

Convergence rate of incremental aggregated gradient algorithms

M. Gürbüzbalaban ^{*}, A. Ozdaglar ^{*}, P. Parrilo ^{*}

May 25, 2015

Abstract

Motivated by applications to distributed asynchronous optimization and large-scale data processing, we analyze the incremental aggregated gradient method for minimizing a sum of strongly convex functions from a novel perspective, simplifying the global convergence proofs considerably and proving a linear rate result. We also consider an aggregated method with momentum and show its linear convergence. We conclude by discussing extensions of our results to the problems with an additional convex, possibly non-smooth function.

1 Introduction

We consider the following unconstrained optimization problem where the objective function is the sum of component functions:

$$\begin{aligned} \text{minimize} \quad & f(x) = \sum_{i=1}^m f_i(x) \\ \text{subject to} \quad & x \in \mathbb{R}^n, \end{aligned} \tag{1}$$

where each $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex, continuously differentiable function where \mathbb{R}^n is the n -dimensional Euclidean space, referred to as a *component function*. This problem arises in many applications including least square problems [MYF03, Ber96] or more general parameter estimation problems where f_i is the corresponding loss function of the i -th data block [GOP14], distributed optimization in wireless sensor networks [BHG07], machine learning problems [RSB12, BLC05] and minimization of expected value of a function (where the expectation is taken over a finite probability distribution or approximated by an m -sample average).

^{*}Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02140, USA.

One widely studied approach is the (deterministic) incremental gradient (IG) method, which cycles through the component functions using a deterministic order and updates the iterates using the gradient of a single component function [Ber99]. This method can be faster than non-incremental methods since each step is relatively cheaper (one gradient computation instead of m gradient computations in the non-incremental case) and each step makes reasonable progress on average [Ber99]. However, IG requires the stepsize to go to zero to obtain convergence to the optimal solution of problem (1) even if it is applied to smooth and strongly convex component functions ([Ber11]) unless a restrictive “gradient growth condition” holds [Sol98]. As a consequence, with a decaying stepsize, IG has typical sublinear convergence rate properties. The same observation applies both to stochastic gradient methods (which uses a random order for cycling through the component functions) [SR13] and to incremental Newton methods that are of second-order [GOP14].

Another interesting class of methods includes the *incremental aggregated gradient* (IAG) method of Blatt *et al.* (see [BHG07, TY14]) and closely-related stochastic methods including the *stochastic average gradient* (SAG) method [RSB12], the SAGA method [DBLJ14] and the MISO method [Mai13]. These methods process a single component function at a time as in incremental methods, but keeps a memory of the most recent gradients of all component functions so that a full gradient step is taken at each iteration. They might require an excessive amount of memory when m is large, however they have linear convergence properties on strongly convex functions with constant stepsize without requiring the restrictive gradient growth condition. Furthermore, IAG forms approximations to the gradient of the objective function $\nabla f(x)$ at each step and this provides an accurate and efficient stopping criterion (stop if the norm of the approximate gradient is below a certain threshold) whereas it is often not clear “when to stop” with IG.

The IAG method was first proposed in a pioneer work by [BHG07] where its global convergence under some assumptions is shown. It is also shown that in the special case when each f_i is a quadratic, IAG exhibits global linear convergence if the stepsize is small enough, although neither a convergence rate result nor an explicit upper bound on the stepsize that can lead to linear convergence is given. This result is based on a perturbation analysis of the eigenvalues of a periodic dynamic linear system which is of independent interest in terms of the techniques used but is also highly technical and computationally demanding as it requires estimating the derivatives of the eigenvalues of a one-parameter matrix family. Furthermore, it only applies to quadratic functions. More recently, [TY14] proved global convergence under less restrictive conditions and local linear convergence in a more general setting when each component function satisfies a local Lipschitzian error condition, a condition satisfied by locally strongly convex functions (around an optimal solution). Although the results are more general than those of [BHG07] as they apply beyond quadratics, the proofs are still involved and do not contain any explicit rate estimates because (i) the constants involved in the analysis are implicit and hard to compute/approximate, (ii) the results are asymptotic (they hold when the stepsize is small enough but bounds on the stepsize are not available). See Remark 3.4 for more detail.

