# First Order Methods with Markovian Noise: from Acceleration to Variational Inequalities 

Aleksandr Beznosikov<br>Innopolis University, Skoltech, MIPT, Yandex

Sergey Samsonov<br>Marina Sheshukova<br>HSE University<br>HSE University

Alexander Gasnikov<br>MIPT, Skoltech, IITP RAS

Alexey Naumov<br>HSE University

Eric Moulines<br>Ecole polytechnique


#### Abstract

This paper delves into stochastic optimization problems that involve Markovian noise. We present a unified approach for the theoretical analysis of first-order gradient methods for stochastic optimization and variational inequalities. Our approach covers scenarios for both non-convex and strongly convex minimization problems. To achieve an optimal (linear) dependence on the mixing time of the underlying noise sequence, we use the randomized batching scheme, which is based on the multilevel Monte Carlo method. Moreover, our technique allows us to eliminate the limiting assumptions of previous research on Markov noise, such as the need for a bounded domain and uniformly bounded stochastic gradients. Our extension to variational inequalities under Markovian noise is original. Additionally, we provide lower bounds that match the oracle complexity of our method in the case of strongly convex optimization problems.


## 1 Introduction

Stochastic gradient methods are an essential ingredient for solving various optimization problems, with a wide range of applications in various fields such as machine learning [36, 37], empirical risk minimization problems [96], and reinforcement learning [93, 85, 69]. Various stochastic gradient descent methods (SGD) and their accelerated versions [75, 31] have been extensively studied under different statistical frameworks [17, 97]. The standard assumption for stochastic optimization algorithms is to consider independent and identically distributed noise variables. However, the growing usage of stochastic optimization methods in reinforcement learning [10, 87, 25] and distributed optimization [63, 18, 65] has led to increased interest in problems with Markovian noise. Despite this, existing theoretical works that consider Markov noise have significant limitations, and their analysis often results in suboptimal finite-time error bounds.
Our research aims to fill the gap in the existing literature on the first-order Markovian setting. By focusing on uniformly geometrically ergodic Markov chains, we obtain finite-time complexity bounds for achieving $\varepsilon$-accurate solutions that scale linearly with the mixing time of the underlying Markov chain. Our approach is based on careful applications of randomized batch size schemes and provides a unified view on both non-convex and strongly convex minimization problems, as well as variational inequalities.
Our contributions. Our main contributions are the following:
$\diamond$ Accelerated SGD. We provide the first analysis of SGD, including the Nesterov accelerated SGD method, with Markov noise without the assumption of bounded domain and uniformly bounded stochastic gradient estimates. Our results are summarised in Table 1 and Section 2.1 and cover both strongly convex and non-convex scenarios. Our findings for non-convex minimization problems complement the results obtained in [21].
$\diamond$ Lower bounds. In Section 2.2 we give the lower bounds showing that the presence of mixing time in the upper complexity bounds is not an artefact of the proof. This is consistent with the results reported in [71].
$\diamond$ Extensions. In Section 2.4 we provide, as far as we know, the first analysis for variational inequalities with general stochastic Markov oracle, arbitrary optimization set, and arbitrary composite term. Our finite-time performance analysis provides complexity bounds in terms of oracle calls that scale linearly with the mixing time of the underlying chain, which is an improvement over the bounds obtained in [99] for the Markov setting.
Related works. Next, we briefly summarize the related works.
$\diamond$ Stochastic gradient methods. Numerous research papers have reported significant improvements achieved by accelerated methods for stochastic optimization with stochastic gradient oracles involving independent and identically distributed (i.i.d.) noise. These methods have been extensively studied in theory [44, 14, 16, 58, 61, 26, 30, 97, 94, 3, 39, 102] and have shown practical success [55, 91]. The finite-time analysis of first-order methods in i.i.d. noise settings has been extensively studied by many authors, as discussed in [59] and references therein. In Table 1] we include only some important results because i.i.d. setting is not in the interest of this paper.
While the literature on i.i.d. noise is extensive, existing research on the first-order Markovian setting is relatively sparse. In this study, we focus on Markov chains that are uniformly geometrically ergodic, and we refer the reader to Section 2 for detailed definitions. We note that the complexity bounds which scale linearly with the mixing time of the underlying general Markov chain are currently available only for general convex and non-convex minimization problems. Namely, [23] has investigated a version of the ergodic mirror descent algorithm that yields optimal convergence rates for Lipschitz, general convex and non-convex problems. Recently, [21] proposed a random batch size algorithm that adapts to the mixing time of the underlying chain for non-convex optimization with a compact domain. In particular, [21] Theorem 4.3] yields optimal complexity rates in terms of the number of oracle calls required for non-convex problems, which is consistent with the results obtained in [23]. Unlike previous studies, this method is insensitive to the mixing time of the noise sequence.
For the general case of Markovian noise the finite-time analysis of non-accelerated SGD-type algorithms was carried out in [90] and [19]. However, [90] heavily relies on the bounded domain assumption and uniformly bounded stochastic gradient oracles, while its bound in [90, Theorem 5] has a suboptimal dependence on the mixing time of the underlying chain, see Table 1 Additionally, [90] does not cover the strongly convex setting. On the other hand, [19] covers both non-convex and strongly convex settings, but the bounds of [19. Theorem 1] has terms that are exponential in the mixing time, and a careful examination reveals suboptimal dependence on the initial condition for strongly convex problems when SGD is applied.
In the study of Nesterov-accelerated SGD with Markovian noise, the authors of [20] considered the use of a batch size of 1 and achieved a rate of forgetting the initial condition that matches that of the i.i.d. noise setting. However, their result is suboptimal in terms of the variance terms in both non-convex and strongly convex settings, as detailed in Table 1. We emphasize that the case of unbounded gradient oracles with Markov noise is not treated in contrast to the i.i.d. setup [97, 62].

The above papers deal with general Markovian noise optimization. But there are also results that deal with Markovian stochasticity with a finite state space. Here we can highlight the work [28], where the author gives quite extensive results and achieves linear scaling by mixing time in the non-convex as well as strongly convex cases. Recently, numerous papers have appeared dealing with the special scenario of distributed optimization [89]. [99] investigates the generalization and stability of Markov SGD with special attention to the excess variance guarantees. We note that first, these algorithms only need to deal with a very special case of Markov gradients, and second, the corresponding dependence on the mixing time of the Markov chain is again quadratic. At the same time, there exist particular results, e.g. [71], which provide a lower bound for the particular finite sum problems in the Markovian setting.
$\diamond$ Variational inequalities. Variational inequalities [29] have been an active area of research in applied mathematics for more than half a century [78, 41, 86]. VI cover important special cases, e.g., minimization over a convex domain, saddle point or min-max and fixed point problems. computational game theory [29], robust [7] and nonsmooth [73, 72] optimization, supervised [51, 4] and unsupervised [103, 5] learning, image denoising [27, 11]. In the last 5 years, variational inequalities and their special cases have attracted much interest in the machine learning community
due to new connections to reinforcement learning [79, 50], adversarial training [64], and GANs [15, 33, 66, 12, 60, 82].
Variational inequalities (VI) and saddle point problems have their own well-established theory and methods. Unlike minimization problems, solving variational inequalities doesn't rely on (accelerated) gradient descent. Instead, the extragradient method [57], various modified versions [72, 42], or similar techniques [95] are recommended as the basic and theoretically optimal methods. While deterministic methods have long been used for solving variational inequalities, stochastic methods have gained importance only in the last 15 years, following pioneering works by [49, 52]. We summarise the results on methods for stochastic variational inequalities with the Lipschitz operator and smooth stochastic saddle point problems in Table 2. The number of papers dealing with stochastic VIs and saddle point problems is small compared to those dealing with stochastic optimization, we include in Table 2 papers with the i.i.d. noise (which we do not do for stochastic optimization). The only competing work dealing with Markovian noise in saddle point problems consider the finite sum problem and thus the finite Markov chain [99], therefore we do not include it in Table 2 Moreover, the results from [99] has much worse oracle complexity guarantees $\mathcal{O}\left(\tau^{2} / \varepsilon^{2}\right)$ in terms of $\tau$. There are more papers dealing with stochastic finite-sum variational inequalities or saddle point problems, but in the i.i.d. setting [12, 80, 104, 2, 8]. We also do not consider in Table 2 because of the difference in the stochastic oracle structure. It is important to note that, unlike most previous works, we consider the most general formulation of VI itself for an arbitrary optimization set and composite term.
Notations and definitions. Let $\left(Z, d_{z}\right)$ be a complete separable metric space endowed with its Borel $\sigma$-field $\mathcal{Z}$. Let $\left(Z^{\mathbb{N}}, \mathcal{Z}^{\otimes \mathbb{N}}\right)$ be the corresponding canonical process. Consider the Markov kernel Q defined on $\mathrm{Z} \times \mathcal{Z}$, and denote by $\mathbb{P}_{\xi}$ and $\mathbb{E}_{\xi}$ the corresponding probability distribution and the expected value with initial distribution $\xi$. Without loss of generality, we assume that $\left(Z_{k}\right)_{k \in \mathbb{N}}$ is the corresponding canonical process. By construction, for any $A \in \mathcal{Z}$, it holds that $\mathbb{P}_{\xi}\left(Z_{k} \in A \mid Z_{k-1}\right)=\mathrm{Q}\left(Z_{k-1}, A\right), \mathbb{P}_{\xi}$-a.s. If $\xi=\delta_{z}, z \in Z$, we write $\mathbb{P}_{z}$ and $\mathbb{E}_{z}$ instead of $\mathbb{P}_{\delta_{z}}$ and $\mathbb{E}_{\delta_{z}}$, respectively. For $x^{1}, \ldots, x^{k}$ being the iterates of any stochastic first-order method, we denote $\mathcal{F}_{k}=\sigma\left(x^{j}, j \leq k\right)$ and write $\mathbb{E}_{k}$ as an alias for $\mathbb{E}\left[\cdot \mid \mathcal{F}_{k}\right]$. We also write $\mathbb{N}^{*}:=\mathbb{N} \backslash\{0\}$. For the sequences $\left(a_{n}\right)_{n \in \mathbb{N}}$ and $\left(b_{n}\right)_{n \in \mathbb{N}}$ we write $a_{n} \lesssim b_{n}$ if there exists a constant $c$ such that that $a_{n} \leq c b_{n}$ for all $n \in \mathbb{N}$.
Table 1: This table summarizes our results on first-order method with Markovian noise. The columns of the table indicate whether the authors consider optimization over bounded domain, potentially unbounded gradients, and whether or not they assume additional restrictions on the Markovian noise (finite state space or reversibility). For ease of comparison we provide the respective results on SGD and ASGD (accelerated SGD) in the i.i.d. setting.

