

Convergence and Stability of the Stochastic Proximal Point Algorithm with Momentum

Junhyung Lyle Kim

Department of Computer Science, Rice University

JLYLEKIM@RICE.EDU

Panos Toulis

Booth School of Business, University of Chicago

PANOS.TOULIS@CHICAGOBOOTH.EDU

Anastasios Kyrillidis

Department of Computer Science, Rice University

ANASTASIOS@RICE.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Stochastic gradient descent with momentum (SGDM) is the dominant algorithm in many optimization scenarios, including convex optimization instances and non-convex neural network training. Yet, in the stochastic setting, momentum interferes with gradient noise, often leading to specific step size and momentum choices in order to guarantee convergence, set aside acceleration. Proximal point methods, on the other hand, have gained much attention due to their numerical stability and elasticity against imperfect tuning. Their stochastic accelerated variants though have received limited attention: how momentum interacts with the stability of (stochastic) proximal point methods remains largely unstudied. To address this, we focus on the convergence and stability of the stochastic proximal point algorithm with momentum (SPPAM), and show that SPPAM allows a faster linear convergence to a neighborhood compared to stochastic proximal point algorithm (SPPA) with a better contraction factor, under proper hyperparameter tuning. In terms of stability, we show that SPPAM depends on problem constants more favorably than SGDM, allowing a wider range of step size and momentum that lead to convergence.

Keywords: empirical risk minimization, stochastic proximal point algorithm, momentum, stability.

1. Introduction

Background. We focus on unconstrained empirical risk minimization instances (Robbins and Monro, 1951; Polyak and Juditsky, 1992; Bottou, 2012; Bottou and Bousquet, 2007; Shalev-Shwartz et al., 2011; Nemirovski et al., 2009; Moulines and Bach, 2011; Bach and Moulines, 2013), as in:

$$\min_{x \in \mathbb{R}^p} f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x). \quad (1)$$

To solve (1), stochastic gradient descent (SGD) has become the de facto method used by the machine learning community, mainly due to its computational efficiency (Zhang, 2004; Bottou, 2012; Bottou et al., 2018). For completeness, SGD iterates as follows:

$$x_{t+1} = x_t - \eta \nabla f_{i_t}(x_t), \quad (2)$$

where η is the step size, and ∇f_i is the (stochastic) gradient computed at the i -th data point.

Properties of SGD and Its Momentum Extension. While computationally efficient, stochastic methods often suffer from two major limitations: (i) slow convergence, and (ii) numerical instability. Due to the gradient noise, SGD could take longer to converge in terms of iterations (Moulines and Bach, 2011; Gower et al., 2019). Moreover, SGD suffers from numerical instabilities both in theory (Nemirovski et al., 2009) and practice (Bottou, 2012), allowing only a small range of η values that lead to convergence, which often depend on unknown quantities (Moulines and Bach, 2011).

With respect to slow convergence, many variants of accelerated methods have been proposed, along with analyses (Su et al., 2014; Defazio, 2019; Laborde and Oberman, 2020; Allen-Zhu and Orecchia, 2017; Lessard et al., 2016; Hu and Lessard, 2017; Wibisono et al., 2016; Bubeck et al., 2015). Most notable cases include the Polyak’s momentum (Polyak, 1964, 1987) and the Nesterov’s acceleration methods (Nesterov et al., 2018; Ahn and Sra, 2022; Nesterov, 1983). These methods allow faster (sometimes optimal) convergence rates, while having virtually the same computational cost as SGD. In particular, SGD with momentum (SGDM) (Polyak, 1964, 1987) iterates as follows:

$$x_{t+1} = x_t - \eta \nabla f_{i_t}(x_t) + \beta(x_t - x_{t-1}), \quad (3)$$

where $\beta \in [0, 1)$ is the momentum parameter. The intuition is that, if the direction from x_{t-1} to x_t was “correct,” then SGDM utilizes this inertia weighted by the momentum parameter β , instead of just relying on the current point x_t . Much of the state-of-the-art performance has been achieved with SGDM (Huang et al., 2017; Howard et al., 2017; He et al., 2016).

Yet, SGDM could be hard to tune: SGDM adds another hyperparameter—momentum β —to an already sensitive stochastic procedure of SGD. As such, various works have found that such motions could aggravate the instability of SGD. For instance, Liu and Belkin (2018) and Kidambi et al. (2018) show that accelerated SGD does not in general provide any acceleration over SGD, regardless of careful tuning; further, accelerated SGD may diverge for the step sizes that SGD converges. Assran and Rabbat (2020) also show that, even with finite-sum of quadratic functions, accelerated SGD may diverge under the usual choices of step size and momentum. See also Loizou and Richtárik (2020); Devolder et al. (2014); d’Aspremont (2008) for more discussions on this topic.

Stability via Proximal Updates. With respect to numerical stability, variants of SGD that utilize proximal updates have recently been proposed (Ryu and Boyd, 2017; Toulis et al., 2014; Toulis and Airoldi, 2017; Toulis et al., 2021; Asi and Duchi, 2019; Asi et al., 2020). In particular, Toulis et al. (2021) introduced stochastic errors in proximal point algorithms (SPPA) and analyzed its convergence and stability, which iterates similar to:

$$x_{t+1} = x_t - \eta (\nabla f(x_{t+1}) + \varepsilon_{t+1}). \quad (4)$$

Without stochastic errors, (4) is known as the proximal point algorithm (PPA) (Rockafellar, 1976; Güler, 1991) or the implicit gradient descent (IGD). PPA/IGD is known to converge with minimal assumptions on hyperparameter tuning, by improving the conditioning of the optimization problem; more details in Section 2. In the stochastic setting, Toulis et al. (2021) show that SPPA enjoys an exponential discount of the initial condition, regardless of the step size η and the smoothness parameter L . On the contrary, for SGD, both η and L show up within an exponential term, amplifying the initial conditions, leading to even divergence if misspecified (Moulines and Bach, 2011).

