



HAL
open science

SGD with Coordinate Sampling: Theory and Practice

Rémi Leluc, François Portier

► **To cite this version:**

Rémi Leluc, François Portier. SGD with Coordinate Sampling: Theory and Practice. Journal of Machine Learning Research, 2022, 23, (342), pp.1-47. hal-04044494

HAL Id: hal-04044494

<https://telecom-paris.hal.science/hal-04044494v1>

Submitted on 24 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SGD with Coordinate Sampling: Theory and Practice

Rémi Leluc

REMI.LELUC@GMAIL.COM

*Department of Statistics, LTCI
Télécom Paris, Institut Polytechnique de Paris
91120 Palaiseau, France*

François Portier

FRANCOIS.PORTIER@GMAIL.COM

*Department of Statistics, CREST
Ecole Nationale de la Statistique et de l'Analyse de l'Information (ENSAI)
35170 Bruz, France*

Editor: Karthik Sridharan

Abstract

While classical forms of stochastic gradient descent algorithm treat the different coordinates in the same way, a framework allowing for adaptive (*non uniform*) coordinate sampling is developed to leverage structure in data. In a non-convex setting and including zeroth-order gradient estimate, almost sure convergence as well as non-asymptotic bounds are established. Within the proposed framework, we develop an algorithm, MUSKETEER, based on a reinforcement strategy: after collecting information on the noisy gradients, it samples the most promising coordinate (*all for one*); then it moves along the one direction yielding an important decrease of the objective (*one for all*). Numerical experiments on both synthetic and real data examples confirm the effectiveness of MUSKETEER in large scale problems.

Keywords: stochastic optimization, stochastic gradient algorithms, zeroth-order optimization, coordinate descent, adaptive methods.

1. Introduction

Coordinate Descent (CD) algorithms have become unavoidable in modern machine learning because they are tractable (Nesterov, 2012) and competitive to other methods when dealing with key problems such as support vector machines, logistic regression, LASSO regression and other ℓ_1 -regularized learning problems (Wu et al., 2008; Friedman et al., 2010). They are applied in a wide variety of problems ranging from linear systems (Lee and Sidford, 2013; Beck and Tetrushvili, 2013) to finite sum optimization (Necoara et al., 2014; Lu and Xiao, 2015) and composite functions (Richtárik and Takáč, 2014) with parallel (Fercoq and Richtárik, 2015; Richtárik and Takáč, 2016b), distributed (Fercoq et al., 2014; Qu et al., 2015) and dual (Shalev-Shwartz and Zhang, 2013; Csiba et al., 2015; Perekrestenko et al., 2017) variants. In many contributions (Loshchilov et al., 2011; Richtárik and Takáč, 2016a; Glasmachers and Dogan, 2013; Qu and Richtárik, 2016; Allen-Zhu et al., 2016; Namkoong et al., 2017), the choice of the coordinate sampling policy is conducted through some optimality criterion estimated along the algorithm. On the one hand, efficient forms of CD methods rely on a deterministic procedure (Nutini et al., 2015) which adapts to the underlying structure in data at the expense of higher calculation and thus, may be costly. On the other hand, stochastic gradient descent (SGD) methods are computationally efficient but often treat all coordinates

equally and thus, may be sub-optimal. In the spirit of adaptive schemes, we tend to bridge the gap between the best of both worlds by developing, within a noisy gradient framework, a general stochastic coordinate descent method with a particular selection strategy.

We are interested in solving unconstrained optimization problems of the form $\min_{\theta \in \mathbb{R}^p} f(\theta)$, where the objective function f may be either known exactly or accessed through noisy observations. When f is differentiable, a common approach is to rely on the gradient of f . However, in many scenarios and particularly in large-scale learning, the gradient may be hard to evaluate or even intractable. Hence, one usually approximates the gradient using zeroth or first order estimates (Ghadimi and Lan, 2013; Lian et al., 2016). The former constructs pseudo-gradients by sampling some perturbed points or using finite differences (Flaxman et al., 2005; Duchi et al., 2012; Nesterov and Spokoiny, 2017; Shamir, 2017) (see Liu et al. (2020) for a recent survey and numerous references) leading to biased gradient estimates while the latter often relies on data sampling techniques (Needell et al., 2014; Papa et al., 2015) to obtain unbiased gradient estimates. In both cases, a random gradient estimate is available at a cheap computing cost and the method consists in moving along this estimate at each iteration. Early seminal works on such stochastic algorithms include Robbins and Monro (1951); Kiefer et al. (1952) and a recent review dealing with large scale learning problems is given in Bottou et al. (2018).