|  |  | Unbounded |  |  |  | Oracle complexity (Smooth and non-convex) | Oracle complexity(Smooth and strongly convex) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Method | Domain | $\begin{gathered} \text { Gradient } \\ \text { noise } \end{gathered}$ | General | Acceleration |  |  |
| : | SGD [84.70. 321 | $\checkmark$ | $x$ | N/A | $x$ | $\tilde{\mathcal{O}}\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)\left[\frac{1}{\varepsilon^{2}}+\frac{\frac{\sigma}{}^{2}}{\varepsilon^{4}}\right]\right)$ | $\tilde{\mathcal{O}}\left(\frac{L}{\mu} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)$ |
|  | ASGD $97.13{ }^{(1)}$ | $\checkmark$ | $\checkmark$ | N/A | $\checkmark$ | $\tilde{\mathcal{O}}\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)\left[\frac{1+\delta^{2}}{\varepsilon^{2}}+\frac{\frac{\sigma}{}^{2}}{\varepsilon^{4}}\right]\right)$ | $\tilde{\mathcal{O}}\left(\left(1+\delta^{2}\right) \sqrt{\frac{L}{\mu}} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)$ |
|  | EMD [ 23$]^{2+}$ | $\times$ | $x$ | $\checkmark$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{\tau G^{2} D^{2}}{\varepsilon^{4}}\right)$ | $x$ |
|  | MC SGD [901 ${ }^{\text {3 }}$ | $\checkmark$ | $x$ | $\times$ | $x$ | $\tilde{\mathcal{O}}\left(h(G, L)\left(\frac{t}{\varepsilon^{2}}\right)^{1 /(1-q)}\right)$ | $\chi^{x}$ |
|  | mC Sgd [191) | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{\tau L^{2}\left(1+\left\\|x^{*}\right\\|^{2}+\left\\|x^{0}-x^{*}\right\\|^{2}\right)}{\varepsilon^{4}}\right)$ | $\tilde{\mathcal{O}}\left(e^{\tau(L / \mu)^{2}}\left[\frac{\left.\left.h\left(\frac{L}{\mu}\right) \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\tau^{2} L^{2}\left(1+\left\\|x^{*}\right\\|^{2}\right)}{\mu^{2} \varepsilon}\right]\right)}{}\right.\right.$ |
|  | ASGD [20] ${ }^{\text {5 }}$ | $x$ | $x$ | $x$ | $\checkmark$ | $\tilde{\mathcal{O}}\left(\frac{1}{\epsilon^{4}}\left[B^{2}+G^{6}\left(L^{2} \tau^{2}+1\right)\right]\right)$ | $\tilde{\mathcal{O}}\left(\sqrt{\frac{L}{\mu}} \frac{\left\\|x x^{0}-x^{*}\right\\|^{2}}{\varepsilon^{1 / 2}}+\frac{\tau^{2}\left(G^{2}+\mu G D+\mu L D^{2}\right)}{\mu^{2} \varepsilon}\right)$ |
|  | MAG $21{ }^{\text {a }}$ | $\checkmark$ | x | $\checkmark$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{\tau(G+L+B)^{2} G^{2}}{\varepsilon^{4}}\right)$ | $x$ |
|  | Mc sgd [28] (Sec. 5.1) ${ }^{\text {c] }}$ | $\checkmark$ | $x$ | $x$ | x | $\mathcal{O}\left(\frac{\tau\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)+\sigma^{2}\right)}{\varepsilon^{2}}+\frac{\tau\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)+\sigma^{2}\right) \sigma^{2}}{\varepsilon^{4}}\right)$ | $\mathcal{O}\left(\frac{\tau L}{\mu} \log \frac{\left(f\left(x^{0}\right)-f\left(x^{*}\right) / \mu+\sigma^{2} /(\mu L)\right.}{\varepsilon}+\frac{\frac{\tau}{2}}{\mu^{2} \varepsilon}\right)$ |
|  | Mc SgD [28] (Sec. 5.2) ${ }^{\text {(8) }}$ | $\checkmark$ | $x$ | $x$ | $x$ | $x$ | $O\left(\frac{L}{\mu} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{L \tau \sigma_{x}^{2}}{\mu^{3} \varepsilon}\right)$ |
|  | RASGD (ours) | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\tilde{\mathcal{O}}\left(\tau L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)\left[\frac{1+\delta^{2}}{\varepsilon^{2}}+\frac{\frac{\sigma}{}^{2}}{\varepsilon^{\top}}\right]\right)$ | $\tilde{\mathcal{O}}\left(\tau\left[\left(1+\delta^{2}\right) \sqrt{\frac{L}{\mu}} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right]\right)$ |

notation: $\mu$ and $L$ are as in A 1 and A $2, G=\sup _{x, z}\|\nabla F(x, z)\|$. Note that $G \geq L$ and $G^{2} \geq \sigma^{2}$ under A4. We also set $B=\sup _{x}|f(x)| ; x^{0}$ - starting point, $x^{*}$ - solution, $\mathcal{D}$ optimisation domain; $D=\sup _{x \in \mathcal{D}}\left\|x-x^{*}\right\|, \sigma$ and $\delta$-stochasticity parameters (see A 4 ; $\sigma_{*}$ - stochasticity parameter in $x^{*} ; \tau$-mixing time of the chain (see A 3 , $\varepsilon$ - accuracy of the solution, measured as $\mathbb{E}\left[\|\nabla f(x)\|^{2}\right] \lesssim \varepsilon^{2}$ for non-convex problems and $\mathbb{E}\left[\left\|x-x^{*}\right\|^{2}\right] \lesssim \varepsilon$ for the strongly convex ones. Functions $h(L / \mu)$ and $h(G, L)$ stands for an implicit dependence of the respective parameters.
${ }^{(1)}$ gives results with stepsize as a parameter, we choose it the close way as in our Corollary $1 .{ }^{(2)}$ covers more general noise setting, then just Markovian. ${ }^{(3)}$ for general state-space Markov noise the analysis of 90 requires reversibility. Parameter $q \in(1 / 2 ; 1)$ refers to the step size $\sim 1 / k^{q}$. ${ }^{(4)}$ The fluctuation terms in 19. Theorem 1,3] contain hidden dependence on the initial error and $\left\|x^{*}\right\|$ in the fluctuation terms, making the result comparison complicated. They also contain hidden factors, which are exponential in $C=\tau / \log 4$ in the notations of our paper. Moreover, the analysis of 19 requires that $F(x, z)$ is Lipschitz w.r.t. $x$ for any $z \in \mathrm{Z}$. ${ }^{(5)}$ considers Markovian noise with finite state space and a specifically decreasing step size. Moreover, in the proof of 20 Theorem 3] (equations (64) - (66)) the authors lost the factor $C^{2}$, with $C=\tau / \log 4$. The result in the table accounts for this lost factor. ${ }^{(6)}$ considers the adaptive tuning of batch size, which is oblivious to $\tau .{ }^{(7)}$ considers Markov noise with finite state space and additionally assumes that all stochastic realization $F(\cdot, Z)$ are $L$-smooth. ${ }^{(8)}$ considers Markovian noise with finite state space, $\sigma_{*}$ bounds noise only in $x^{*}$, but additionally assumes that all stochastic realization $F(\cdot, Z)$ are $L$-smooth and $\mu$-strongly convex.