Our Focus and Contributions. Stochastic accelerated variants of PPA have received limited attention: how momentum interacts with the stability that PPA provides remains unstudied. To the best of our knowledge, *no momentum has been considered for stochastic proximal point updates that, beyond convergence, also studies the stability of the acceleration motions.* This is the aim of this work. Our contributions are summarized as:

- We introduce stochastic PPA with momentum (SPPAM), and study its convergence and stability behavior. SPPAM directly incorporates the momentum term akin to (3) into (4):

$$x_{t+1} = x_t - \eta (\nabla f(x_{t+1}) + \varepsilon_{t+1}) + \beta(x_t - x_{t-1}). \quad (5)$$

We study whether adding momentum β results in faster convergence akin to SGDM, while preserving the numerical stability, inherited by utilizing proximal updates akin to SPPA.

- We show that SPPAM enjoys linear convergence to a neighborhood (Theorem 9) with a better contraction factor than SPPA (Lemma 6). We further characterize the conditions on η and β that result in acceleration (Corollary 7). Finally, we characterize the condition that leads to the exponential discount of initial conditions for SPPAM (Theorem 10), which is significantly easier to satisfy compared to SGDM.
- Empirically, we confirm our theory with experiments on generalized linear models (GLM), including linear and Poisson regressions with different condition numbers. As expected, SGD and SGDM converge only for specific choices of η and β , while SPPA converges for a much wider range of η . SPPAM enjoys the advantages of both acceleration from the momentum and stability from the proximal step: it converges for the range of η that SPPA converges but with faster rate, which improves or matches that of SGDM, when the latter converges.

2. Preliminaries

Proximal Point Algorithm (PPA). The proximal point algorithm (PPA) (Rockafellar, 1976; Güler, 1991) obtains the next iterate for minimizing f by solving the following optimization problem:

$$x_{t+1} = \arg \min_{x \in \mathbb{R}^p} \left\{ f(x) + \frac{1}{2\eta} \|x - x_t\|_2^2 \right\}, \quad (6)$$

which is equivalent to implicit gradient descent (IGD) by the first-order optimality condition:

$$x_{t+1} = x_t - \eta \nabla f(x_{t+1}). \quad (7)$$

In words, instead of minimizing f directly, PPA minimizes f with an additional quadratic term. This small change brings a major advantage to PPA: if f is convex, the added quadratic term can make the problem strongly convex; if f is non-convex, PPA can make it convex (Ahn and Sra, 2022). Thanks to this conditioning property, PPA exhibits different behavior compared to GD in the deterministic setting. Güler (1991) proved that for a convex function f , PPA satisfies:

$$f(x_T) - f(x^*) \leq O\left(\frac{1}{\sum_{t=1}^T \eta_t}\right), \quad (8)$$

after T iterations. By setting the step size η_t to be large, PPA can converge “arbitrarily” fast.

PPA was soon considered in the stochastic setting. In Ryu and Boyd (2017), a stochastic version of PPA, dubbed as stochastic proximal iterations (SPI), was analyzed, where an approximation of f using a single data f_i was considered. The same algorithm was (statistically) analyzed under the name of implicit stochastic gradient descent (ISGD) (Toulis et al., 2014; Toulis and Airolidi, 2017), and was extended to the Robbins-Monro procedure in Toulis et al. (2021). Similar algorithms were analyzed recently in Asi and Duchi (2019); Asi et al. (2020) where each f_i was further approximated by simpler surrogate functions. These works generally indicate that, in the asymptotic regime, SGD and SPI/ISGD have the same convergence behavior; but in the non-asymptotic regime, SPI/ISGD outperforms SGD due to numerical stability provided by utilizing proximal updates.

Table 1: Comparison of different algorithms in Section 2. $(\cdot)^\alpha$ is Rockafellar (1976); Güler (1991); $(\cdot)^\beta$ is Güler (1992); Lin et al. (2015, 2018); $(\cdot)^\gamma$ is Toulis et al. (2014); Toulis and Airoldi (2017); Ryu and Boyd (2017); $(\cdot)^\delta$ is Asi and Duchi (2019); Asi et al. (2020); $(\cdot)^\epsilon$ is Kulunchakov and Mairal (2019); $(\cdot)^\zeta$ is Chadha et al. (2021); $(\cdot)^\eta$ is Deng and Gao (2021). We highlight with color the algorithms that include momentum motions.