In this paper, we present a novel convergence analysis for the IAG method with several advantages and implications. First, our analysis is based on a careful choice of a Lyapunov

function which leads to simple global and linear convergence proofs. Furthermore, our proofs give more insight into the behavior of IAG compared to previous approaches, showing that IAG can be treated as a perturbed gradient descent method where gradient errors can be interpreted as shocks with a finite duration that are fading away as one gets closer to an optimal solution (in a way we will precise). Second, to our knowledge, our analysis is the first to provide explicit rate estimates for (deterministic) IAG methods. Third, we discuss an “IAG method with momentum” and show its global and linear convergence. To our knowledge, this is the first global convergence and linear convergence result for an aggregated gradient method with memory.

In problems where there is no favorable deterministic ordering of the component functions, stochastic gradient methods may have a better worst-case complexity as they do not favor one ordering than the other [Ber15], [Ber99, Example 1.5.6] avoiding some circular behavior that might arise in corner cases. Furthermore, the analysis of the convergence of stochastic methods are often given in the expected cost which is a weaker averaged sense of convergence compared to probability one convergence or deterministic convergence. As a consequence, it is expected that the resulting convergence rate results in expected cost for stochastic methods are no worse than the rate of deterministic incremental gradient methods. Indeed, the stochastic version of the IAG method when the order is stochastic, SAG, has a faster convergence rate than the rate we prove for IAG (see Theorem 3.3 and [RSB12]). Nevertheless, there are many applications when stochastic algorithms are not applicable and incremental algorithms are appropriate due to the structure of the problem. For example, in sensor networks or source localization problems, sensors are a part of a big network structure and are only able to communicate with their neighbors subject to certain constraints in terms of geography and distance. In these scenarios, randomized sampling of the order may not be possible or practical due to communication constraints requiring to follow a deterministic order (see e.g. [BHG07]) which motivates a study of the deterministic incremental methods such as IAG further. To our knowledge, our analysis gives the first explicit linear rate results for the deterministic IAG method. Furthermore, our analysis admits simple extensions to the generalization of (1) that contains an additive convex possibly non-smooth term (such as the indicator of a set). This generalization also allows one to handle convex constraints (see Section 5).

In the next section, Section 2, we describe the IAG method. Section 3 introduces the main assumptions and estimates for our convergence analysis and the linear rate result. In Section 4, we develop a new IAG method with momentum and provide a linear rate result for it. In Section 5, we conclude by discussing summary and future work.

Notation We use $\|\cdot\|$ to denote the standard Euclidean norm on \mathbb{R}^n . For a scalar x , we define $(x)_+ = \max(x, 0)$. The gradient and the Hessian matrix of f at a point $x \in \mathbb{R}^n$ are denoted by $\nabla f(x)$ and $\nabla^2 f(x)$ respectively.

2 IAG method

For a constant stepsize $\gamma > 0$, an integer $K \geq 0$ and arbitrary initial points $x^0, x^{-1}, \dots, x^{-K} \in \mathbb{R}^n$, the IAG method consists of the following iterations,

$$g^k = \sum_{i=1}^m \nabla f_i(x^{\tau_i^k}), \quad (2)$$

$$x^{k+1} = x^k - \gamma g^k, \quad k = 0, 1, 2, \dots, \quad (3)$$

where the gradient sampling times $\{\tau_i^k\}_{i=1}^m$ can be arbitrary as long as they are sampled at least once in the last K iterations, i.e.

$$k \geq \tau_i^k \geq k - K, \quad i = 1, 2, \dots, m. \quad (4)$$

In other words, K is an upper bound on the delay encountered by the gradients of the component functions. The update (2) determines the direction of motion $-g^k$ by approximating the steepest descent direction $-\nabla f(x^k)$ at (iterate) time k from the recently computed gradients of the component functions (at times $\{\tau_i^k\}_{i=1}^m$). For example, if the component functions are processed one by one using a deterministic cyclic order, we have $\tau_i^0 = 0$ for $i = 1, 2, \dots, m$ and the recursion

$$\tau_i^k = \begin{cases} k, & \text{if } i = (k)_m + 1 \\ \tau_i^{k-1}, & \text{else,} \end{cases}, \quad i = 1, 2, \dots, m, \quad k = 0, 1, \dots,$$

which satisfies $\tau_i^k \geq k - (m - 1)$ for all i, k where $K = m - 1$. This is the original IAG method introduced by Blatt *et al.*[BHG07]. Later on, Tseng *et al* [TY14] generalized Blatt *et al.*'s work allowing more general gradient sampling times $\{\tau_i^k\}$ with bounded delays, i.e. satisfying (4).