## 2 Main results

Assumptions. In this paper we study the minimization problem

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{d}} f(x):=\mathbb{E}_{Z \sim \pi}[F(x, Z)], \tag{1}
\end{equation*}
$$

where the access to the function $f$ and its gradient is available only through the (unbiased) noisy oracle $F(x, Z)$ and $\nabla F(x, Z)$, respectively. In the following presentation we impose at least one of the following regularity constraint on the underlying function $f$ itself:

Table 2: This table summarizes the findings on methods for solving stochastic (strongly) monotone variational inequalities with a Lipschitz operator and (un)bounded stochasticity. The columns of the table indicate whether the authors consider variational inequalities or only certain saddle point problems, the arbitrariness of the sets, and the use of additional composite terms. The columns on stochasticity provide information on the assumptions made with respect to the stochastic operator, such as bounded variance and the Markovian noise setting. Note that with the exception of our work, all other studies assume the independent noise.

|  |  | Statement |  |  | Stochasticity |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Method | VI? | Any set? | Composite? | Unbounded? | Markovian? | Oracle complexity |
|  | SPEG 33, 42 | $\checkmark$ | $\checkmark$ | $x$ | $x$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{L}{\mu} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)^{(1)}$ |
|  | SEG 53 | $\checkmark$ | $\checkmark$ | $x$ | $x$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{B^{2}+\sigma^{2}+(B+\sigma)(1+L D)}{\sigma^{2} \varepsilon}\right)$ |
|  | SS-SEG 68.381 | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{L}{\mu} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma_{*}^{2}}{\mu^{2} \varepsilon}\right)$ |
|  | SEG 9 | $x$ | $\checkmark$ | $x$ | $x$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{L}{\mu} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)$ |
|  | DSEG 431 | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $x$ | $\mathcal{O}\left(\left[\frac{L^{2} \sigma^{2}}{\mu^{4} \varepsilon}\right]^{3}\right)^{(2)}$ |
|  | UEG 38 | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $x$ | $\mathcal{O}\left(\left(\frac{L+\Delta}{\mu}+\frac{\Delta^{2}}{\mu^{2}}\right) \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)$ |
|  | SGDA 81 | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $x$ | $\mathcal{O}\left(\frac{L^{2}}{\mu^{2}} \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right)^{(3)}$ |
|  | REG (ours) | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\tilde{\mathcal{O}}\left(\tau \cdot\left[\left(\frac{L+\Delta}{\mu}+\frac{\Delta^{2}}{\mu^{2}}\right) \log \frac{\left\\|x^{0}-x^{*}\right\\|^{2}}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right]\right)$ |
| $\begin{aligned} & 0 \\ & 0 \\ & 0 \\ & E \end{aligned}$ | SMP 52 | $\checkmark$ | $\checkmark$ | $x$ | $x$ | $x$ | $\mathcal{O}\left(\frac{L D^{2}}{\varepsilon}+\frac{\sigma^{2} \Delta^{2}}{\varepsilon^{2}}\right)$ |
|  | VR-SEG 45 | $\checkmark$ | $\checkmark$ | $x$ | $\checkmark$ | $x$ | $\mathcal{O}\left(\frac{(\sigma+\Delta)^{8} D^{4}}{\varepsilon^{2}}+\frac{D^{4}}{\varepsilon^{2}}\right)$ |
|  | IPM 46 | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $\mathcal{O}\left(\tilde{\mathcal{O}}\left(\frac{L^{4} D^{4}}{\varepsilon^{2}}+\frac{\sigma_{*}^{2} D^{4}}{\varepsilon^{2}}\right)\right)$ |
|  | SS-SEG 68 | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $\tilde{\mathcal{O}}\left(\frac{L^{2} D^{4}}{\varepsilon}+\frac{\sigma_{*}^{4}}{L^{2} \varepsilon^{2}}\right)$ |
|  | SEG 9 | $x$ | $\checkmark$ | $x$ | $x$ | $x$ | $\mathcal{O}\left(\frac{L D^{2}}{\varepsilon}+\frac{\sigma^{2} \Delta^{2}}{\varepsilon^{2}}\right)$ |
|  | REG (ours) | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\tilde{\mathcal{O}}\left(\tau \cdot\left[\frac{L D^{2}}{\varepsilon}+\frac{\sigma^{2} D^{2}}{\varepsilon^{2}}+\frac{\Delta^{2} D^{4}}{\varepsilon^{2}}\right]\right)$ |

notation: $\mu=$ constant of strong monotonicity of operator $F, L=$ Lipschitz constant of $F, B=$ uniform bound of $F, D=$ uniform bound of iterations $x^{k}, x^{0}=$ starting point, $x^{*}=$ solution, $\Delta$ and $\sigma=\operatorname{stochasticity~parameters~(see~A~7,~[52.~33,~42,~53,~9.~8]~}$ take $\Delta=0$ ), $\sigma_{*}=$ stochasticity parameter in $x^{*}$ (see 68$]$ ), $\tau=$ mixing time of the chain (see A 3$], \varepsilon=$ accuracy of the solution. ${ }^{(1)}$ give results with stepsize as a parameter, we choose it according to Section 3 from [88]. ${ }^{(2)}$ consider A 7, but do not provide explicit rates if $\Delta \neq 0$ (see also [38. Table 1]). ${ }^{(3)}$ consider the cocoercive case, for which in general $\ell=L^{2} / \mu$.

A 1. The function $f$ is $L$-smooth on $\mathbb{R}^{d}$ with $L>0$, i.e., it is differentiable and there is a constant $L>0$ such that the following inequality holds for all $x, y \in \mathbb{R}^{d}$ :

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

A 2. The function $f$ is $\mu$-strongly convex on $\mathbb{R}^{d}$, i.e., it is continuously differentiable and there is a constant $\mu>0$ such that the following inequality holds for all $x, y \in \mathbb{R}^{d}$ :

$$
\begin{equation*}
(\mu / 2)\|x-y\|^{2} \leq f(x)-f(y)-\langle\nabla f(y), x-y\rangle . \tag{2}
\end{equation*}
$$

Next we specify our assumptions on the sequence of noise variables $\left\{Z_{i}\right\}_{i=0}^{\infty}$. We consider here the general setting of $\left\{Z_{i}\right\}_{i=0}^{\infty}$ being a time-homogeneous Markov chain. Such problems naturally arise in stochastic optimization. In the empirical risk minimization problems it naturally appears in the context of non-random minibatch choice. Indeed, a random choice of a batch number may lose to a non-random one, see [67, [56]. A wide range of problems dealing with Markovian noise is spawned by the reinforcement learning methods. The usual MDP setting falls naturally inside this paradigm, moreover, the analysis of non-tabular RL problems requires to deal with the general state-space Markov noise. Here the potential range of applications include the policy evaluation methods, such as the temporal difference methods [92], and policy optimization algorithms, such as policy gradient family, e.g. the celebrated REINFORCE algorithm [100].

We denote by Q the Markov kernel corresponding to the sequence $\left\{Z_{i}\right\}_{i=0}^{\infty}$ and impose the following assumption on the mixing properties of Q :
A 3. $\left\{Z_{i}\right\}_{i=0}^{\infty}$ is a stationary Markov chain on $(Z, \mathcal{Z})$ with Markov kernel Q and unique invariant distribution $\pi$. Moreover, Q is uniformly geometrically ergodic with mixing time $\tau \in \mathbb{N}$, i.e., for every $k \in \mathbb{N}$,

$$
\begin{equation*}
\Delta\left(\mathrm{Q}^{k}\right)=\sup _{z, z^{\prime} \in \mathrm{Z}}(1 / 2)\left\|\mathrm{Q}^{k}(z, \cdot)-\mathrm{Q}^{k}\left(z^{\prime}, \cdot\right)\right\|_{\mathrm{TV}} \leq(1 / 4)^{\lfloor k / \tau\rfloor} \tag{3}
\end{equation*}
$$

The assumption A 3 is classical in the literature on optimization methods with Markovian noise and has been considered in particular in recent works [90, 21, 20]. In particular, this assumption covers finite state-space Markov chains with irreducible and aperiodic transition matrix considered in
[28]. Yet our definition of the mixing time $\tau$ is more classical in the probability literature [81], and is slightly different from the one considered e.g. in [28, 65]. Next we specify our assumptions on stochastic gradient:
A 4. For all $x \in \mathbb{R}^{d}$ it holds that $\mathbb{E}_{\pi}[\nabla F(x, Z)]=\nabla f(x)$. Moreover, for all $z \in \mathrm{Z}$ and $x \in \mathbb{R}^{d}$ it holds that

$$
\begin{equation*}
\|\nabla F(x, z)-\nabla f(x)\|^{2} \leq \sigma^{2}+\delta^{2}\|\nabla f(x)\|^{2} \tag{4}
\end{equation*}
$$

The assumption A 4 resembles the strong growth condition [97], which is classical for the overparametrized learning setup [97, 98]. The main difference is that A4 concerns the almost sure bound in (4), which is unavoidable when dealing with uniformly geometrically ergodic Markovian noise A 3 . Note that it is possible that the quantity $\delta^{2}$ in (4) is not instance-independent and scales with the ratio $L / \mu$ from A 1 -A 2 in the particular problems. With the assumptions A 3 and A 4 we can prove the result on the mean squred error of the stochastic gradient estimate computed over batch size $n$ under arbitrary initial distribution. This result is summarized below in Lemma 1
Lemma 1. Assume $A 3$ and $A 4$ Then, for any $n \geq 1$ and $x \in \mathbb{R}^{d}$, it holds that

$$
\begin{equation*}
\mathbb{E}_{\pi}\left[\left\|n^{-1} \sum_{i=1}^{n} \nabla F\left(x, Z_{i}\right)-\nabla f(x)\right\|^{2}\right] \leq \frac{8 \tau}{n}\left(\sigma^{2}+\delta^{2}\|\nabla f(x)\|^{2}\right) \tag{5}
\end{equation*}
$$

Moreover, for any initial distribution $\xi$ on $(\mathrm{Z}, \mathcal{Z})$, that

$$
\begin{equation*}
\mathbb{E}_{\xi}\left[\left\|n^{-1} \sum_{i=1}^{n} \nabla F\left(x, Z_{i}\right)-\nabla f(x)\right\|^{2}\right] \leq \frac{C_{1} \tau}{n}\left(\sigma^{2}+\delta^{2}\|\nabla f(x)\|^{2}\right) \tag{6}
\end{equation*}
$$

where $C_{1}=16\left(1+\frac{1}{\ln ^{2} 4}\right)$.
Proof. We first prove (5). Note that due to [81, Proposition 3.4] the Markov kernel Q under A 3 admits a positive pseudospectral gap $\gamma_{p s}>0$ such that $1 / \gamma_{p s} \leq 2 \tau$. Thus, applying the statement of [81, Theorem 3.2], we get under A 4 that

$$
\mathbb{E}_{\pi}\left[\left\|n^{-1} \sum_{i=1}^{n} \nabla F\left(x, Z_{i}\right)-\nabla f(x)\right\|^{2}\right] \leq \frac{4 \mathbb{E}_{\pi}\left[\left\|\nabla F\left(x, Z_{1}\right)-\nabla f(x)\right\|^{2}\right]}{n \gamma_{p s}} \leq \frac{8 \tau}{n}\left(\sigma^{2}+\delta^{2}\|\nabla f(x)\|^{2}\right)
$$

To prove the second part we use the maximal exact coupling construction and follow, e.g., [24, Theorem 1]. The complete proof is given in Appendix B. 1
The proof of Lemma 1 simplifies the arguments in [21, Lemma 4] and allows us to obtain tighter values for the constants when dealing with the randomized batch size. Note that it is especially important to have the result for MSE under arbitrary initial distribution $\xi$, since in the proofs of our main results we will inevitably deal with the conditional expectations w.r.t. the previous iterate. We provide more details on the bias and variance of the Markov SGD gradients in the next section.