Method	Deterministic
PPA/IGD ^{α}	$x_{t+1} = \arg \min_x \left\{ f(x) + \frac{1}{2\eta_t} \ x - x_t\ _2^2 \right\}$ $\Leftrightarrow x_{t+1} = x_t - \eta_t \nabla f(x_{t+1})$
Acc. PPA/Catalyst ^{β}	$x_{t+1} \approx \arg \min_x \left\{ f(x) + \frac{\kappa}{2} \ x - y_t\ _2^2 \right\}$ $y_t = x_t + \beta_t(x_t - x_{t-1})$ where $\alpha_t^2 = (1 - \alpha_t)\alpha_{t-1}^2 + \frac{\mu}{\mu + \kappa}\alpha_t$, $\beta_t = \frac{\alpha_{t-1}(1 - \alpha_{t-1})}{\alpha_{t-1}^2 + \alpha_t}$
Stochastic	
SPI/ISGD ^{γ}	$x_{t+1} = \arg \min_x \left\{ f_{i_t}(x) + \frac{1}{2\eta_t} \ x - x_t\ _2^2 \right\}$ $\Leftrightarrow x_{t+1} = x_t - \eta_t \nabla f_{i_t}(x_{t+1})$
APROX ^{δ}	Set $f_{i_t}(x) := \max \{ f_{i_t}(x_t) + \langle \nabla f_{i_t}(x_t), x - x_t \rangle, \inf_z f_{i_t}(z) \}$ from SPI
Stochastic Catalyst ^{ϵ}	Set $f(x) := f(y_t) + \langle g_t, x - y_t \rangle + \frac{\kappa + \mu}{2} \ x - y_t\ _2^2$ from Catalyst
Acc. APROX ^{ζ}	$y_t = (1 - \beta_t)x_t + \beta_t z_t$ $z_t = \arg \min_x \left\{ f_{i_t}(x) + \frac{1}{\eta_t} \ x - z_t\ _2^2 \right\}$ $x_{t+1} = (1 - \beta_t)x_t + \beta_t z_{t+1}$ where $f_{i_t}(x) := \max \{ f_{i_t}(x) + \langle \nabla f_{i_t}(x), y - x \rangle, \inf_z f_{i_t}(z) \}$
SEMOD ^{η}	$y_t = x_t + \beta(x_t - x_{t-1})$ $x_{t+1} = \arg \min_x \left\{ f_{i_t}(x) + \frac{1}{2\eta_t} \ x - y_t\ _2^2 \right\}$
SPPAM (this work)	$x_{t+1} = x_t - \eta (\nabla f(x_{t+1}) + \varepsilon_{t+1}) + \beta(x_t - x_{t-1})$

Accelerated PPA. Accelerated PPA was first proposed in deterministic setting in Güler (1992), where Nesterov’s acceleration was applied *after* solving the proximal step in (6). This yields the convergence rate of the form:

$$f(x_T) - f(x^*) \leq O \left(\frac{1}{(\sum_{t=1}^T \sqrt{\eta_t})^2} \right), \quad (9)$$

which is faster than the rate in (8). This bound is based on Nesterov’s momentum schedules, but does not study the effect in stability different tuning pairs (η, β) might have. Moreover, as can be seen in (8), PPA can already achieve arbitrarily fast convergence, given it is implemented exactly. Hence, following works focused on studying the conditions under which the proximal step in (6) can be computed inexactly, while still exhibiting some acceleration (Lin et al., 2015, 2018); similar analyses were later extended to the stochastic setting in Kulunchakov and Mairal (2019).

Chadha et al. (2021) and Deng and Gao (2021) also considered accelerated SPPA. Both of these works apply a convoluted 2- or 3-step Nesterov’s procedure after the proximal step, where f_i was further approximated with auxiliary functions. Yet, stability arguments via proximal updates are

less apparent due to the auxiliary functions, requiring specific step size and momentum schedules, which might involve an additional one-dimensional optimization per iteration; see also Theorem 10. A summary of these algorithms is provided in Table 1.

Intuition of SPPAM in (5). In contrast to the aforementioned works, we include Polyak’s momentum (Polyak, 1964) directly to SPPA, yielding (5). Apart from the similarity between SPPAM in (5) and SGDM in (3), SPPAM shares the same geometric intuition as Polyak’s momentum for SGDM. Disregarding the stochastic errors, the update in (5) follows from the solution of:

$$\arg \min_{x \in \mathbb{R}^p} \left\{ f(x) + \frac{1}{2\eta} \|x - x_t\|_2^2 - \frac{\beta}{\eta} \langle x_t - x_{t-1}, x \rangle \right\}.$$

We can get a sense of the behavior of SPPAM from this expression. First, for large η , the algorithm is minimizing the original f . For small η , the algorithm not only tries to stay local by minimizing the quadratic term, but also tries to minimize $-\frac{\beta}{\eta} \langle x_t - x_{t-1}, x \rangle$. By the definition of inner product, this means that x , on top of minimizing f and staying close to x_t , also tries to move along the direction from x_{t-1} to x_t . This intuition aligns with that of Polyak’s momentum.

3. The Quadratic Model Case




Figure 1: We generate $A \in \mathbb{R}^{p \times p}$ and $b, x^* \in \mathbb{R}^p$ from $\mathcal{N}(0, I)$, where $p = 100$ and the condition number of A is 10. We sweep η and β from -5 to 5 , with 0.2 interval. We plot the accuracy $\|x_t - x^*\|_2^2$ after 100 iterations, with the maximum replaced by 10.

For simplicity, we first consider the convex quadratic optimization problem under the deterministic setting. Specifically, we consider the objective function:

$$f(x) = \frac{1}{2} x^\top A x - b^\top x, \quad (10)$$

where $A \in \mathbb{R}^{p \times p}$ is positive semi-definite with eigenvalues $[\lambda_1, \dots, \lambda_p]$. Under this scenario, we can study how the step size η and the momentum β affect each other, by deriving exact conditions that lead to convergence for each algorithm. The comparison list includes gradient descent (GD), gradient descent with momentum (GDM), the PPA, and PPA with momentum (PPAM). Propositions 1 and 3 for GD and GDM are from Goh (2017), and included for completeness. Proofs for PPA and PPAM in Propositions 2 and 4 can be found in the extended version of this work (Kim et al., 2021).