Starting from an initial point $\theta_0 \in \mathbb{R}^p$, the SGD algorithm is defined by the update rule

$$\forall t \geq 0, \quad \theta_{t+1} = \theta_t - \gamma_{t+1} g_t$$

where $g_t \in \mathbb{R}^p$ is a gradient estimate at θ_t (possibly biased) and $(\gamma_t)_{t \geq 1}$ is some learning rate sequence that should decrease throughout the algorithm. While the computation of g_t may be cheap, it still requires the computation of a vector of size p which may be a critical issue in high-dimensional problems. To address this difficulty, we rely on sampling well-chosen coordinates of the gradient estimate at each iteration.

We consider the framework of stochastic coordinate gradient descent (SCGD) which modifies standard stochastic gradient descent methods by adding a selection step to perform random coordinate descent. The SCGD algorithm is defined by the following iteration

$$\begin{cases} \theta_{t+1}^{(k)} = \theta_t^{(k)} & \text{if } k \neq \zeta_{t+1} \\ \theta_{t+1}^{(k)} = \theta_t^{(k)} - \gamma_{t+1} g_t^{(k)} & \text{if } k = \zeta_{t+1} \end{cases}$$

where ζ_{t+1} is a random variable valued in $\llbracket 1, p \rrbracket$ which selects a coordinate of the gradient estimate. The distribution of ζ_t is called the *coordinate sampling policy*. Note that the SCGD framework is very general as it contains as many methods as there are ways to generate both the gradient estimate g_t and the random variables ζ_t .

Contributions. The main contributions are as follows

(i) (*Theory*) We show the almost-sure convergence of the SCGD iterates $(\theta_t)_{t \in \mathbb{N}}$ towards stationary points in the sense that $\nabla f(\theta_t) \rightarrow 0$ almost surely as well as non-asymptotic bounds on the optimality gap $\mathbb{E}[f(\theta_t) - f^*]$ where f^* is a lower bound of f . The working conditions are relatively weak as the function f is only required to be L -smooth (classical in non-convex problems) and the stochastic gradients are possibly biased with unbounded variance, using a growth condition related to *expected smoothness* (Gower et al., 2019).

(ii)(Practice) We develop a new algorithm, called MUSKETEER, for *Multivariate Stochastic Knowledge Extraction Through Exploration Exploitation Reinforcement*. In the image of the motto 'all for one and one for all', this procedure belongs to the SCGD framework with a particular design for the *coordinate sampling policy*. It compares the value of all past gradient estimates g_t to select a descent direction (*all for one*) and then moves the current iterate according to the chosen direction (*one for all*). The heuristic is the one of reinforcement learning in the sense that large gradient coordinates represent large decrease of the objective and can be seen as high rewards. The resulting directions should be favored compared to the path associated to small gradient coordinates. By updating the *coordinate sampling policy*, the algorithm is able to detect when a direction becomes rewarding and when another one stops being engaging.

Related work. The authors of (Nutini et al., 2015) investigate the deterministic Gauss-Southwell rule which consists of picking the coordinate with maximum gradient value. In trusting large gradients, this rule looks like the one of MUSKETEER except that no stochastic noise -neither in the gradient evaluation nor in the coordinate selection- is present in their algorithm. In that aspect, our method differs from all the previous CD studies (Loshchilov et al., 2011; Richtárik and Takáč, 2016a; Glasmachers and Dogan, 2013; Qu and Richtárik, 2016; Allen-Zhu et al., 2016; Namkoong et al., 2017) which rely on ∇f . Among the SGD literature, compression and sparsification methods (Alistarh et al., 2017; Wangni et al., 2018) were developed for communication efficiency. The former use compression operators to select a few components of the gradient estimates at the cost of full gradient computation and coordinate sorting. The latter use a gradient estimate g which is sparsified using probability weights to reach an unbiased estimate of the gradient. In contrast, the SCGD framework allows the gradient to be biased as no importance re-weighting is performed. Note also that, to cover zeroth-order methods, the gradient estimate itself g_t is allowed to be biased as for instance in the recent study of Ajalloeian and Stich (2020). The proofs of the asymptotic convergence results are based on ideas from Bertsekas and Tsitsiklis (2000) with particular extensions in the framework of biased gradient estimates. Finally, the non-asymptotic bounds are inspired from Moulines and Bach (2011) where the authors provide a non-asymptotic analysis for standard SGD.