### 2.1 Accelerated method

We begin with a version of Nesterov accelerated SGD with randomized batch size, described in Algorithm 1. Due to the unboundedness of the stochastic gradient variance (see A 47, using of the classical Nesterov accelerated method [76, Section 2.2.] does not give the desired result, it is necessary to introduce an additional momentum [74, 97]. We use our own version, but partially similar to [74, 97]. The main feature of Algorithm[1] is that the number of samples used during the $k$-th gradient computation scales as $2^{J_{k}}$, where $J_{k}$ comes from a truncated geometric distribution. The truncation parameter needs to be adopted (see Theorem (1) in order to control the computational complexity of the algorithm.
Randomized batch size allows for efficient bias reduction in the stochastic gradient estimates and can be seen as a particular case of the so called multilevel MCMC [35, 34]. In the optimization context this approach was successfully used by [21] for the non-convex problems. Indeed, this bias naturally appears under the Markovian stochastic gradients oracles. It is easy to see that, with the counter $T^{k}$ defined in Line 9 we have

$$
\mathbb{E}_{k}\left[\nabla F\left(x^{k}, Z_{T^{k}+i}\right)\right] \neq \nabla f\left(x^{k}\right)
$$

Below we show how the bias of the gradient estimate scales with the truncation parameter $M$. The statement of Lemma 2 yields that the gradient estimates $g_{k}$ introduced above have the bias, which decreases quadratically with $M$.
Lemma 2. Assume A 3 and A 4. Then for the gradient estimates $g^{k}$ from Algorithm 1 it holds that $\mathbb{E}_{k}\left[g^{k}\right]=\mathbb{E}_{k}\left[g_{\left[\log _{2} M\right\rfloor}^{k}\right]$. Moreover,

$$
\begin{aligned}
& \mathbb{E}_{k}\left[\left\|\nabla f\left(x^{k}\right)-g^{k}\right\|^{2}\right] \lesssim\left(\tau B^{-1} \log _{2} M+\tau^{2} B^{-2}\right)\left(\sigma^{2}+\delta^{2}\left\|\nabla f\left(x^{k}\right)\right\|^{2}\right), \\
& \left\|\nabla f\left(x^{k}\right)-\mathbb{E}_{k}\left[g^{k}\right]\right\|^{2} \lesssim \tau^{2} M^{-2} B^{-2}\left(\sigma^{2}+\delta^{2}\left\|\nabla f\left(x^{k}\right)\right\|^{2}\right)
\end{aligned}
$$

```
Algorithm 1 Randomized Accelerated GD
    Parameters: stepsize \(\gamma>0\), momentums \(\theta, \eta, \beta, p\), number of iterations \(N\), batchsize limit \(M\)
    Initialization: choose \(x^{0}=x_{f}^{0}\)
    for \(k=0,1,2, \ldots, N-1\) do
        \(x_{g}^{k}=\theta x_{f}^{k}+(1-\theta) x^{k}\)
        Sample \(J_{k} \sim \operatorname{Geom}(1 / 2)\)
        \(g^{k}=g_{0}^{k}+\left\{\begin{array}{ll}2^{J_{k}}\left(g_{J_{k}}^{k}-g_{J_{k}-1}^{k}\right), & \text { if } 2^{J_{k}} \leq M \\ 0, & \text { otherwise }\end{array} \quad\right.\) with \(\quad g_{j}^{k}=2^{-j} B^{-1} \sum_{i=1}^{2^{j} B} \nabla f\left(x_{g}^{k}, Z_{T^{k}+i}\right)\)
        \(x_{f}^{k+1}=x_{g}^{k}-p \gamma g^{k}\)
        \(x^{k+1}=\eta x_{f}^{k+1}+(p-\eta) x_{f}^{k}+(1-p)(1-\beta) x^{k}+(1-p) \beta x_{g}^{k}\)
        \(T^{k+1}=T^{k}+2^{J_{k}} B\)
    end for
```

The proof and the statement with explicit constants are given in Appendix B.2. Note that the Lemma 2 is a natural counterpart of the deterministic bound Lemma 1. Moreover, it gives the idea of the trade-off between the parameters $B$ and $M$. Namely, the expected number of oracle calls to compute $g_{k}$ is $\mathcal{O}\left(B \log _{2}(M)\right)$ with the bias scaling as $M^{-2}$. Thus the increase of $M$ drastically reduced the bias with only a logarithmic payment in variance. At the same time, gradient variance scales as $(\tau / B)^{2}$, but the increase of $B$ is much more expensive for the computational cost of the whole procedure. Taking into account the considerations above, we can prove the following result:
Theorem 1. Assume A 1 - A 4 Let problem (1) be solved by Algorithm $\mathbb{1}$. Then for any $b \in \mathbb{N}^{*}$, $\gamma \in\left(0 ; \frac{3}{4 L}\right]$, and $\beta, \theta, \eta, p, M, \vec{B}$ satisfying

$$
\begin{aligned}
& p \simeq\left(1+(1+\gamma L)\left[\delta^{2} \tau b^{-1}+\delta^{2} \tau^{2} b^{-2}\right]\right)^{-1}, \quad \beta \simeq \sqrt{p^{2} \mu \gamma}, \quad \eta \simeq \sqrt{\frac{1}{\mu \gamma}}, \\
& \theta \simeq \frac{p \eta^{-1}-1}{\beta p \eta^{-1}-1}, \quad M \simeq \max \left\{2 ; \sqrt{p^{-1}(1+p / \beta)}\right\}, \quad B=\left\lceil b \log _{2} M\right\rceil
\end{aligned}
$$

it holds that

$$
\begin{align*}
& \mathbb{E}\left[\left\|x^{N}-x^{*}\right\|^{2}+\frac{6}{\mu}\left(f\left(x_{f}^{N}\right)-f\left(x^{*}\right)\right)\right] \lesssim \exp \left(-N \sqrt{\frac{p^{2} \mu \gamma}{3}}\right) {\left[\left\|x^{0}-x^{*}\right\|^{2}+\frac{6}{\mu}\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)\right] } \\
&+\frac{p \sqrt{\gamma}}{\mu^{3 / 2}}\left(\sigma^{2} \tau b^{-1}+\sigma^{2} \tau^{2} b^{-2}\right) \tag{7}
\end{align*}
$$

The proof is provided in Appendix B.3. The result of Theorem 1 can be rewritten as an upper complexity bound under an appropriate choice of the remaining free parameter $b$ :
Corollary 1. Under the conditions of Theorem 1 choosing $b=\tau$ and $\gamma$ as

$$
\gamma \simeq \min \left\{\frac{1}{L} ; \frac{1}{p^{2} \mu N^{2}} \ln \left(\max \left\{2 ; \frac{\mu^{2} N\left[\left\|x^{0}-x^{*}\right\|^{2}+6 \mu^{-1}\left(f\left(x_{f}^{0}\right)-f\left(x^{*}\right)\right)\right]}{\sigma^{2}}\right\}\right)\right\}
$$

in order to achieve $\varepsilon$-approximate solution (in terms of $\mathbb{E}\left[\left\|x-x^{*}\right\|^{2}\right] \lesssim \varepsilon$ ) it takes

$$
\begin{equation*}
\tilde{\mathcal{O}}\left(\tau\left[\left(1+\delta^{2}\right) \sqrt{\frac{L}{\mu}} \log \frac{1}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right]\right) \quad \text { oracle calls } \tag{8}
\end{equation*}
$$

The results of Corollary 1 are obtained with fixed parameters of the method. In Corollary 1 these parameters are selected a bit artificially, e.g., the stepsize $\gamma$ depends on the iteration horizon $N$. In Appendix B. 4 we show how one can similar results, but with a decreasing stepsize.
Comparison. Running the procedure above requires to know the mixing time $\tau$. Estimating the mixing time from a single trajectory of the running Markov chain is known to be computationally hard problem, see e.g. [101] and references therein. At the same time, methods, which share the same (optimal) linear scaling of the sample complexity w.r.t. the mixing time also share the same drawback as our method. In particular, it holds true for the EMD algorithm [23], SGD-DD algorithm [71], and usual SGD with Markovian data [28]. At the same time, in the non-convex scanario the paper [21] is truly oblivious to mixing time, allowing to obtain sample complexity rates for non-convex problems, which are homogeneous w.r.t. $\tau$ with AdaGrad-type learning rate. An interesting direction for the future work to suggest a procedure that would allow to generalize the results of [21] to accelerated SGD setting.

