Here, f corresponds to the first two terms, and g is the indicator of the nonnegative orthant, i.e., $\mathbf{1}_{\{\mathbf{x} \geq 0\}}$. This problem is nonconvex due to the negative sign and satisfies Assumption 1. In particular, it satisfies the KŁ property since it is quadratic.

For the experiment, we set $n = 2000, \gamma = 10^{-3}$ and randomly generate the samples \mathbf{z}_i from normal distribution.

All samples are then normalized to have unit norm. The initialization is randomly generated, and is applied to all the methods. We then compare the function values versus the number of effective passes through n samples.

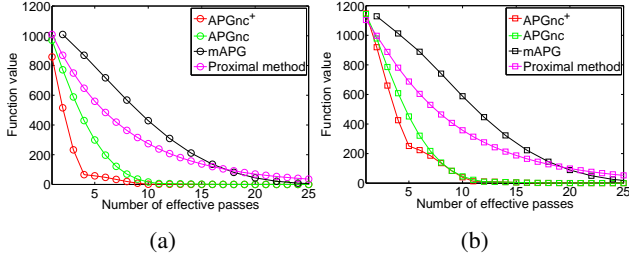


Figure 1. Performance comparison of APGnc+, APGnc, mAPG, and the traditional proximal method. (a) Error free. (b) There exists the proximal error.

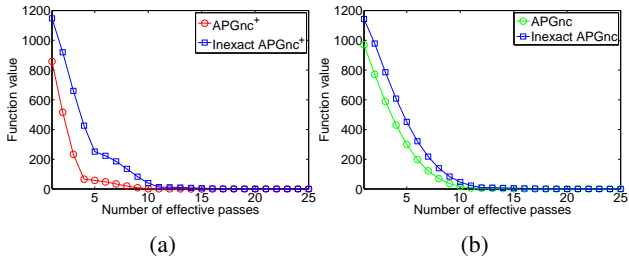


Figure 2. Performance comparison of the same algorithm with and without the proximal error. (a) Performance comparison of APGnc+. (b) Performance comparison of APGnc.

4.1. Comparison among APG variants

We first compare among the deterministic APG-like methods in Algorithms 2 - 4 and the standard proximal gradient method. The original APG in Algorithm 1 is not considered since it is not a descent method and does not have convergence guarantee in nonconvex cases. We tuned a fixed step size $\eta = 0.05/L$, where L is the spectral norm of the sample matrix $\sum_{i=1}^n \mathbf{z}_i \mathbf{z}_i^T$. We set $t = 1/2$ for APGnc+.

The results are shown in Figures 1 and 2. In Figure 1 (a), we show the performance comparison of the methods when there is no error in gradient or proximal calculation. One can see that APGnc and APGnc+ outperform all other APG variants. In particular, APGnc+ performs the best with our adaptive momentum strategy, justifying its empirical advantage. We note that the mAPG requires two passes over all samples at each iteration, and is, therefore, less data efficient compared to other APG variants. We further note that other choices of stepsize less than the standard choice $0.5/L$ does not change the relative comparison of the performance among the algorithms, and we observed that the adaptive momentum performs practically well for

$t \in (1/3, 2/3)$.

We further study the inexact case in Figure 1 (b), where we introduce the proximal error $\epsilon_k = \frac{1}{100k^3}$ at the k th iteration. One can see that inexact APGnc+ and inexact APGnc also outperform other two inexact algorithms. Furthermore, in Figure 2 (a) and (b), we compare exact and inexact algorithms respectively for APGnc+ and APGnc. It can be seen that even with a reasonable amount of inexactness, both methods converge comparably to their corresponding exact methods. Although initially the function value drops faster in exact algorithms, both exact and inexact algorithms converge to the optimal point almost at the same time. Such a fact demonstrates the robustness of the algorithms. We note that the relative comparison of the performance among the algorithms does not change under other choices like $\epsilon_k = 1/100k^2, 1/100k^4$.




Figure 3. Performance comparison of SVRG-APGnc+, SVRG-APGnc, and the traditional proximal SVRG. (a) Error free. (b) There exists the proximal error.




Figure 4. Performance comparison of the same algorithm with and without the proximal error. (a) Performance comparison of SVRG-APGnc+. (b) Performance comparison of SVRG-APGnc.

4.2. Comparison among SVRG-APG variants

We then compare the performance among SVRG-APGnc, SVRG-APGnc+ and the original proximal SVRG methods, and pick the stepsize $\eta = 1/8mL$ with $m = n$. The results are presented in Figures 3 and 4. In the error free case in Figure 3 (a), one can see that SVRG-APGnc+ method outperforms the others due to the adaptive momentum, and the SVRG-APGnc method also performs better than the original proximal SVRG method.

For the inexact case, we set the proximal error as $\epsilon_k =$

$\min(\frac{1}{100k^3}, 10^{-7})$, where 10^{-7} is chosen to suppress the large inexactness during the initial few iterations. One can see from Figure 3 (b) that the performance is degraded compared to the exact case, and converges to a different local minimum. In this result, all the methods are no longer monotone due to the inexactness and the stochastic nature of SVRG. Nevertheless, the SVRG-APGnc⁺ still yields the best performance.