Proposition 1 (GD (Goh, 2017)) *To minimize (10) with gradient descent, the step size η needs to satisfy $0 < \eta < \frac{2}{\lambda_i} \forall i$, where λ_i is the i -th eigenvalue of A .*

Proposition 2 (PPA) *To minimize (10) with PPA, the step size η needs to satisfy $\left| \frac{1}{1+\eta\lambda_i} \right| < 1 \forall i$.*

Proposition 3 (GDM (Goh, 2017)) *To minimize (10) with gradient descent with momentum, the step size η needs to satisfy $0 < \eta\lambda_i < 2 + 2\beta \ \forall i$, where $0 \leq \beta < 1$.*

Proposition 4 (PPAM) *Let $\delta_i = \left(\frac{\beta+1}{1+\eta\lambda_i}\right)^2 - \frac{4\beta}{1+\eta\lambda_i}$. To minimize (10) with PPAM, the step size η and momentum β need to satisfy $\forall i$:*

- $\eta > \frac{\beta-1}{\lambda_i}$, *if $\delta_i \leq 0$;*
- $\frac{\beta+1}{1+\eta\lambda_i} + \sqrt{\delta_i} < 2$, *if $\delta_i > 0$ and $\frac{\beta+1}{1+\eta\lambda_i} \geq 0$;*
- $\frac{\beta+1}{1+\eta\lambda_i} - \sqrt{\delta_i} > -2$, *otherwise.*

Given the above propositions, we can study the stability with respect to the step size η and the momentum β for the considered algorithms. Numerical simulations support the above propositions, and are illustrated in Figure 1, matching the theoretical conditions exhibited above. In particular, for GD (1st), only a small range of step sizes η leads to convergence (small white band); this “white band” corresponds to the restriction that η has to satisfy $\eta < \frac{2}{\lambda_i}$ for all i . On the other hand, PPA/IGD (2nd) converges in much wider choices of η ; this is apparent from Proposition 2, since $\left|\frac{1}{1+\eta\lambda_i}\right|$ can be arbitrarily small for larger values of η . GDM (3rd) requires both η and β to be in a small region to converge, following Proposition 3. Finally, PPAM (4th) converges in much wider choices of η and β ; e.g., the conditions in Proposition 4 define different regions of the pair (η, β) that lead to convergence, some of which set both η and β to be negative. Note that the empirical convergence region of PPAM almost exactly matches the theoretical region that leads to convergence in Proposition 4 (5th). In the remainder of the paper, we study how such pattern translates to a general strongly convex function f , with stochasticity.

4. Theoretical Results

In this section, we theoretically characterize the convergence and stability behavior of SPPAM.¹ We follow the stochastic errors of PPA, as set up in [Toulis et al. \(2021\)](#); we can thus express (5) as:²

$$\begin{aligned} x_{t+1}^+ &= x_t - \eta \nabla f(x_{t+1}^+) + \beta(x_t - x_{t-1}) \\ x_{t+1} &= x_{t+1}^+ - \eta \varepsilon_{t+1}. \end{aligned}$$

We further assume the following:

Assumption 1 *f is a μ -strongly convex function. That is, for some fixed $\mu > 0$ and for all x and y ,*

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \mu \|x - y\|_2^2.$$

Assumption 2 *There exists fixed $\sigma^2 > 0$ such that, given the natural filtration \mathcal{F}_{t-1} ,*

$$\mathbb{E}[\varepsilon_t \mid \mathcal{F}_{t-1}] = 0 \text{ and } \mathbb{E}[\|\varepsilon_t\|^2 \mid \mathcal{F}_{t-1}] \leq \sigma^2 \text{ for all } t.$$

We now study whether SPPAM enjoys faster convergence than SPPA in (4). We start with the iteration invariant bound:

-
1. All proofs are provided in the extended version of this work ([Kim et al., 2021](#)).
 2. x_{t+1}^+ is an auxiliary intermediate variable that is only used for the analysis.

Theorem 5 For μ -strongly convex f , SPPAM satisfies the following iteration invariant bound:

$$\mathbb{E} [\|x_{t+1} - x^*\|_2^2] \leq \frac{4}{(1+\eta\mu)^2} \mathbb{E} [\|x_t - x^*\|_2^2] + \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)} \mathbb{E} [\|x_{t-1} - x^*\|_2^2] + \eta^2 \sigma^2. \quad (11)$$

Notice that all terms –except the last one– are divided by $(1 + \eta\mu)^2$. Thus, large step size η help convergence (to a neighborhood), reminiscent of the convergence behavior of PPA in (8). Based on (11), we can write the following 2×2 system that characterizes the progress of SPPAM:

$$\begin{bmatrix} \mathbb{E} [\|x_{t+1} - x^*\|_2^2] \\ \mathbb{E} [\|x_t - x^*\|_2^2] \end{bmatrix} \leq \underbrace{\begin{bmatrix} \frac{4}{(1+\eta\mu)^2} & \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)} \\ 1 & 0 \end{bmatrix}}_A \cdot \begin{bmatrix} \mathbb{E} [\|x_t - x^*\|_2^2] \\ \mathbb{E} [\|x_{t-1} - x^*\|_2^2] \end{bmatrix} + \begin{bmatrix} \eta^2 \sigma^2 \\ 0 \end{bmatrix}. \quad (12)$$