It is possible that the sample complexity bound $(8)$ is worse than the respective bounds for nonaccelerated SGD with Markov data, provided that $\delta^{2}$ grows quickly with $L / \mu$. At the same time, this drawback is shared by the classical results on learning under the strong growth condition, see e.g. [97]. As it is shown in [62], the respective rates can be worse than the ones obtained by usual SGD even under the i.i.d. noise setting, see Appendix F. 3 in [62]. Making the analysis of accelerated SGD 'backward compatible' w.r.t. the rates of usual SGD requires to perform analysis in terms of additional problem-specific quantities, see [47, 62].

The closest equivalent of the result Corollary 11 is given by [20, Theorem 3]. However, the corresponding bound of [20, Theorem 3] is incomplete, since the factor $\tau^{2}$ is lost in the proof (see equations $(64-66)$ ). With this completion, the bound of [20, Theorem 3] yields a variance term of order $\tilde{\mathcal{O}}\left(\frac{\sigma^{2} \tau^{2}}{\mu^{2} \varepsilon}\right)$, which is suboptimal with respect to $\tau$. Moreover, the corresponding analysis relies heavily on the assumption of a bounded domain. In [28], the author considers Markovian noise with a finite number of states and manages to obtain a rather interesting result of the form $O\left(\frac{L}{\mu} \log \frac{1}{\varepsilon}+\frac{L \tau \sigma_{*}^{2}}{\mu^{3} \varepsilon}\right)$. Here the first term does not depend on $\tau$, and the second consists only $\sigma^{*}$ (stochasticity in $x^{*}$ ), but the price for this is an additional factor $L / \mu$ in the second term and more strict assumption that all realizations $F(\cdot, z)$ are smooth and strongly convex. In the context of overparameterized learning, our results are almost consistent with the bound of [97], Theorem 1] under i.i.d. sampling. The difference is that the term $\delta^{2}$ in A 4 can be more pessimistic than the expectation bound in [97].

### 2.2 Lower bounds

We start with a lower bound for the complexity of Markovian stochastic optimization under the assumptions A1-A4. Below we provide a result that highlights that the bound of Theorem 1] is tight provided that $\delta$ does not scale with the instance-dependent quantities, e.g., condition number $L / \mu$.
Theorem 2. There exists an instance of the optimization problem satisfying assumptions $A[\square-A[4$ with $\delta=1$ and arbitrary $\sigma \geq 0, L, \mu>0, \tau \in \mathbb{N}^{*}$, such that for any first-order gradient method it takes at least

$$
N=\Omega\left(\tau \sqrt{\frac{L}{\mu}} \log \frac{1}{\varepsilon}+\frac{\tau \sigma^{2}}{\mu^{2} \varepsilon}\right)
$$

oracle calls in order to achieve $\mathbb{E}\left[\left\|x^{N}-x^{*}\right\|^{2}\right] \leq \varepsilon$.
The proof is provided in Appendix B.5. The idea of the constructed lower for deterministic part bound $\Omega\left(\tau \sqrt{\frac{L}{\mu}} \log \frac{1}{\varepsilon}\right)$ goes back to [76, Theorem 2.1.13]. The stochastic part lower bound goes back to the classical statistical reasoning, and is well explained for i.i.d. noise in [59, Chapter 4.1]. Our adaptation for Markovian setting is based on Le Cam's theory, see [1]. Theorem 8], and also [105]. For the case of Markov noise this lower bound is, to the best of our knowledge, original. The closest result to ours is the stochastic term lower bound in [28, Proposition 1], but it is valid only for the vanilla stochastic gradient methods. Below we provide another lower bound showing that the dependence of the sample complexity Corollary 1 on $\delta$ is not an artefact of the proof.
Proposition 1. There exists an instance of the optimization problem satisfying assumptions $A[\square-A 4$ with arbitrary $L, \mu>0, \tau \in \mathbb{N}^{*}, \delta=\frac{L}{\mu}$, and $\sigma=0$, such that for any first-order gradient method it takes at least

$$
N=\Omega\left(\tau \frac{L}{\mu} \log \frac{1}{\varepsilon}\right)
$$

gradient calls in order to achieve $\mathbb{E}\left[\left\|x^{N}-x^{*}\right\|^{2}\right] \leq \varepsilon$.
This lower bound is adapted from the information-theoretic lower bound [71]. The detailed proof can be found in Appendix B.5. Recent studies [54, 71, 13] have revealed the impossibility of accelerating stochastic gradient descent (SGD) for online linear regression problems with specific noise structures. To address this issue, researchers have proposed various solutions, such as the MaSS algorithm [62] and the approach presented in [48]. However, these methods rely heavily on the particular structure of the online regression setup. Another question that naturally arises is whether one can get rid of the dependence on $\tau$ in the deterministic part of 8 if $\delta=0$. The following counterexample shows that this is not the case in general.
Proposition 2. There exists an instance of the optimisation problem satisfying assumptions $A$ 1]-A 4 with with arbitrary $L, \mu>0, \tau \in \mathbb{N}^{*}, \sigma=1, \delta=0$, such that for any first-order gradient method it takes at least

$$
N=\Omega\left(\left(\tau+\sqrt{\frac{L}{\mu}}\right) \log \frac{1}{\varepsilon}\right)
$$

oracle calls in order to achieve $\mathbb{E}\left[\left\|x^{N}-x^{*}\right\|^{2}\right] \leq \varepsilon$.
The proof is provided in Appendix B. 5

### 2.3 Non-convex problems

Now we proceed with a randomized batch size version of the simple SGD algorithm. It is summarized in Algorithm 2 and can be shown to achieve optimal rates of convergence for smooth non-convex problems. For the case of non-convex problems with Markov noise similar analysis appeared in [21] Theorem 4].

```
Algorithm 2 Randomized GD
    Parameters: stepsize \(\gamma>0\), number of iterations \(K\), bound on batchsize \(B\), mixing time \(\tau\);
    Initialization: choose \(x^{0} \in \mathcal{X}\)
    for \(k=0,1,2, \ldots, N-1\) do
        Sample \(J_{k} \sim \operatorname{Geom}\left(\frac{1}{2}\right)\)
        \(g^{k}=g_{0}^{k}+\left\{\begin{array}{ll}2^{J_{k}}\left(g_{J_{k}}^{k}-g_{J_{k}-1}^{k}\right), & \text { if } 2^{J_{k}} \leq M \\ 0, & \text { otherwise }\end{array} \quad\right.\) with \(\quad g_{j}^{k}=2^{-j} B^{-1} \sum_{i=1}^{2^{j} B} \nabla f\left(x^{k}, Z_{T^{k}+i}\right)\)
        \(x^{k+1}=x^{k}-\gamma g^{k}\)
        \(T^{k+1}=T^{k}+2^{J_{k}} B\)
    end for
```

By balancing the values of $B$ and $M$ with Lemma 2, we establish the following result:
Theorem 3. Assume A 1] A 3] A 4. Let problem (1) be solved by Algorithm 2. Let $f^{*}$ be a global (maybe not unique) minimum of $f$. Then for any $b \in \mathbb{N}^{*}$, and $\gamma, M$ satisfying

$$
\gamma \lesssim\left(L\left[1+\delta^{2} \tau b^{-1}+\delta^{2} \tau^{2} b^{-2}\right]\right)^{-1}, \quad M \simeq \max \left\{2 ; \sqrt{\gamma^{-1} L^{-1}}\right\}, \quad B=\left\lceil b \log _{2} M\right\rceil
$$

it holds that

$$
\mathbb{E}\left[\frac{1}{N} \sum_{k=0}^{N-1}\left\|\nabla f\left(x^{k}\right)\right\|^{2}\right] \lesssim \frac{f\left(x^{0}\right)-f^{*}}{\gamma N}+L \gamma \cdot\left[\sigma^{2} \tau b^{-1}+\sigma^{2} \tau^{2} b^{-2}\right]
$$

The proof is provided in Appendix B.6 The next corollary immediately follows from the theorem.
Corollary 2. Under the conditions of Theorem 3 if we choose $b=\tau$ and $\gamma$ given by

$$
\gamma \simeq \min \left\{\frac{1}{L\left(1+\delta^{2}\right)} ; \sqrt{\frac{f\left(x^{0}\right)-f^{*}}{L N \sigma^{2}}}\right\}
$$

then to achieve $\varepsilon$-solution (in terms of $\mathbb{E}\left[\|\nabla f(x)\|^{2}\right] \lesssim \varepsilon^{2}$ ) we need

$$
\tilde{\mathcal{O}}\left(\tau \cdot\left[\frac{\left(1+\delta^{2}\right) L\left(f\left(x^{0}\right)-f^{*}\right)}{\varepsilon^{2}}+\frac{L\left(f\left(x^{0}\right)-f^{*}\right) \sigma^{2}}{\varepsilon^{4}}\right]\right) \quad \text { oracle calls. }
$$

Comparison. The respective bound for the non-convex setting provided in [20, Theorem 1] yields the sample complexity of order $\tilde{\mathcal{O}}\left(\frac{\tau^{2} L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right) \sigma^{2}}{\varepsilon^{4}}\right)$. Also we can note the results of [28], Theorem 2] with the following estimate $O\left(\frac{\tau\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)+\sigma^{2}\right)}{\varepsilon^{2}}+\frac{\tau\left(L\left(f\left(x^{0}\right)-f\left(x^{*}\right)\right)+\sigma^{2}\right) \sigma^{2}}{\varepsilon^{4}}\right)$.
To achieve linear convergence rates in the non-convex setting we can use the Polyak-Lojasiewicz (PL) condition [83]. The respective result is provided in Appendix B.7.