As IAG takes an approximate steepest gradient descent direction, it is natural to analyze it as a perturbed steepest gradient descent method. In fact, in the special case, when $K = 0$, all the gradients are up-to-date and IAG reduces to the classical (non-incremental) gradient descent (GD) method which is well known. Therefore, the more interesting case that we will analyze is when $K > 0$.

For simplicity of notation in our analysis, we will take

$$x^0 = x^{-1} = \dots = x^{-K} \quad (5)$$

but the analysis would extend to other (arbitrary) choices of initial points $\{x^j\}_{j=-K}^0$ in a straightforward manner.

3 Convergence analysis

3.1 Preliminaries

We will make the following assumptions that have appeared in a number of papers analyzing incremental methods including [Sol98], [TY14], [MYF03].

Assumption 3.1. (Strong convexity and Lipschitz gradients)

(i) Each f_i has Lipschitz continuous gradients on \mathbb{R}^n satisfying

$$\|\nabla f_i(y) - \nabla f_i(z)\| \leq L_i \|y - z\|, \quad \forall y, z \in \mathbb{R}^n,$$

where $L_i \geq 0$ is the Lipschitz constant, $i = 1, 2, \dots, m$. Let $L = \sum_{i=1}^m L_i$. It follows by the triangle inequality that f has Lipschitz continuous gradients with a Lipschitz constant L .

(ii) The function f is strongly convex on \mathbb{R}^n with parameter $\mu > 0$.

We define the condition number of f as

$$Q = \frac{L}{\mu} \geq 1 \tag{6}$$

where L is a Lipschitz constant for ∇f (see e.g. [Nes04])¹. As an example, in the special case when each f_i is a quadratic function, $\nabla^2 f_i(x)$ is a constant (matrix) for each i and we can take L_i to be its largest eigenvalue. The strong convexity constant c can be taken as the smallest eigenvalue of the Hessian of f .

A consequence of Assumption 3.1 on the strong convexity of f is that there exists a unique optimal solution of the problem (1) which we denote by x^* . Furthermore, we have the following well-known inequality

$$\langle \nabla f(x), x - x^* \rangle \geq \mu \|x - x^*\|^2, \quad \forall x \in \mathbb{R}^n, \tag{7}$$

(see [Nes04, Theorem 2.1.9]). By the gradient Lipschitzness, we also have

$$\|\nabla f(x)\| = \|\nabla f(x) - \nabla f(x^*)\| \leq L \|x - x^*\|, \quad \forall x \in \mathbb{R}^n. \tag{8}$$

We finally introduce the following lemma for proving the linear rate for IAG. We omit the proof due to space considerations, for a proof see Feyzmadvanian *et al.* [FAJ14] where this lemma is used to analyze the effect of delays in a first-order method. The intuition behind this lemma is the following: If a non-negative sequence $\{V_k\}$ that decays to zero linearly obeying $V_{k+1} \leq pV_k$ for some $p < 1$ is perturbed with an additive (noise) shock term that depends on the recent history, i.e. the shock at step k is on the order of V_ℓ where $\ell \in [k - d(k), k]$ and $d(k)$ is the time interval (duration) of the shock, the linear convergence property can be preserved if the shocks are decaying “fast” enough but this comes at the expense of a degraded rate $r > p$ which is determined by the “amplitude” of the shocks and the “duration” of the shocks.

Lemma 3.2. *Let $\{V_k\}$ be a sequence of real numbers satisfying*

$$V_{k+1} \leq pV_k + q \max_{k-d(k) \leq \ell \leq k} V_\ell, \quad k \geq 0,$$

¹The Lipschitz constant $L = \sum_i L_i$ of f might not be the best Lipschitz constant as Lipschitz constants L_i of f_i might be subadditive.

for some non-negative constants p and q . If $p + q < 1$ and

$$0 \leq d(k) \leq d_{\max}$$

for some $d_{\max} \geq 0$, then

$$V_k \leq r^k V_0, \quad k \geq 0,$$

where $r = (p + q)^{\frac{1}{1+d_{\max}}}$.