We also compare the results corresponding to SVRG-APGnc⁺ and SVRG-APGnc, with and without the proximal error, in Figure 4 (a) and (b), respectively. It is clear that the SVRG-based algorithms are much more sensitive to the error comparing with APG-based ones. Even though the error is set to be smaller than in the inexact case with APG-based methods, one can observe more significant performance gaps than those in Figure 2.

5. Conclusion

In this paper, we provided comprehensive analysis of the convergence properties of APGnc as well as its inexact and stochastic variance reduced forms by exploiting the KŁ property. We also proposed an improved algorithm APGnc⁺ by adapting the momentum parameter. We showed that APGnc shares the same convergence guarantee and the same order of convergence rate as the mAPG, but is computationally more efficient and more amenable to adaptive momentum. In order to exploit the KŁ property for accelerated algorithms in the situations with inexact errors and/or with stochastic variance reduced gradients, we developed novel convergence analysis techniques, which can be useful for exploring other algorithms for nonconvex problems.

Acknowledgements

This work was supported in part by the NSF grant ECCS 1609916.

References

Attouch, H. and Bolte, J. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming*, 116(1-2):5–16, 2009. ISSN 0025-5610.

Attouch, H., Bolte, J., Redont, P., and Soubeyran, A. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Łojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.

Attouch, H., Bolte, J., and Svaiter, B. Convergence of descent methods for semi-algebraic and tame problems:

proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. *Mathematical Programming*, 137(1-2):91–129, 2013.

Beck, A. and Teboulle, M. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *Transactions on Image Processing*, 18(11):2419–2434, November 2009a.

Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Image Science.*, 2(1):183–202, 2009b.

Bolte, J., Daniilidis, A., and Lewis, A. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17:1205–1223, 2007.

Bolte, J., Daniilidis, A., Ley, O., and Mazet, L. Characterizations of Łojasiewicz inequalities and applications: Subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362(6):3319–3363, 2010.

Bolte, J., Sabach, S., and Teboulle, M. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.

Frankel, P., Garrigos, G., and Peypouquet, J. Splitting methods with variable metric for kurdyka-Łojasiewicz functions and general convergence rates. *Journal of Optimization Theory and Applications*, 165(3):874–900, 2015.

Ghadimi, S. and Lan, G. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming*, 156(1-2):59–99, 2016.

Gong, P., Zhang, C., Lu, Z., Huang, J. Z., and Ye, J. A general iterative shrinkage and thresholding algorithm for non-convex regularized optimization problems. In *International Conference on Machine Learning*, 2013.

Hardt, M. and Ma, T. Identity matters in deep learning. *Arxiv preprint*, 2016. URL <http://arxiv.org/abs/1611.04231>.

Johnson, R. and Zhang, T. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pp. 315–323, 2013.

Karimi, H., Nutini, J., and Schmidt, M. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. *Machine Learning and Knowledge Discovery in Databases: European Conference*, pp. 795–811, 2016.

- Kurdyka, K. On gradients of functions definable in o-minimal structures. *Annales de l'institut Fourier*, 48(3): 769–783, 1998.
- Kurdyka, K. and Spodzieja, S. *Separation of real algebraic sets and the Łojasiewicz exponent*. Wydział Matematyki Informatyki. Uniwersytet Łódzki, 2011. URL <https://books.google.com/books?id=CL6PMwEACAAJ>.
- Li, G. and Kei, T. Calculus of the exponent of Kurdyka-Łojasiewicz inequality and its applications to linear convergence of first-order methods. *ArXiv preprint*, February 2016. URL <https://arxiv.org/abs/1602.02915>.
- Li, H. and Lin, Z. Accelerated proximal gradient methods for nonconvex programming. In *Advances in Neural Information Processing Systems*, pp. 379–387. 2015.
- Łojasiewicz, S. *Ensembles semi-analytiques*. Institut des Hautes Etudes Scientifiques, 1965.
- Nesterov, Y. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27:372–376, 1983.
- Reddi, S., Hefny, A., Sra, S., Póczos, B., and Smola, A. Stochastic variance reduction for nonconvex optimization. *ArXiv preprint*, 2016a. URL <https://arxiv.org/abs/1603.06160>.
- Reddi, S., Sra, S., Póczos, B., and Smola, A. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. In *Advances in Neural Information Processing Systems 29*, pp. 1145–1153. 2016b.
- Rockafellar, R.T. and Wets, R.J.B. *Variational Analysis*. Springer, 1997.
- Schmidt, M., Roux, N.L., and Bach, F.R. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in Neural Information Processing Systems 24*, pp. 1458–1466. 2011.
- Tseng, P. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Mathematical Programming*, 125(2):263–295, 2010.
- Yao, Q. and Kwok, J.T. More efficient accelerated proximal algorithm for nonconvex problems. *ArXiv preprint*, December 2016. URL <https://arxiv.org/abs/1612.09069>.
- Zhou, Y., Zhang, H., and Liang, Y. Geometrical properties and accelerated gradient solvers of non-convex phase retrieval. *The 54th Annual Allerton Conference*, 2016.
- Zhu, Z. and Yuan, Y. Improved SVRG for non-strongly-convex or sum-of-non-convex objectives. In *International Conference on Machine Learning*, 2016.