### 2.4 Variational inequalities

In this section, we are interested in the following problem:

$$
\begin{equation*}
\text { Find } x^{*} \in \mathcal{X} \text { such that }\left\langle F\left(x^{*}\right), x-x^{*}\right\rangle+r(x)-r\left(x^{*}\right) \geq 0 \text { for all } x \in \mathcal{X} \tag{9}
\end{equation*}
$$

Here $F: \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ an operator, $\mathcal{X}$ a convex set, and $r: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is a regularization term (a suitable lower semicontinuous convex function) which is assumed to have a simple structure. As mentioned earlier, this problem is quite general and covers a wide range of possible problem formulations. For example, if the operator $F$ is the gradient of a convex function $f$, then the problem ( 9 ) is equivalent to the composite minimization problem [6], i.e., minimization of $f(x)+r(x)$. In the meantime, (9) is also a reformulation of the min-max problem

$$
\begin{equation*}
\min _{x_{1} \in \mathcal{X}_{1}} \max _{x_{2} \in \mathcal{X}_{2}} r_{1}\left(x_{1}\right)+g\left(x_{1}, x_{2}\right)-r_{2}\left(x_{2}\right) \tag{10}
\end{equation*}
$$

with convex-concave continuously differentiable $g$, convex sets $\mathcal{X}_{1}, \mathcal{X}_{2}$ and convex functions $r_{1}$, $r_{2}$. Using the first-order optimality conditions, it is easy to verify that (10) is equivalent to (9) with $x=\left(x_{1}^{T}, x_{2}^{T}\right)^{T}, F(x)=\left(\nabla_{x_{1}} f\left(x_{1}, x_{2}\right)^{T},-\nabla_{x_{2}} f\left(x_{1}, x_{2}\right)^{T}\right)^{T}$, and $r(x)=r_{1}\left(x_{1}\right)+r_{2}\left(x_{2}\right)$.

A 5. The operator $F$ is L-Lipschitz continuous on $\mathcal{X}$ with $L>0$, i.e., the following inequality holds for all $x, y \in \mathcal{X}$ :

$$
\|F(x)-F(y)\| \leq L\|x-y\|, \quad \forall x, y \in \mathcal{X}
$$

A 6. The operator $F$ is $\mu_{F}$-strongly monotone on $\mathcal{X}$, i.e., the following inequality holds for all $x, y \in \mathcal{X}$ :

$$
\begin{equation*}
\langle F(x)-F(y), x-y\rangle \geq \mu_{F}\|x-y\|^{2} . \tag{11}
\end{equation*}
$$

The function $r$ is $\mu_{r}$-strongly convex on $\mathcal{X}$, i.e. for all $x, y \in \mathcal{X}$ and any $r^{\prime}(x) \in \partial r(x)$ we have

$$
\begin{equation*}
r(y) \geq r(x)+\left\langle r^{\prime}(x), y-x\right\rangle+\left(\mu_{r} / 2\right)\|x-y\|^{2} . \tag{12}
\end{equation*}
$$

These two assumptions are more than standard for the study of variational inequalities and are found in all the papers from Table 2 We consider two cases: strongly monotone/convex with $\mu_{F}+\mu_{r}>0$ and monotone/convex with $\mu_{F}+\mu_{r}=0$.
A 7. For all $x \in \mathbb{R}^{d}$ it holds that $\mathbb{E}_{\pi}[F(x, Z)]=F(x)$. Moreover, for all $z \in \mathrm{Z}$ and $x \in \mathcal{X}$ it holds that

$$
\begin{equation*}
\|F(x, z)-F(x)\|^{2} \leq \sigma^{2}+\Delta^{2}\left\|x-x^{*}\right\|^{2}, \tag{13}
\end{equation*}
$$

where $x^{*}$ is some point from the solution set.
A 7 is found in the literature on variational inequalities [43, 45, 38] and is considered to be analog to A4 on overparametrized learning.
Just as the Nesterov accelerated method is optimal for smooth convex minimization problems, the ExtraGradient method [57, 72, 52] is optimal for monotone variational inequalities. Therefore, we take it as a base. On the extrapolation step (Line 4) of Algorithm 3, we simply collect a batch of size $B$, but on the main step (Line 8) we use the randomization as in Algorithm 1. The next theorem gives the convergence of our method.
Theorem 4. Assume A 5 A 6 with $\mu_{F}+\mu_{r}>0, A 3$ A 7 Let problem (9) be solved by Algorithm 3 Then for any $b \in \mathbb{N}^{*}$, and $\gamma, M$ satisfying

$$
\begin{gathered}
\gamma \lesssim \min \left\{\left(\mu_{F}+\mu_{r}\right)^{-1} ; L^{-1} ;\left(\mu_{F}+\mu_{r}\right)\left(\Delta^{2} \tau b^{-1}+\Delta^{2} \tau^{2} b^{-2}\right)^{-1} ; \sqrt{\Delta^{-2} \tau^{-1} b}\right\} \\
M \simeq \max \left\{2 ; \sqrt{\gamma^{-1}\left(\mu_{F}+\mu_{r}\right)^{-1}}\right\}, \quad B=\left\lceil b \log _{2} M\right\rceil
\end{gathered}
$$

it holds that

$$
\mathbb{E}\left[\left\|x^{N}-x^{*}\right\|^{2}\right] \lesssim \exp \left(-\frac{N\left(\mu_{F}+\mu_{r}\right) \gamma}{2}\right)\left\|x^{0}-x^{*}\right\|^{2}+\frac{\gamma}{\mu}\left(\sigma^{2} \tau b^{-1}+\sigma^{2} \tau^{2} b^{-2}\right)
$$

The proof is postponed to Appendix B.8. One can get an estimate on oracle complexity.

```
Algorithm 3 Randomized ExtraGradient
    Parameters: stepsize \(\gamma>0\), number of iterations \(N\)
    Initialization: choose \(x^{0} \in \mathcal{X}\)
    for \(k=0,1,2, \ldots, N-1\) do
        \(x^{k+1 / 2}=\operatorname{prox}_{\gamma r}\left(x^{k}-\gamma B^{-1} \sum_{i=1}^{B} F\left(x^{k}, Z_{T^{k}+i}\right)\right)\)
        \(T^{k+1 / 2}=T^{k}+B\)
        Sample \(J_{k} \sim \operatorname{Geom}\left(\frac{1}{2}\right)\)
            \(g^{k}=g_{0}^{k}+\left\{\begin{array}{ll}2^{J_{k}}\left(g_{J_{k}}^{k}-g_{J_{k}-1}^{k}\right), & \text { if } 2^{J_{k}} \leq M \\ 0, & \text { otherwise }\end{array} \quad\right.\) with \(\quad g_{j}^{k}=2^{-j} B^{-1} \sum_{i=1}^{2^{j} \cdot B} F\left(x^{k+1 / 2}, Z_{T^{k+1 / 2}+i}\right)\)
        \(x^{k+1}=\operatorname{prox}_{\gamma r}\left(x^{k}-\gamma g^{k}\right)\)
        \(T^{k+1}=T^{k+1 / 2}+2^{J_{k}} B\)
    end for
```

Corollary 3. Under the conditions of Theorem 4 if we choose $b=\tau$ and $\gamma$ as follows

$$
\gamma \simeq \min \left\{\frac{1}{\mu_{F}+\mu_{r}} ; \frac{1}{L} ; \frac{\mu_{F}+\mu_{r}}{\Delta^{2}} ; \frac{1}{\Delta} ; \frac{1}{N\left(\mu_{F}+\mu_{r}\right)} \ln \left(\max \left\{2 ; \frac{\mu N\left\|x^{0}-x^{*}\right\|^{2}}{\sigma^{2}}\right\}\right)\right\}
$$

then to achieve $\varepsilon$-solution (in terms of $\mathbb{E}\left[\left\|x-x^{*}\right\|^{2}\right] \lesssim \varepsilon$ ) we need

$$
\tilde{\mathcal{O}}\left(\tau \cdot\left[\left(1+\frac{L}{\mu_{F}+\mu_{r}}+\frac{\Delta}{\mu_{F}+\mu_{r}}+\frac{\Delta^{2}}{\left(\mu_{F}+\mu_{r}\right)^{2}}\right) \log \frac{1}{\varepsilon}+\frac{\sigma^{2}}{\mu^{2} \varepsilon}\right]\right) \quad \text { oracle calls. }
$$

Note that one provide an (almost) matching lower complexity bounds for variational inequalities via lower bounds for saddle point problems, which are a special case of variational inequalities. The method for obtaining lower bounds for saddle point problems is reduced to obtaining estimates for the strongly convex minimization problem (see [106, 40] for respective deterministic lower bounds), which we provide in Section 2.2. Similarly, the question of constructing a lower bound which is tight w.r.t. $\Delta$ remains open.