It is clear that the spectrum of the contraction matrix A determines the convergence rate to a neighborhood, as in (Goh, 2017). This is summarized in the following lemma:

Lemma 6 The maximum eigenvalue of A , which determines the convergence rate of SPPAM, is:

$$\frac{2}{(1+\eta\mu)^2} + \sqrt{\frac{4}{(1+\eta\mu)^4} + \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)}}. \quad (13)$$

Notice the one-step contraction factor in (13) is of order $O(1/\eta^2)$, exhibiting acceleration compared to that of SPPA for strongly convex objectives (Toulis et al., 2021): $1/(1 + 2\eta\mu) \approx O(1/\eta)$. However, due to the additional terms, it is not immediately obvious when SPPAM enjoys faster contraction than SPPA. We thus characterize this condition more precisely in the following corollary:

Corollary 7 For μ -strongly convex f , SPPAM enjoys smaller contraction factor than SPPA if:

$$\frac{4\beta^2}{4 - (1 + \beta)^2} < \frac{\eta^2 \mu^2 - 6\eta\mu - 3}{(1 + \eta\mu)^2}.$$

In words, for a fixed step size η and given a strongly convex parameter μ , there is a range of momentum parameters β that exhibits acceleration compared to SPPA.

Remark 8 In contrast to (stochastic) gradient method analyses in convex optimization, where acceleration is usually shown by improving the dependency on the condition number from $\kappa = \frac{L}{\mu}$ to $\sqrt{\kappa}$, such a claim can hardly be made for stochastic proximal point methods. This is also the case in deterministic setting; see (8) and (9). As shown in Theorem 5, our convergence analysis of SPPAM does not depend on L -smoothness at all. This robustness of SPPAM is also confirmed in numerical simulations in Section 5, where SPPAM exhibits the fastest convergence rate, virtually independent of the different settings considered.

We now formalize the convergence behavior of SPPAM. In particular, we characterize the condition that leads to an exponential discount of the initial conditions. By unrolling the recursion of SPPAM in (12) for T iterations, we obtain:

$$\begin{bmatrix} \mathbb{E} [\|x_T - x^*\|_2^2] \\ \mathbb{E} [\|x_{T-1} - x^*\|_2^2] \end{bmatrix} \leq A^T \cdot \begin{bmatrix} \|x_0 - x^*\|_2^2 \\ \|x_{-1} - x^*\|_2^2 \end{bmatrix} + \left(\sum_{i=1}^{T-1} A^i \right) \begin{bmatrix} 1 \\ 0 \end{bmatrix} \eta^2 \sigma^2.$$

It is clear from the above that the convergence is determined by A^T and $\left(\sum_{i=1}^{T-1} A^i \right)$, where A was defined in (12). Our next theorem derives convergence to a neighborhood based on the spectrum of these quantities, akin to Assran and Rabbat (2020, Theorem 1) and Toulis et al. (2021, Theorem 3).

Theorem 9 For μ -strongly convex f , assume SPPAM is initialized with $x_0 = x_{-1}$. Then, after T iterations, we have:

$$\mathbb{E} [\|x_T - x^*\|_2^2] \leq \frac{2\sigma_1^T}{\sigma_1 - \sigma_2} \left(\left(\|x_0 - x^*\|_2^2 + \frac{\eta^2 \sigma^2}{1-\theta} \right) \cdot (1 + \theta) \right) + \frac{\eta^2 \sigma^2}{1-\theta}, \quad (14)$$

where $\theta = \frac{4}{(1+\eta\mu)^2} + \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)}$. Here, $\sigma_{1,2}$ are the eigenvalues of A , and

$$\frac{2\sigma_1^T}{\sigma_1 - \sigma_2} = \tau^{-1} \cdot \left(\frac{2}{(1+\eta\mu)^2} + \tau \right)^T \quad \text{with} \quad \tau = \sqrt{\frac{4}{(1+\eta\mu)^4} + \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)}}. \quad (15)$$

The above theorem states that the term in (15) determines the discounting rate of the initial conditions. In particular, the condition that leads to an exponential discount of the initial conditions is characterized by the following theorem:

Theorem 10 Let the following condition hold:

$$\tau = \sqrt{\frac{4}{(1+\eta\mu)^4} + \frac{4\beta^2}{(1+\eta\mu)^2(4-(1+\beta)^2)}} < \frac{1}{2}. \quad (16)$$

Then, for μ -strongly convex f , initial conditions of SPPAM exponentially discount: i.e., in (14),

$$\frac{2\sigma_1^T}{\sigma_1 - \sigma_2} = \tau^{-1} \cdot \left(\frac{2}{(1+\eta\mu)^2} + \tau \right)^T = C^T, \quad \text{where} \quad C \in (0, 1).$$

Remark 11 The condition in (16) is much easier to satisfy than accelerated SGD. E.g., as described below, the required condition for accelerated SGD to converge linearly to a neighborhood in strongly convex quadratic objective relies on knowing $\eta = \frac{1}{L}$ and momentum $\beta = \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ (Assran and Rabbat, 2020), where both L and κ are unknown in practice. While this is also true for SPPAM (i.e., μ is an unknown quantity), (16) suggests that one can essentially set η sufficiently large to ensure the exponential discount, even without knowing μ exactly.