3.2 Bounding gradient error

We denote the distance to the optimal solution at iterate k by

$$\text{dist}_k = \|x^k - x^*\| \tag{9}$$

and the gradient error by

$$e^k = g^k - \nabla f(x^k). \tag{10}$$

We will show that the gradient error can be bounded in terms of a finite sum involving distances of iterates to the optimal solution. Using the triangle inequality and the Lipschitzness of the gradients, for any $k \geq 0$,

$$\begin{aligned} \|e^k\| &\leq \sum_{i=1}^m \|f_i(x^{\tau_i^k}) - f_i(x^k)\| \\ &\leq \sum_{i=1}^m L_i \|x^{\tau_i^k} - x^k\|. \end{aligned} \tag{11}$$

As the gradient delays are bounded by K (see (4)), by a repetitive application of the triangle inequality, we obtain for any $k \geq 0$,

$$\begin{aligned} \|e^k\| &\leq \sum_{i=1}^m L_i \sum_{j=\tau_i^k}^{k-1} \|x^{j+1} - x^j\| \\ &\leq L \sum_{j=(k-K)_+}^{k-1} \|x^{j+1} - x^j\| \end{aligned} \tag{12}$$

$$\begin{aligned} &= \gamma L \sum_{j=(k-K)_+}^{k-1} \|g^j\| \\ &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(\|\nabla f(x^j)\| + \|e^j\| \right) \end{aligned} \tag{13}$$

(with the convention that $\|e^0\| = 0$ which is implied by (5)). The inequality (13) provides a recursive upper bound for the gradient error that relates the gradient error $\|e^k\|$ to the

previous gradient errors $\|e^j\|$ with $j < k$. Bounding the $\|e^j\|$ term in (13) with the previously obtained bound (11) on the gradient error and by the triangle inequality,

$$\begin{aligned}
\|e^k\| &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(\|\nabla f(x^j)\| + \sum_{i=1}^m L_i \|x^{\tau_i^j} - x^j\| \right) \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(\|\nabla f(x^j)\| + \sum_{i=1}^m L_i (\|x^{\tau_i^j} - x^*\| + \|x^* - x^j\|) \right) \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(\|\nabla f(x^j)\| + \sum_{i=1}^m L_i \left(\max_{\ell \in \{\tau_i^j\}_{i=1}^m} \text{dist}_\ell + \text{dist}_j \right) \right).
\end{aligned} \tag{14}$$

Invoking (4) on the boundedness of the gradient delays by K once more and using (8) on gradient Lipschitzness to bound the norm of the gradient, we finally get

$$\begin{aligned}
\|e^k\| &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(\|\nabla f(x^j)\| + 2L \max_{(j-K)_+ \leq \ell \leq j} \text{dist}_\ell \right) \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(L \text{dist}_j + 2L \max_{(j-K)_+ \leq \ell \leq j} \text{dist}_\ell \right) \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \left(3L \max_{(j-K)_+ \leq \ell \leq j} \text{dist}_\ell \right) \\
&\leq 3\gamma L^2 K \max_{(k-2K)_+ \leq \ell \leq k-1} \text{dist}_\ell.
\end{aligned} \tag{15}$$

3.3 Linear convergence analysis

From the IAG update formulæ (2)–(3), it follows directly by taking norm squares that

$$\begin{aligned}
\text{dist}_{k+1}^2 &= \text{dist}_k^2 - 2\gamma \langle x^k - x^*, g^k \rangle + \gamma^2 \|g_k\|^2 \\
&= \text{dist}_k^2 - 2\gamma \langle x^k - x^*, \nabla f(x^k) \rangle - 2\gamma \langle x^k - x^*, e^k \rangle + \gamma^2 \|g_k\|^2 \\
&\leq (1 - 2\gamma\mu) \text{dist}_k^2 - 2\gamma \langle x^k - x^*, e^k \rangle + \gamma^2 \|g_k\|^2
\end{aligned}$$

where in the last step we used the inequality (7) due to the strong convexity of f . Plugging the identity $\|g^k\|^2 = \|\nabla f(x^k) + e^k\|^2$ and using the Cauchy-Schwartz inequality we obtain

$$\begin{aligned}
\text{dist}_{k+1}^2 &\leq (1 - 2\gamma\mu) \text{dist}_k^2 - 2\gamma \langle x^k - x^*, e^k \rangle + \gamma^2 \|\nabla f(x^k) + e^k\|^2 \\
&\leq (1 - 2\gamma\mu) \text{dist}_k^2 + 2\gamma \text{dist}_k \|e^k\| + 2\gamma^2 \|\nabla f(x^k)\|^2 + 2\gamma^2 \|e^k\|^2 \\
&\leq (1 - 2\gamma\mu + 2\gamma^2 L^2) \text{dist}_k^2 + 2\gamma \text{dist}_k \|e^k\| + 2\gamma^2 \|e^k\|^2
\end{aligned} \tag{16}$$

where we used (8) again to bound the norm of the gradient in the last step. Note that when $K = 0$, we have $e^k = 0$ and IAG reduces to the classical gradient descent method which

is well-known to be globally linearly convergent when a constant stepsize small enough is used. This is also apparent from (16). The next theorem shows that in the more interesting general case when there are gradient errors, i.e. when $K > 0$ and $e^k \neq 0$, a similar linear convergence argument can be done. The main idea is to eliminate the gradient error $\|e^k\|$ terms in (16) by replacing them with terms involving only distances. This can be done by invoking (15) which essentially provides an upper bound for the gradient errors in terms of distances. Then, Lemma 3.2 with $V_k = \text{dist}_k^2$ applies and provides the convergence rate.