For the monotone case, we use the gap function as a convergence criterion:

$$
\begin{equation*}
\operatorname{Gap}(x)=\sup _{y \in \mathcal{X}}[\langle F(y), x-y\rangle+r(x)-r(y)] . \tag{14}
\end{equation*}
$$

Such a criterion is standard and classical for monotone variational inequalities [72, 52]. An important assumption for the gap function is the boundedness of the set $\mathcal{X}$.
A 8. The set $\mathcal{X}$ is bounded and has a diameter $D$, i.e., for all $x, y \in \mathcal{X}:\|x-y\|^{2} \leq D^{2}$.
A 8 can be slightly relaxed. We need to use a simple trick from [77]. In particular, we need to consider $\mathcal{C}$ - a compact subset of $\mathcal{X}$ and change $\mathcal{X}$ to $\mathcal{C}$ in (14). But such a technique is rather technical and does not change the essence. Finally, the following result holds.
Theorem 5. Assume A 5] A 6 with $\mu_{F}+\mu_{r}=0, A 8$ A 3, A 7 Let problem (9) be solved by Algorithm 3 Then for any $B \in \mathbb{N}^{*}$, and $\gamma, M$ satisfying $\gamma \lesssim L^{-1}, M=\sqrt{N}$, it holds that

$$
\mathbb{E}\left[\operatorname{Gap}\left(\bar{x}^{N}\right)\right] \lesssim \frac{D^{2}}{\gamma N}+\gamma\left(\tau B^{-1} \log _{2} N+\tau^{2} B^{-2}\right)\left(\sigma^{2}+\Delta^{2} D^{2}\right) \text { where } \bar{x}^{N}=\frac{1}{N} \sum_{k=0}^{N-1} x^{k+1 / 2}
$$

The proof is postponed to Appendix B.9. The following corollary holds.
Corollary 4. Under the conditions of Theorem 5 if we choose $B=\tau$ and $\gamma$ as follows

$$
\gamma \simeq \min \left\{\frac{1}{L} ; \sqrt{\frac{D^{2}}{\left(\sigma^{2}+\Delta^{2} D^{2}\right) N}}\right\}
$$

then to achieve $\varepsilon$-solution (in terms of $\mathbb{E}[\operatorname{Gap}(x)] \lesssim \varepsilon$ ) we need

$$
\tilde{\mathcal{O}}\left(\tau\left[\frac{L D^{2}}{\varepsilon}+\frac{\sigma^{2} D^{2}+\Delta^{2} D^{4}}{\varepsilon^{2}}\right]\right) \quad \text { oracle calls }
$$

Comparison. These results is the first for variational inequalities with Markovian stochasticity, either in the strongly monotone or monotone cases. The only close work is [99]. The authors work with convex-concave saddle point problems and provide the following estimate on the oracle complexity $\mathcal{O}\left(\tau^{2} \cdot \frac{G^{4}}{\varepsilon^{2}}+\frac{D^{2}}{\varepsilon^{2}}\right)$ (with $G$ - the uniform bound of the operator), which is worse than ours at least in terms of $\tau$. Moreover, the authors consider the case of a finite Markov chain, which is a special case of our setup.

## 3 Conclusion

In this paper, we present a unified random batch size framework that achieves optimal finite-time performance for non-convex and strongly convex optimization problems with Markov noise, as well as for variational inequalities. Unlike existing methods, our framework relaxes the assumptions typically imposed on the domain and stochastic gradient oracle. We also provide a variety of lower bounds, which are to the best of our knowledge original in the Markov setting.

## Acknowledgments

This research of A. Beznosikov has been supported by The Analytical Center for the Government of the Russian Federation (Agreement No. 70-2021-00143 dd. 01.11.2021, IGK 000000D730321P5Q0002). E. Moulines received support from the grant ANR-19-CHIA-002 SCAI and parts of his work has been done under the auspices of Lagrange Center for maths and computing.