Outline. Section 2 introduces the mathematical framework with the different sampling strategies and Section 3 contains our main theoretical results. Section 4 is dedicated to MUSKETEER algorithm and a numerical analysis is performed in Section 5. Proofs, technical details and additional experiments may be found in the appendix.

2. Mathematical Background

2.1 Notation and problem set-up

Notation. Denote by (e_1, \dots, e_p) the canonical basis of \mathbb{R}^p and for $k \in \llbracket 1, p \rrbracket$, $D(k) = e_k e_k^T \in \{0, 1\}^{p \times p}$ is a diagonal matrix with a 1 in position k . $\|\cdot\|_2$ and $\|\cdot\|_\infty$ are respectively the Euclidian and infinity norm. For any $u \in \mathbb{R}^p$, $u^{(k)}$ is the k -th coordinate of u ; $\mathbb{1}_A$ is the indicator function of the event A , i.e., $\mathbb{1}_A = 1$ if A is true and $\mathbb{1}_A = 0$ otherwise. Denote by $\mathcal{U}(\llbracket 1, p \rrbracket)$ the uniform distribution over $\llbracket 1, p \rrbracket$. For a vector of probability weights $d = (d^{(1)}, \dots, d^{(p)})$ with $\sum_{k=1}^p d^{(k)} = 1$, denote by $Q(d)$ the associated categorical distribution.

Problem set-up. Consider the classical stochastic optimization problem

$$\min_{\theta \in \mathbb{R}^p} \{f(\theta) = \mathbb{E}_\xi[f(\theta, \xi)]\},$$

where ξ is a random variable. In many scenarios, *e.g.* empirical risk minimization or reinforcement learning, the gradient ∇f cannot be computed in a reasonable time and only a stochastic version, possibly biased, is available. The distribution of ξ is called the *data sampling policy* as it refers to the sampling mechanism in the empirical risk minimization (ERM) framework. This running example is presented below and shall be considered throughout the paper. Other classical optimization problems where stochastic gradients are available include adaptive importance sampling (Delyon and Portier, 2018), policy gradient methods (Hanna et al., 2019) and optimal transport (Genevay et al., 2016).

Running Example (ERM). Given some observed data $z_1, \dots, z_n \subset \mathcal{Z}$ and a loss function $\ell : \mathbb{R}^p \times \mathcal{Z} \rightarrow \mathbb{R}$, the objective function f approximates the risk $\mathbb{E}_z[\ell(\theta, z)]$ by the so-called empirical risk defined as

$$\forall \theta \in \mathbb{R}^p, \quad f(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(\theta, z_i).$$

Evaluating f or its gradient is prohibitive in large scale machine learning as it requires seeing all the samples in the dataset. Instead, after picking at random an index $j = \xi$, uniformly distributed over $\llbracket 1, n \rrbracket$, the k -th coordinate of the gradient estimate may be computed as $(\ell(\theta + he_k, z_j) - \ell(\theta, z_j))/h$. When differentiation is possible, another gradient estimate is offered by $\nabla_\theta \ell(\theta, z_j)$. These two gradient estimates are of a different nature: the first one, often referred to as zeroth-order estimate, is biased whereas the second one, often referred to as first order estimate, is unbiased.

2.2 Gradient estimates

Throughout the paper, the gradient generator is denoted by $g_h(\cdot, \xi)$ where the parameter $h \geq 0$ represents the underlying bias as claimed in the next assumption. This level of generality allows to include zeroth-order estimate as discussed right after the assumption.

Assumption 1 (Biased gradient) *There exists a constant $c \geq 0$ such that:*

$$\forall h > 0, \forall \theta \in \mathbb{R}^p, \quad \|\mathbb{E}_\xi[g_h(\theta, \xi)] - \nabla f(\theta)\|_2 \leq ch.$$

This general assumption enables to work with classical unbiased gradient in the framework of first order estimates by taking $c = 0$. Furthermore, Assumption 1 is satisfied for the following well-spread zeroth-order estimates.

Example 1 (smoothing). The smoothed gradient estimate (Nesterov and Spokoiny, 2017) is given for all $\theta \in \mathbb{R}^p$ by $g_h(\theta, \xi) = h^{-1}[f(\theta + hU, \xi) - f(\theta, \xi)]U$ where U is a standard Gaussian vector (independent from ξ). An alternative version consists in taking U uniformly distributed over the unit sphere.

Example 2 (finite differences). The finite differences gradient estimate is given for all $\theta \in \mathbb{R}^p$ by $g_h(\theta, \xi) = \sum_{k=1}^p g_h(\theta, \xi)^{(k)} e_k$ where for all $k = 1, \dots, p$ the coordinates are $g_h(\theta, \xi)^{(k)} = h^{-1}[f(\theta + he_k, \xi) - f(\theta, \xi)]$.