Theorem 3.3. *Suppose that Assumption 3.1 holds. Consider the IAG iterations (2)–(3) with $K > 0$ and constant stepsize $0 < \gamma < 2\bar{\gamma}$ where*

$$\bar{\gamma} = \left(\frac{1}{11K + 2} \right) \frac{\mu}{L^2}.$$

Then, IAG iterates $\{x^k\}$ are globally linearly convergent. Furthermore, when $\gamma = \bar{\gamma}$ we have

$$\|x^k - x^*\| \leq \left(1 - \frac{c_K}{Q^2}\right)^k \|x^0 - x^*\|, \quad \text{for } k = 0, 1, 2, \dots,$$

where $c_K = \frac{1}{2}[(11K + 2)(2K + 1)]^{-1}$.

Proof. Combining the upper bound (15) on the gradient error and the recursive inequality (16) on the distance to the optimal solution we obtain

$$\begin{aligned} \text{dist}_{k+1}^2 &\leq (1 - 2\gamma\mu + 2\gamma^2L^2)\text{dist}_k^2 + 6\gamma^2L^2K \max_{(k-2K)_+ \leq \ell \leq k-1} \text{dist}_\ell \text{dist}_k \\ &\quad + 18\gamma^4L^4K^2 \left(\max_{(k-2K)_+ \leq \ell \leq k-1} \text{dist}_\ell \right)^2 \\ &\leq (1 - 2\gamma\mu + 2\gamma^2L^2)\text{dist}_k^2 + 6\gamma^2L^2K \max_{(k-2K)_+ \leq \ell \leq k-1} \left(\frac{\text{dist}_\ell^2 + \text{dist}_k^2}{2} \right) \\ &\quad + 18\gamma^4L^4K^2 \max_{(k-2K)_+ \leq \ell \leq k-1} \text{dist}_\ell^2 \\ &\leq p(\gamma) \text{dist}_k^2 + q(\gamma) \max_{(k-2K)_+ \leq \ell \leq k} \text{dist}_\ell^2 \end{aligned}$$

with

$$p(\gamma) := 1 - 2\gamma\mu + 2\gamma^2L^2, \quad q(\gamma) := 6\gamma^2L^2K + 18\gamma^4L^4K^2. \quad (17)$$

Assume $0 < \gamma < 2\bar{\gamma}$. Using the fact that $Q \geq 1$, it is a straightforward computation to check that $18\gamma^4L^4K^2 \leq 1$ and this can be used to show that

$$s(\gamma) := p(\gamma) + q(\gamma) \leq 1 - 2\gamma\mu + \gamma^2L^2(11K + 2) < 1. \quad (18)$$

Then by Lemma 3.2, we have global linear convergence of the sequence $V_k = \text{dist}_k^2$ to zero with rate $\rho(\gamma) = s(\gamma)^{1/(2K+1)} < 1$. This shows the global linear convergence of IAG. It remains

to show the claimed convergence rate for $\gamma = \bar{\gamma}$. Let $\gamma = \bar{\gamma}$. Note that this minimizes the quadratic with respect to γ in (18) leading to

$$s(\bar{\gamma}) \leq 1 - \frac{1}{Q^2(11K + 2)}. \quad (19)$$

Then, as the linear convergence rate of the sequence $\{\text{dist}_k^2\}$ is $\rho(\bar{\gamma}) < 1$, by taking square roots, the sequence $\{\text{dist}_k\}$ is linearly convergent with rate $\bar{\rho} = \rho(\bar{\gamma})^{1/2}$ satisfying

$$\text{dist}_k \leq \bar{\rho}^k \text{dist}_0, \quad \bar{\rho} = \rho(\bar{\gamma})^{1/2} = (s(\bar{\gamma})^{1/(2K+1)})^{1/2} \leq 1 - \frac{c_K}{Q^2},$$

where we used (19) and the inequality $(1 - x)^a \leq 1 - ax$ for $x, a \in [0, 1]$ to get an upper bound for $\bar{\rho}$. This completes the proof. \square