## References

[1] Eddie Aamari and Clément Levrard. Nonasymptotic rates for manifold, tangent space and curvature estimation. The Annals of Statistics, 47(1):177-204, 2019.
[2] Ahmet Alacaoglu and Yura Malitsky. Stochastic variance reduction for variational inequality methods. In Conference on Learning Theory, pages 778-816. PMLR, 2022.
[3] Necdet Serhat Aybat, Alireza Fallah, Mert Gurbuzbalaban, and Asuman Ozdaglar. A universally optimal multistage accelerated stochastic gradient method. Advances in neural information processing systems, 32, 2019.
[4] Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. arXiv preprint arXiv:1108.0775, 2011.
[5] Francis Bach, Julien Mairal, and Jean Ponce. Convex sparse matrix factorizations. arXiv preprint arXiv:0812.1869, 2008.
[6] Amir Beck. First-order methods in optimization. Society for Industrial and Applied Mathematics (SIAM), 2017.
[7] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. Robust Optimization. Princeton University Press, 2009.
[8] Aleksandr Beznosikov, Eduard Gorbunov, Hugo Berard, and Nicolas Loizou. Stochastic gradient descent-ascent: Unified theory and new efficient methods. In International Conference on Artificial Intelligence and Statistics, pages 172-235. PMLR, 2023.
[9] Aleksandr Beznosikov, Valentin Samokhin, and Alexander Gasnikov. Distributed saddle-point problems: Lower bounds, optimal and robust algorithms. arXiv preprint arXiv:2010.13112, 2020.
[10] Jalaj Bhandari, Daniel Russo, and Raghav Singal. A finite time analysis of temporal difference learning with linear function approximation. In Conference on learning theory, pages 16911692. PMLR, 2018.
[11] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. Journal of mathematical imaging and vision, 40(1):120-145, 2011.
[12] Tatjana Chavdarova, Gauthier Gidel, François Fleuret, and Simon Lacoste-Julien. Reducing noise in gan training with variance reduced extragradient. Advances in Neural Information Processing Systems, 32, 2019.
[13] You-Lin Chen, Sen Na, and Mladen Kolar. Convergence analysis of accelerated stochastic gradient descent under the growth condition. arXiv preprint arXiv:2006.06782, 2020.
[14] Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. Advances in neural information processing systems, 24, 2011.
[15] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. arXiv preprint arXiv:1711.00141, 2017.
[16] Olivier Devolder et al. Stochastic first order methods in smooth convex optimization. Technical report, CORE, 2011.
[17] Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger convergence rates for least-squares regression. The Journal of Machine Learning Research, 18(1):3520-3570, 2017.
[18] Alexandros G. Dimakis, Soummya Kar, José M. F. Moura, Michael G. Rabbat, and Anna Scaglione. Gossip algorithms for distributed signal processing. Proceedings of the IEEE, 98(11):1847-1864, 2010.
[19] Thinh T. Doan. Finite-time analysis of markov gradient descent. IEEE Transactions on Automatic Control, 68(4):2140-2153, 2023.
[20] Thinh T Doan, Lam M Nguyen, Nhan H Pham, and Justin Romberg. Convergence rates of accelerated markov gradient descent with applications in reinforcement learning. arXiv preprint arXiv:2002.02873, 2020.
[21] Ron Dorfman and Kfir Yehuda Levy. Adapting to mixing time in stochastic optimization with markovian data. In International Conference on Machine Learning, pages 5429-5446. PMLR, 2022.
[22] R. Douc, E. Moulines, P. Priouret, and P. Soulier. Markov chains. Springer Series in Operations Research and Financial Engineering. Springer, 2018.
[23] John C Duchi, Alekh Agarwal, Mikael Johansson, and Michael I Jordan. Ergodic mirror descent. SIAM Journal on Optimization, 22(4):1549-1578, 2012.
[24] Alain Durmus, Eric Moulines, Alexey Naumov, Sergey Samsonov, and Marina Sheshukova. Rosenthal-type inequalities for linear statistics of markov chains. arXiv preprint arXiv:2303.05838, 2023.
[25] Alain Durmus, Eric Moulines, Alexey Naumov, Sergey Samsonov, and Hoi-To Wai. On the stability of random matrix product with markovian noise: Application to linear stochastic approximation and td learning. In Conference on Learning Theory, pages 1711-1752. PMLR, 2021.
[26] Pavel Dvurechensky and Alexander Gasnikov. Stochastic intermediate gradient method for convex problems with stochastic inexact oracle. Journal of Optimization Theory and Applications, 171:121-145, 2016.
[27] Ernie Esser, Xiaoqun Zhang, and Tony F Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. SIAM Journal on Imaging Sciences, 3(4):1015-1046, 2010.
[28] Mathieu Even. Stochastic gradient descent under Markovian sampling schemes. arXiv preprint arXiv:2302.14428, 2023.
[29] F. Facchinei and J.S. Pang. Finite-Dimensional Variational Inequalities and Complementarity Problems. Springer Series in Operations Research and Financial Engineering. Springer New York, 2007.
[30] Alexander Vladimirovich Gasnikov and Yu E Nesterov. Universal method for stochastic composite optimization problems. Computational Mathematics and Mathematical Physics, 58:48-64, 2018.
[31] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. SIAM Journal on Optimization, 23(4):2341-2368, 2013.
[32] Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. Mathematical Programming, 155(1-2):267-305, 2016.
[33] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. arXiv preprint arXiv:1802.10551, 2018.
[34] Michael B. Giles. Multilevel monte carlo path simulation. Operations Research, 56(3):607617, 2008.
[35] Peter W Glynn and Chang-han Rhee. Exact estimation for markov chain equilibrium expectations. Journal of Applied Probability, 51(A):377-389, 2014.
[36] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in Neural Information Processing Systems (NIPS), 2014.
[37] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, Cambridge, MA, USA, 2016. http://www.deeplearningbook.org
[38] Eduard Gorbunov, Hugo Berard, Gauthier Gidel, and Nicolas Loizou. Stochastic extragradient: General analysis and improved rates. In International Conference on Artificial Intelligence and Statistics, pages 7865-7901. PMLR, 2022.
[39] Eduard Gorbunov, Marina Danilova, Innokentiy Shibaev, Pavel Dvurechensky, and Alexander Gasnikov. Near-optimal high probability complexity bounds for non-smooth stochastic optimization with heavy-tailed noise. arXiv preprint arXiv:2106.05958, 2021.
[40] Yuze Han, Guangzeng Xie, and Zhihua Zhang. Lower complexity bounds of finite-sum optimization problems: The results and construction. arXiv preprint arXiv:2103.08280, 2021.
[41] P. T. Harker and J.-S. Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. Mathematical programming, 1990.
[42] Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. On the convergence of single-call stochastic extra-gradient methods. Advances in Neural Information Processing Systems, 32, 2019.
[43] Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling. Advances in Neural Information Processing Systems, 33:16223-16234, 2020.
[44] Chonghai Hu, Weike Pan, and James Kwok. Accelerated gradient methods for stochastic optimization and online learning. Advances in Neural Information Processing Systems, 22, 2009.
[45] A. N. Iusem, A. Jofré, R. I. Oliveira, and P. Thompson. Extragradient method with variance reduction for stochastic variational inequalities. SIAM Journal on Optimization, 27(2):686-724, 2017.
[46] Alfredo N Iusem, Alejandro Jofré, and Philip Thompson. Incremental constraint projection methods for monotone stochastic variational inequalities. Mathematics of Operations Research, 44(1):236-263, 2019.
[47] Prateek Jain, Sham M Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Accelerating stochastic gradient descent for least squares regression. In Conference On Learning Theory, pages 545-604. PMLR, 2018.
[48] Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Parallelizing stochastic gradient descent for least squares regression: Mini-batching, averaging, and model misspecification. Journal of Machine Learning Research, 18(223):1-42, 2018.
[49] Houyuan Jiang and Huifu Xu. Stochastic approximation approaches to the stochastic variational inequality problem. IEEE Transactions on Automatic Control, 53(6):1462-1475, 2008.
[50] Yujia Jin and Aaron Sidford. Efficiently solving MDPs with stochastic mirror descent. In Proceedings of the 37th International Conference on Machine Learning (ICML), volume 119, pages 4890-4900. PMLR, 2020.
[51] Thorsten Joachims. A support vector method for multivariate performance measures. pages 377-384, 012005.
[52] Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. Stochastic Systems, 1(1):17-58, 2011.
[53] Aswin Kannan and Uday V Shanbhag. Optimal stochastic extragradient schemes for pseudomonotone stochastic variational inequality problems and their variants. Computational Optimization and Applications, 74(3):779-820, 2019.
[54] 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.
[55] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
[56] Anastasia Koloskova, Nikita Doikov, Sebastian U Stich, and Martin Jaggi. Shuffle SGD is always better than SGD: Improved analysis of SGD with arbitrary data orders. arXiv preprint arXiv:2305.19259, 2023.
[57] G. M. Korpelevich. The extragradient method for finding saddle points and other problems. Matecon, 12:35-49, 1977.
[58] Guanghui Lan. An optimal method for stochastic composite optimization. Mathematical Programming, 133(1-2):365-397, 2012.
[59] Guanghui Lan. First-order and Stochastic Optimization Methods for Machine Learning. 01 2020.
[60] Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In Kamalika Chaudhuri and Masashi Sugiyama, editors, Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, volume 89 of Proceedings of Machine Learning Research, pages 907-915. PMLR, 16-18 Apr 2019.
[61] Qihang Lin, Xi Chen, and Javier Pena. A smoothing stochastic gradient method for composite optimization. Optimization Methods and Software, 29(6):1281-1301, 2014.
[62] Chaoyue Liu and Mikhail Belkin. Accelerating sgd with momentum for over-parameterized learning. arXiv preprint arXiv:1810.13395, 2018.
[63] Cassio G. Lopes and Ali H. Sayed. Incremental adaptive strategies over distributed networks. IEEE Transactions on Signal Processing, 55(8):4064-4077, 2007.
[64] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In International Conference on Learning Representations (ICLR), 2018.
[65] Xianghui Mao, Kun Yuan, Yubin Hu, Yuantao Gu, Ali H. Sayed, and Wotao Yin. Walkman: A communication-efficient random-walk algorithm for decentralized optimization. IEEE Transactions on Signal Processing, 68:2513-2528, 2020.
[66] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. arXiv preprint arXiv:1807.02629, 2018.
[67] Konstantin Mishchenko, Ahmed Khaled, and Peter Richtárik. Random reshuffling: Simple analysis with vast improvements. Advances in Neural Information Processing Systems, 33:17309-17320, 2020.
[68] Konstantin Mishchenko, Dmitry Kovalev, Egor Shulgin, Peter Richtárik, and Yura Malitsky. Revisiting stochastic extragradient. In International Conference on Artificial Intelligence and Statistics, pages 4573-4582. PMLR, 2020.
[69] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. Nature, 518(7540):529-533, 2015.
[70] Eric Moulines and Francis Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 24. Curran Associates, Inc., 2011.
[71] Dheeraj Nagaraj, Xian Wu, Guy Bresler, Prateek Jain, and Praneeth Netrapalli. Least squares regression with markovian data: Fundamental limits and algorithms. Advances in neural information processing systems, 33:16666-16676, 2020.
[72] Arkadi Nemirovski. Prox-method with rate of convergence $O(1 / t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. SIAM Journal on Optimization, 15(1):229-251, 2004.
[73] Yu Nesterov. Smooth minimization of non-smooth functions. Mathematical programming, 103(1):127-152, 2005.
[74] Yu. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. SIAM Journal on Optimization, 22(2):341-362, 2012.
[75] Yu. E. Nesterov. A method for solving the convex programming problem with convergence rate $O\left(1 / k^{2}\right)$. Dokl. Akad. Nauk SSSR, 269(3):543-547, 1983.
[76] Yurii Nesterov. Introductory lectures on convex optimization: A basic course, volume 87 Springer Science \& Business Media, 2003.
[77] Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. Mathematical Programming, 109(2):319-344, 2007.
[78] J. Von Neumann and O. Morgenstern. Theory of games and economic behavior. Princeton University Press, 1944.
[79] Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P. How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In Proceedings of the 34th International Conference on Machine Learning (ICML), volume 70, pages 2681-2690. PMLR, 2017.
[80] Balamurugan Palaniappan and Francis Bach. Stochastic variance reduction methods for saddle-point problems. Advances in Neural Information Processing Systems, 29, 2016.
[81] Daniel Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. Electronic Journal of Probability, 20(none):1-32, 2015.
[82] Wei Peng, Yu-Hong Dai, Hui Zhang, and Lizhi Cheng. Training gans with centripetal acceleration. Optimization Methods and Software, 35(5):955-973, 2020.
[83] B.T. Polyak. Gradient methods for the minimisation of functionals. USSR Computational Mathematics and Mathematical Physics, 3(4):864-878, 1963.
[84] Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. The Annals of Mathematical Statistics, 22(3):400 - 407, 1951.
[85] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In International conference on machine learning, pages 1889-1897. PMLR, 2015.
[86] Gesualdo Scutari, Daniel Palomar, Francisco Facchinei, and Jong-shi Pang. Convex optimization, game theory, and variational inequality theory. Signal Processing Magazine, IEEE, 27:35 -49, 062010.
[87] Rayadurgam Srikant and Lei Ying. Finite-time error bounds for linear stochastic approximation andtd learning. In Conference on Learning Theory, pages 2803-2830. PMLR, 2019.
[88] Sebastian U Stich. Unified optimal analysis of the (stochastic) gradient method. arXiv preprint arXiv:1907.04232, 2019.
[89] Tao Sun, Dongsheng Li, and Bao Wang. Adaptive Random Walk Gradient Descent for Decentralized Optimization. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pages 20790-20809. PMLR, 17-23 Jul 2022.
[90] Tao Sun, Yuejiao Sun, and Wotao Yin. On Markov chain gradient descent. Advances in neural information processing systems, 31, 2018.
[91] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In Sanjoy Dasgupta and David McAllester, editors, Proceedings of the 30th International Conference on Machine Learning, volume 28 of Proceedings of Machine Learning Research, pages 1139-1147, Atlanta, Georgia, USA, 17-19 Jun 2013. PMLR.
[92] Richard S Sutton. Learning to predict by the methods of temporal differences. Machine learning, 3:9-44, 1988.
[93] Richard S. Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction. The MIT Press, second edition, 2018.
[94] Adrien Taylor and Francis Bach. Stochastic first-order methods: non-asymptotic and computeraided analyses via potential functions. In Conference on Learning Theory, pages 2934-2992. PMLR, 2019.
[95] Paul Tseng. A modified forward-backward splitting method for maximal monotone mappings. SIAM Journal on Control and Optimization, 38(2):431-446, 2000.
[96] Aad W Van der Vaart. Asymptotic statistics, volume 3. Cambridge university press, 2000.
[97] Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In The 22nd international conference on artificial intelligence and statistics, pages 1195-1204. PMLR, 2019.
[98] Sharan Vaswani, Aaron Mishkin, Issam H. Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates. CoRR, abs/1905.09997, 2019.
[99] Puyu Wang, Yunwen Lei, Yiming Ying, and Ding-Xuan Zhou. Stability and generalization for markov chain stochastic gradient methods. arXiv preprint arXiv:2209.08005, 2022.
[100] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine learning, 8:229-256, 1992.
[101] Geoffrey Wolfer and Aryeh Kontorovich. Estimating the mixing time of ergodic markov chains. In Conference on Learning Theory, pages 3120-3159. PMLR, 2019.
[102] Blake E Woodworth and Nathan Srebro. An even more optimal stochastic optimization algorithm: minibatching and interpolation learning. Advances in Neural Information Processing Systems, 34:7333-7345, 2021.
[103] Linli Xu, James Neufeld, Bryce Larson, and Dale Schuurmans. Maximum margin clustering. In L. Saul, Y. Weiss, and L. Bottou, editors, Advances in Neural Information Processing Systems, volume 17. MIT Press, 2005.
[104] Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 1153-1165. Curran Associates, Inc., 2020.
[105] Bin Yu. Assouad, Fano, and Le Cam, pages 423-435. Springer New York, New York, NY, 1997.
[106] Junyu Zhang, Mingyi Hong, and Shuzhong Zhang. On lower iteration complexity bounds for the saddle point problems. arXiv preprint arXiv:1912.07481, 2019.