Remark 12 Other works that study variants of accelerated SPPA (Kulunchakov and Mairal, 2019; Chadha et al., 2021) still require specific choices of step size and momentum (e.g., $\eta_t = \frac{1}{L+c_0\sqrt{t+1}}$, $\beta_t = \frac{2}{t+2}$ for the latter; see Table 1 for the former), similarly to accelerated SGD.

To provide more context of the condition in Theorem 10, we make an “unfair” comparison of (16), which holds for general strongly convex f , to the condition that accelerated SGD requires for strongly convex quadratic objective in (10). Assran and Rabbat (2020, Theorem 1) show that Nesterov’s accelerated SGD converges to a neighborhood at a linear rate for strongly convex quadratic objective if $\max\{\rho_\mu(\eta, \beta), \rho_L(\eta, \beta)\} < 1$, where $\rho_\lambda(\eta, \beta)$ for $\lambda \in \{\mu, L\}$ is defined as:

$$\rho_\lambda(\eta, \beta) = \begin{cases} \frac{|(1+\beta)(1-\eta\lambda)|}{2} + \frac{\sqrt{\Delta_\lambda}}{2} & \text{if } \Delta_\lambda \geq 0, \\ \sqrt{\beta(1-\eta\lambda)} & \text{otherwise,} \end{cases} \quad (17)$$

with $\Delta_\lambda = (1+\beta)^2(1-\eta\lambda)^2 - 4\beta(1-\eta\lambda)$. This condition for convergence can thus be divided into three cases, depending on the range of $\eta\lambda$. Define $\psi_{\beta,\eta,\lambda} = (1+\beta)(1-\eta\lambda)$. Then:

$$\begin{cases} \eta\lambda \geq 1, & \text{Converges if } -\psi_{\beta,\eta,\lambda} + \sqrt{\Delta_\lambda} < 2, \\ \frac{(1-\beta)^2}{(1+\beta)^2} \leq \eta\lambda < 1, & \text{Always converges,} \\ \eta\lambda < \frac{(1-\beta)^2}{(1+\beta)^2}, & \text{Converges if } \psi_{\beta,\eta,\lambda} + \sqrt{\Delta_\lambda} < 2. \end{cases}$$

Now, consider the standard momentum value $\beta = 0.9$. For the first case, the convergence requirement translates to $1 \leq \eta\lambda \leq \frac{24}{19}$. The second range is given by $\frac{1}{361} \leq \eta\lambda < 1$. The third condition is lower bounded by 2 for $\beta = 0.9$, leading to divergence. Combining, accelerated SGD requires $0.0028 \approx \frac{1}{361} \leq \eta\lambda \leq \frac{24}{19} \approx 1.26$ to converge for strongly convex quadratic objectives, set aside that this bound has to satisfy for (unknown) μ or L . Albeit an unfair comparison, for general strongly convex objective, (16) becomes $\eta\mu > 4.81$ for $\beta = 0.9$. Even though μ is unknown, one can see this condition is easy to satisfy, by using a sufficiently large step size η .

5. Experiments

In this section, we perform numerical experiments to study the convergence behaviors of SP-PAM, SPPA, SGDM, and SGD, using generalized linear models (GLM) [Nelder and Wedderburn \(1972\)](#). Let $b_i \in \mathbb{R}$ be the label, $a_i \in \mathbb{R}^p$ be the features, and $x^* \in \mathbb{R}^p$ be the model parameter of interest. GLM assumes that b_i follows an exponential family distribution: $b_i \mid a_i \sim \exp\left(\frac{\gamma b_i - c_1(\gamma)}{\omega} c_2(b_i, \omega)\right)$. Here, $\gamma = \langle a_i, x^* \rangle$ is the linear predictor, ω is the dispersion parameter related to the variance of b_i , and $c_1(\cdot)$ and $c_2(\cdot)$ are known real-valued functions. GLM subsumes a wide family of models including linear, logistic, and Poisson regressions. Different models connects the linear predictor $\gamma = \langle a_i, x^* \rangle$ through different *mean* functions $h(\cdot)$. We focus on linear and Poisson regression models, where mean functions are defined respectively as $h(\gamma) = \gamma$ and $h(\gamma) = e^\gamma$. The former is an “easy” case, where objective is strongly convex, satisfying Assumption 1. The latter is a “hard” case with non-Lipschitz continuous gradients, where SGD and SGDM are expected to suffer.

[Toulis et al. \(2014\)](#) introduced an efficient, exact implementation of SPPA for GLM. We adapt this procedure to SPPAM in (5); see Algorithm 1. Its derivation can be found in [Kim et al. \(2021\)](#).