Remark 3.4. (Comparison with previous results) In a related work, Tseng et al. shows the existence of a positive constant C that if the stepsize γ is small enough then IAG has a K -step linear convergence rate of $\sqrt{1 - C\gamma}$ under a Lipschitzian error assumption which is a strong-convexity-like condition [TY14, Theorem 6.1]. However, in their analysis, there is no explicit rate estimate; as (i) results are asymptotic, holding for γ small enough without giving a precise interval for γ , (ii) the constant C is implicit and hard to compute/approximate as it depends on several other implicit constants and a Lipschitzian error parameter τ . Our analysis is not only simple using basic distance inequalities but also the constants are transparent and explicit.

Remark 3.5. (IAG versus IG) Theorem 3.3 shows that IAG with constant stepsize is globally linearly convergent, however the same is not true for IG. In fact, IG with constant stepsize is linearly convergent to a neighborhood of the solution but does not in general converge to the optimal solution due to the existence of gradient errors that are typically bounded away from zero [Ber99]. As a consequence, achieving global convergence with IG requires to use a stepsize that goes to zero and this results in typical slow convergence [Ber11]. In contrast, the gradient error in IAG is controlled by the distance of recent iterates to the optimal solution. The intuition is that the error fades away eventually as the iterates come closer to the optimal solution, therefore a diminishing stepsize is not needed to control the error.

Remark 3.6. (Local strong convexity implies local linear rate) We note that when f is not globally convex but only locally strongly convex around a stationary point², by a reasoning along the lines of the proof of Theorem 3.3, it is possible to show that IAG is locally linearly convergent.

4 IAG with momentum

An important variant of the GD method is the heavy-ball method [Pol87] which extrapolates the direction implied by the previous two iterates by the following update rule:

$$x^{k+1} = x^k - \gamma \nabla f(x^k) + \beta(x^k - x^{k-1})$$

²such as when the Hessian is not degenerate around a locally optimal solution

where $\beta \geq 0$ is the *momentum* parameter. It can be shown that the heavy-ball method can achieve a faster local convergence than GD when β is in a certain range [Pol87, Section 3.1]. There has also been much interest in understanding its global convergence properties [LRP14, GFJ14]. Accelerated gradient methods introduced by Nesterov [Nes83, N⁺07, Bub14] can also be thought of as *momentum* methods where the momentum parameter is variable and appropriately chosen. There has been a lot of recent interest in these accelerated methods as they have optimal iteration complexity properties under some conditions [Nes04].

In contrast to the recent advances in non-incremental methods with momentum, there has been less progress on incremental methods with momentum. In particular, no deterministic incremental methods with favorable convergence characteristics similar to those of accelerated gradient methods are currently known. However, there is the *IG method with momentum* which consists of the *inner* iterations

$$x_i^{k+1} = x_i^k - \gamma^k \nabla f_i(x^k) + \beta(x^k - x^{k-1}), \quad i = 1, 2, \dots, m \quad k \geq 1,$$

starting from $x_1^1 \in \mathbb{R}^n$ with the convention that $x_1^{k+1} = x_{m+1}^k$ where γ^k is the stepsize [Ber96, MS94, Tse98]. This method can be faster than IG on some problems [Tse98] especially when gradients have oscillatory behavior, however it would still require the stepsize go to zero due to gradient errors, leading to typical sublinear convergence. It is natural to ask whether IAG with such an additional momentum term, which we abbreviate by *IAG-M*,

$$x^{k+1} = x^k - \gamma g^k + \beta(x^k - x^{k-1}), \quad k = 0, 1, 2, \dots, \quad (20)$$

would be globally convergent and if it is, whether it could be faster than IAG in some practical problems. The global linear convergence of the IAG-M method for a certain range of β values can be shown by a similar reasoning along the lines presented in Section 3. Most of the logic in the derivation of the inequalities (11)–(15) apply with the only difference that the $\|x^{j+1} - x^j\|$ terms will now contain an additional momentum term due to the modified update rule (20). We will however provide a sketch of the proof for the sake of completeness. Using the gradient error bound (12) and iterate update equation (20) of IAG-M,

$$\begin{aligned} \|e^k\| = \|g^k - \nabla f(x^k)\| &\leq L \sum_{j=(k-K)_+}^{k-1} \|x^{j+1} - x^j\| = L \sum_{j=(k-K)_+}^{k-1} \|\gamma g^j + \beta(x^j - x^{j-1})\| \\ &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} \|g^j\| + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\| \\ &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} (\|\nabla f(x^j)\| + \|e^j\|) + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\|. \end{aligned}$$