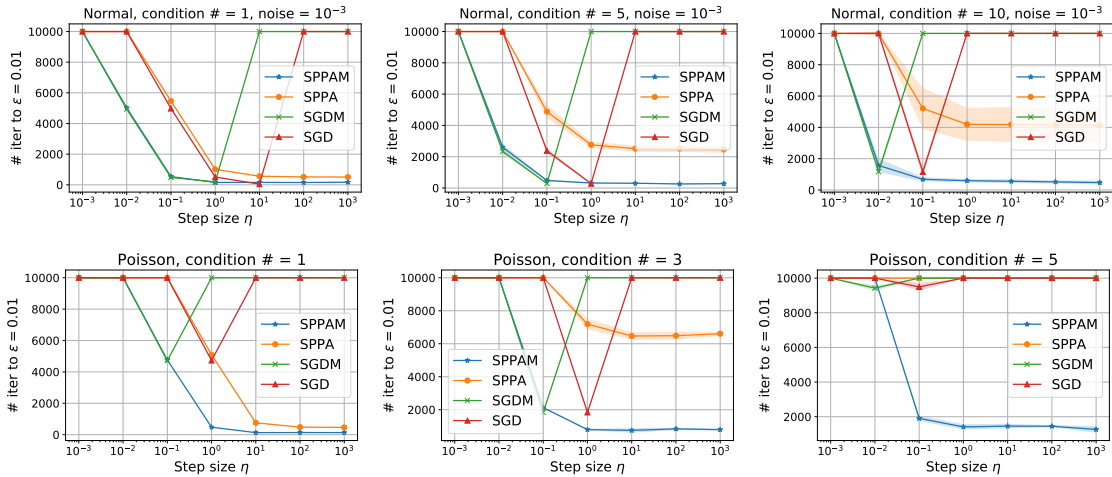


Figure 2: **Top:** Linear regression with condition number $\kappa \in \{1, 5, 10\}$ with gaussian noise level $1e-3$. **Bottom:** Poisson regression with condition number $\kappa \in \{1, 3, 5\}$. We set $p = n = 100$ in both cases. Batch size is 10 for all algorithms. The median number of iterations to reach $\epsilon = 0.01$ is plotted. Shaded area are the standard deviations across 5 experiments.

Algorithm 1 SPPAM for GLM

```

for  $t = 1, 2, \dots$  do
    Sample  $i_t \sim \text{Unif}(1, n)$ 
     $r_t \leftarrow \eta(b_{i_t} - h(\langle a_{i_t}, x_{t-1} \rangle))$ 
     $B_t \leftarrow [0, r_t]$ 
    if  $r_t \leq 0$  then
        |  $B_t \leftarrow [r_t, 0]$ 
    end
     $\xi_t = \eta[b_{i_t} - h((1 + \beta)\langle a_{i_t}, x_{t-1} \rangle$ 
         $- \beta\langle a_{i_t}, x_{t-2} \rangle + \xi_t \cdot \|a_{i_t}\|_2^2)]$ ,  $\xi_t \in B_t$ 
     $x_t \leftarrow x_{t-1} + \xi_t \cdot a_{i_t} + \beta(x_{t-1} - x_{t-2})$ 
end
    
```

We generate the data as follows. $A \in \mathbb{R}^{p \times n}$ and $x^* \in \mathbb{R}^p$ are drawn from $\mathcal{N}(0, I)$. For the normal case, we generate $b_i = \langle a_i, x^* \rangle$, and for the Poisson case, we generate $b_i \sim \text{Poisson}(e^{\langle a_i, x^* \rangle})$ for $i = 1, \dots, n$. For each experimental setup, we run SPPAM (blue), SPPA (orange), SGDM (green), and SGD (red) for 10^4 iterations. We repeat each experiment for 5 independent trials, and plot the median number of iterations to reach precision $\varepsilon \leq 10^{-2}$, along with the standard deviation. We measure the precision $\varepsilon = \frac{\|b - \hat{b}\|_2^2}{\|b\|_2^2}$, where b is the true label and \hat{b} is the predicted label.

In Figure 2 (Top), we present the results for the linear regression with different condition numbers, with gaussian noise level $1e-3$. We run each algorithm constant step size η varying from 10^{-3} to 10^3 with $10\times$ increment, and with $\beta = 0.9$. As expected, SGD and SGDM only converge for specific step size η , while SPPA and SPPAM converge for much wider ranges. In terms of convergence rate, SPPAM converges faster than SPPA in all scenarios, which improves or matches the rate of SGDM, when it converges. As κ increases, the range of η that leads to convergence for SGD and SGDM shrinks; notice the sharper “ \vee ” shape for SGD and SGDM for $\kappa = 10$ (3rd), compared to $\kappa = 5$ (2nd) or $\kappa = 1$ (1st). SPPA also slightly slows down as κ increases, while SPPAM converges essentially in the same manner for all scenarios.

Such trend is much more pronounced for the Poisson regression case presented in Figure 2 (Bottom). Due to the exponential mean function $h(\cdot)$ for Poisson model, the outcomes are extremely sensitive, and its likelihood does not satisfy standard assumptions like L -smoothness. As such, SGD and SGDM struggles with slow convergence even when $\kappa = 1$ (1st), while also exhibiting instability—each method converges only for a single choice of η considered. Similar trend is shown when $\kappa = 3$ (2nd) where SPPA starts slowing down. For $\kappa = 5$ (3rd), all methods except for SPPAM did not make much progress in 10^4 iterations, for the entire range of η and β considered. Quite remarkably, SPPAM still converges in the same manner without sacrificing both the convergence rate and the range of hyperparameters that lead to convergence.

6. Conclusion

We propose the stochastic proximal point algorithm with momentum (SPPAM), which directly incorporates Polyak’s momentum inside the proximal step. We show that SPPAM converges to a neighborhood at a faster rate than stochastic proximal point algorithm (SPPA), and characterize the conditions that result in acceleration. Further, we prove linear convergence of SPPAM to a neighborhood, and provide conditions that lead to an exponential discount of the initial conditions, akin to SPPA. We confirm our theory with numerical simulations on linear and Poisson regression models; SPPAM converges for all the step sizes that SPPA converges, with a faster rate that matches or improves SGDM.