Then, using (8) to bound the norm of the gradient and (12) to bound $\|e^j\|$, this becomes

$$\begin{aligned}
\|e^k\| &\leq \gamma L \sum_{j=(k-K)_+}^{k-1} (L \text{dist}_j + \|e^j\|) + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\| \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} (L \text{dist}_j + L \sum_{\ell=(j-K)_+}^{j-1} \|x^{\ell+1} - x^\ell\|) + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\| \\
&\leq \gamma L \sum_{j=(k-K)_+}^{k-1} (L \text{dist}_j + L \sum_{\ell=(j-K)_+}^{j-1} (\text{dist}_\ell + \text{dist}_{\ell+1})) + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\| \\
&\leq (3\gamma L^2 K) \max_{(k-2K)_+ \leq \ell \leq k-1} \text{dist}_\ell + \beta L \sum_{j=(k-K)_+}^{k-1} \|x^j - x^{j-1}\|. \tag{21}
\end{aligned}$$

Note that when $\beta = 0$, this inequality reduces to (15) obtained for IAG. From the inner update equation (20), we also have

$$\begin{aligned}
\|x^j - x^{j-1}\| &\leq \gamma \|g^{j-1}\| + \beta \|x^{j-1} - x^{j-2}\| \\
&\leq \gamma \|g^{j-1}\| + \beta (\|x^{j-1} - x^*\| + \|x^{j-2} - x^*\|) \\
&\leq \gamma \|\nabla f(x^{j-1})\| + \gamma \|e^{j-1}\| + \beta (\text{dist}_{j-1} + \text{dist}_{j-2}) \\
&\leq \gamma L \text{dist}_{j-1} + \gamma \|e^{j-1}\| + 2\beta \max(\text{dist}_{j-1}, \text{dist}_{j-2}) \tag{22}
\end{aligned}$$

where we used (8) to bound the norm of the gradient in the last inequality. We next bound the gradient error term $\|e^{j-1}\|$ on the right-hand side. A consequence of (11), the triangle inequality and the boundedness of the gradient delays is that gradient error is bounded by

$$\|e^k\| \leq \sum_{i=1}^m L_i (\text{dist}_{\tau_i^k} + \text{dist}_k) \leq 2L \max_{(k-K)_+ \leq \ell \leq k} \text{dist}_\ell, \quad k \geq 0. \tag{23}$$

Combining all the inequalities (21), (22) and (23) together, leads to

$$\|e^k\| \leq \left(3\gamma L^2 K + \beta L K (3\gamma L + 2\beta) \right) \max_{(k-2K-1)_+ \leq \ell \leq k-1} \text{dist}_\ell, \tag{24}$$

which is an analogue of (15) for IAG-M. Then, following the same line of argument with the proof of Theorem 3.3, it is straightforward to show a linear rate for IAG-M as long as the momentum parameter β is not very large. More specifically, it follows that when $\beta = O(\gamma^p)$ as $\gamma \rightarrow 0$ with $p \geq 1/2$, IAG-M is globally linearly convergent. We omit the details for the sake of brevity and leave it to the reader.

We also leave the numerical experiments comparing the performance of IAG-M and IAG in detail as a future work.

5 Summary and future work

We analyzed the IAG method when component functions are strongly convex by viewing it as a gradient descent method with errors. To the best of our knowledge, our analysis provides the first linear rate result. Furthermore, it is different than the existing two approaches [BHG07] and [TY14] in the sense that (i) it is based on simple basic inequalities that makes global convergence analysis simpler, (ii) gives more insight into the behavior of IAG. In particular, our analysis shows that the gradient errors can be treated as shocks with a finite duration which can be bounded in terms of distance of iterates to the optimal solution. Therefore, by choosing the stepsize small enough and using the strong convexity properties we can guarantee that the the distance to the optimal solution shrinks down at each step by a factor less than one.

We also developed a new algorithm, IAG with momentum, and provided a linear convergence and rate analysis. We expect that this algorithm can outperform IAG in problems where the individual gradients show oscillatory behavior because the momentum term provides an extra smoothing/averaging affect on the iterates. We leave numerical experiments for comparing the performance of IAG and IAG-M as a future work.

We note that the extension of IAG to the generalized version of (1)

$$\begin{aligned} \text{minimize} \quad & f(x) = \sum_{i=1}^m f_i(x) + h(x) \\ \text{subject to} \quad & x \in \mathbb{R}^n, \end{aligned}$$

with $h : \mathbb{R}^n \rightarrow \mathbb{R}$ convex and possibly non-smooth (such as the indicator of a function when there are constraints) is simple by an additional (proximal) step, see [TY14]. Our linear rate results would easily extend to this case in a straightforward manner due to the fact that the proximal steps are non-expansive in distances.

References

- [Ber96] D. Bertsekas. Incremental least squares methods and the extended Kalman filter. SIAM Journal on Optimization, 6(3):807–822, 1996.
- [Ber99] D. Bertsekas. Nonlinear programming. Athena Scientific, 1999.
- [Ber11] D. Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: a survey. Optimization for Machine Learning, 2010:1–38, 2011.
- [Ber15] D. Bertsekas. Convex Optimization Algorithms. Athena Scientific, 2015.
- [BHG07] D. Blatt, A. Hero, and H. Gauchman. A convergent incremental gradient method with a constant step size. SIAM Journal on Optimization, 18(1):29–51, 2007.
- [BLC05] L. Bottou and Y. Le Cun. On-line learning for very large data sets. Applied Stochastic Models in Business and Industry, 21(2):137–151, 2005.

- [Bub14] S. Bubeck. Theory of Convex Optimization for Machine Learning. ArXiv e-prints, May 2014.
- [DBLJ14] A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA:A fast incremental gradient method with support for non-strongly convex composite objectives. arXiv preprint arXiv:1407.0202, 2014.
- [FAJ14] Hamid Reza Feysmahdavian, Arda Aytekin, and Mikael Johansson. A delayed proximal gradient method with linear convergence rate. In Machine Learning for Signal Processing (MLSP), 2014 IEEE International Workshop on, pages 1–6. IEEE, 2014.
- [GFJ14] Euhanna Ghadimi, Hamid R. Feysmahdavian, and Mikael Johansson. Global convergence of the heavy-ball method for convex optimization, December 2014.
- [GOP14] M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. A globally convergent incremental Newton method. arXiv preprint arXiv:1410.5284, October 2014.
- [LRP14] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints, December 2014.
- [Mai13] Julien Mairal. Optimization with First-Order Surrogate Functions. In ICML, volume 28 of JMLR Proceedings, pages 783–791, Atlanta, United States, 2013.
- [MS94] O. L. Mangasarian and M. V. Solodov. Serial and parallel backpropagation convergence via nonmonotone perturbed minimization. Optimization Methods and Software, 4(2):103–116, January 1994.
- [MYF03] H. Moriyama, N. Yamashita, and M. Fukushima. The incremental Gauss-Newton algorithm with adaptive stepsize rule. Computational Optimization and Applications, 26(2):107–141, 2003.
- [N⁺07] Yurii Nesterov et al. Gradient methods for minimizing composite objective function, 2007.
- [NBB01] A Nedić, Dimitri P Bertsekas, and Vivek S Borkar. Distributed asynchronous incremental subgradient methods. Studies in Computational Mathematics, 8:381–407, 2001.
- [Nes83] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. In Soviet Mathematics Doklady, volume 27, pages 372–376, 1983.
- [Nes04] Y. Nesterov. Introductory lectures on convex optimization: a basic course. Applied Optimization. Springer, Boston, 2004.
- [Pol87] B. T. Polyak. Introduction to optimization. Translations series in mathematics and engineering. Optimization Software, Publications Division, New York, 1987.

- [RSB12] N. L. Roux, M. Schmidt, and F.R. Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems 25, pages 2663–2671. Curran Associates, Inc., 2012.
- [SCS14] M. Schmidt, A. Clifton, and A. Sarkar. Non-uniform stochastic average gradient method for training conditional random fields. NIPS OPT, 2014.
- [Sol98] M.V. Solodov. Incremental gradient algorithms with stepsizes bounded away from zero. Computational Optimization and Applications, 11(1):23–35, 1998.
- [SR13] M. Schmidt and N.L. Roux. Fast convergence of stochastic gradient descent under a strong growth condition. arXiv preprint arXiv:1308.6370, 2013.
- [Tse98] Paul Tseng. An incremental Gradient(-Projection) method with momentum term and adaptive stepsize rule. SIAM Journal on Optimization, 8(2):506–531, May 1998.
- [TY14] P. Tseng and S. Yun. Incrementally updated gradient methods for constrained and regularized optimization. Journal of Optimization Theory and Applications, 160(3):832–853, 2014.