Acknowledgments

This work is supported by NSF FET:Small no.1907936, NSF MLWiNS CNS no.2003137 (in collaboration with Intel), NSF CMMI no.2037545, NSF CAREER award no.2145629, and Rice Inter-Disciplinary Excellence Award (IDEA).

References

- Kwangjun Ahn and Suvrit Sra. Understanding nesterov’s acceleration via proximal point method. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 117–130. SIAM, 2022.
- Z. Allen-Zhu and L. Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- Hilal Asi and John C Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019.
- Hilal Asi, Karan Chadha, and Gary Cheng. Minibatch Stochastic Approximate Proximal Point Methods. *34th Conference on Neural Information Processing Systems*, page 11, 2020.
- Mahmoud Assran and Michael Rabbat. On the Convergence of Nesterov’s Accelerated Gradient Method in Stochastic Settings. *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate $o(1/n)$. In *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 1*, pages 773–781, 2013.
- Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.
- Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. *Advances in neural information processing systems*, 20, 2007.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- S. Bubeck, Y.-T. Lee, and M. Singh. A geometric alternative to Nesterov’s accelerated gradient descent. *arXiv preprint arXiv:1506.08187*, 2015.
- Karan Chadha, Gary Cheng, and John C Duchi. Accelerated, optimal, and parallel: Some results on model-based stochastic optimization. *arXiv preprint arXiv:2101.02696*, 2021.
- Alexandre d’Aspremont. Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, 2008.
- A. Defazio. On the curved geometry of accelerated optimization. In *Advances in Neural Information Processing Systems*, pages 1764–1773, 2019.

- Qi Deng and Wenzhi Gao. Minibatch and momentum model-based methods for stochastic weakly convex optimization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1):37–75, 2014.
- Gabriel Goh. Why momentum really works. *Distill*, 2(4):e6, 2017.
- Robert M Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtarik. SGD: General Analysis and Improved Rates. *Proceedings of the 36 th International Conference on Machine Learning*, page 10, 2019.
- Osman Güler. On the convergence of the proximal point algorithm for convex minimization. *SIAM journal on control and optimization*, 29(2):403–419, 1991.
- Osman Güler. New proximal point algorithms for convex minimization. *SIAM Journal on Optimization*, 2(4):649–664, 1992.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- B. Hu and L. Lessard. Dissipativity theory for Nesterov’s accelerated method. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1549–1557. JMLR.org, 2017.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- Rahul Kidambi, Praneeth Netrapalli, Prateek Jain, and Sham Kakade. On the insufficiency of existing momentum schemes for stochastic optimization. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–9. IEEE, 2018.
- Junhyung Lyle Kim, Panos Toulis, and Anastasios Kyrillidis. Convergence and stability of the stochastic proximal point algorithm with momentum. *arXiv preprint arXiv:2111.06171*, 2021.
- Andrei Kulunchakov and Julien Mairal. A Generic Acceleration Framework for Stochastic Composite Optimization. *Advances in Neural Information Processing Systems*, 32, October 2019. arXiv: 1906.01164.
- Maxime Laborde and Adam Oberman. A Lyapunov analysis for accelerated gradient methods: from deterministic to stochastic case. In *International Conference on Artificial Intelligence and Statistics*, pages 602–612. PMLR, 2020.
- L. Lessard, B. Recht, and A. Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.

- H. Lin, Julien Mairal, and Zaid Harchaoui. Catalyst Acceleration for First-order Convex Optimization: from Theory to Practice. *Journal of Machine Learning Research*, 18:1–54, 2018.
- Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. *Advances in neural information processing systems*, 28, 2015.
- Chaoyue Liu and Mikhail Belkin. Accelerating sgd with momentum for over-parameterized learning. *arXiv preprint arXiv:1810.13395*, 2018.
- Nicolas Loizou and Peter Richtárik. Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *Computational Optimization and Applications*, 77(3):653–710, 2020.
- Eric Moulines and Francis R. Bach. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 451–459. Curran Associates, Inc., 2011.
- John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384, 1972.
- Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- Yurii Nesterov et al. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Yurii E Nesterov. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. In *Dokl. akad. nauk Sssr*, volume 269, pages 543–547, 1983.
- Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *Ussr computational mathematics and mathematical physics*, 4(5):1–17, 1964.
- Boris T. Polyak. Introduction to optimization. *Inc., Publications Division, New York*, 1, 1987.
- Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- R Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- Ernest K Ryu and Stephen Boyd. Stochastic Proximal Iteration: A Non-Asymptotic Improvement Upon Stochastic Gradient Descent. *Author website*, page 42, 2017.
- Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: Primal estimated sub-gradient solver for SVM. *Mathematical programming*, 127(1):3–30, 2011.

- W. Su, S. Boyd, and E. Candes. A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2014.
- Panagiotis Toulis, Edoardo Airoldi, and Jason Rennie. Statistical analysis of stochastic gradient methods for generalized linear models. In *International Conference on Machine Learning*, pages 667–675. PMLR, 2014.
- Panos Toulis and Edoardo M Airoldi. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *The Annals of Statistics*, 45(4):1694–1727, 2017.
- Panos Toulis, Thibaut Horel, and Edoardo M Airoldi. The proximal robbins–monro method. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83(1):188–212, 2021.
- A. Wibisono, A. Wilson, and M. Jordan. A variational perspective on accelerated methods in optimization. *proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.
- Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, page 116, 2004.