Both previous examples share the following general property. There exists a probability measure ν satisfying $\int_{\mathbb{R}^p} xx^\top \nu(dx) = I_p$ such that,

$$\forall h > 0, \theta \in \mathbb{R}^p, \quad \mathbb{E}_\xi[g_h(\theta, \xi)] = \int_{\mathbb{R}^p} x \left\{ \frac{f(\theta + hx) - f(\theta)}{h} \right\} \nu(dx). \quad (1)$$

The smoothed gradient estimate is recovered when ν is the standard Gaussian measure and taking $\nu = \sum_{k=1}^p \delta_{e_k}/p$ covers the finite differences estimate. As detailed in the next subsection, an interesting framework is to use a measure ν that evolves through time and put different weights on the different directions. As stated in the following proposition, when the function f is L -smooth, i.e., ∇f is L -Lipschitz, the bias of the gradient estimate (1) is of order h and thus satisfies Assumption 1.

Proposition 1 *Under Eq. (1), if f is L -smooth, then Assumption 1 holds true with $c = \sqrt{CL}/2$ where $C = \int_{\mathbb{R}^p} \|x\|_2^6 \nu(dx) < \infty$.*

The previous proposition allows us to cover the two methods: smoothing and finite difference. Note that for the latter, the constant C is equal to 1.

2.3 Coordinate Sampling Policy

Let $(\xi_t)_{t \geq 1}$ be a sequence of independent and identically distributed random variables. Let $(\gamma_t)_{t \geq 1}$ be a sequence of positive numbers called *learning rates*. Let $(h_t)_{t \geq 1}$ be a sequence of positive numbers called *smoothing parameters*. Denote by $g_t = g_{h_{t+1}}(\theta_t, \xi_{t+1})$ the gradient estimate at time t . The classical SGD update rule is given by

$$\theta_{t+1} = \theta_t - \gamma_{t+1} g_t, \quad t \geq 0, \quad (2)$$

For any $t \in \mathbb{N}$, $\mathcal{F}_t = \sigma(\theta_0, \theta_1, \dots, \theta_t)$ is the σ -field associated to the sequence of iterates $(\theta_t)_{t \in \mathbb{N}}$.

The framework of SCGD is introduced thanks to random coordinate sampling. At each step, only one coordinate of the parameter of interest is updated. This coordinate is selected at random according to a distribution valued in $\llbracket 1, p \rrbracket$ which is allowed to evolve during the algorithm. The iteration of the coordinate sampling algorithm is given coordinate-wise by

$$\begin{cases} \theta_{t+1}^{(k)} = \theta_t^{(k)} & \text{if } k \neq \zeta_{t+1} \\ \theta_{t+1}^{(k)} = \theta_t^{(k)} - \gamma_{t+1} g_t^{(k)} & \text{if } k = \zeta_{t+1} \end{cases} \quad (3)$$

where ζ_{t+1} is a random variable valued in $\llbracket 1, p \rrbracket$. Hence ζ_{t+1} selects the coordinate along which the t -th descent shall proceed. The distribution of ζ_{t+1} is called the *coordinate sampling policy* as opposed to the *data sampling policy* governed by the random variable ξ_{t+1} . The distribution of ζ_{t+1} is characterized by the probability weights vector $d_t = (d_t^{(1)}, \dots, d_t^{(p)})$ defined by

$$d_t^{(k)} = \mathbb{P}(\zeta_{t+1} = k | \mathcal{F}_t), \quad k \in \llbracket 1, p \rrbracket.$$

The categorical distribution on $\llbracket 1, p \rrbracket$ associated to d_t is denoted by $Q(d_t)$, i.e., conditionally to \mathcal{F}_t , we have:

$$\forall t \geq 0, \quad \zeta_{t+1} \sim Q(d_t) \quad \text{with} \quad d_t = (d_t^{(1)}, \dots, d_t^{(p)}).$$

Running Example (ERM). The CD algorithm defined by Equation (3) can easily be applied in the ERM framework. The *coordinate sampling* strategy $\zeta \sim Q(d_t)$ combined with the uniform *data sampling* $\xi \sim \mathcal{U}(\llbracket 1, n \rrbracket)$ leads to $\theta_{t+1}^{(\zeta)} = \theta_t^{(\zeta)} - (\gamma_{t+1}/h_{t+1})(\ell(\theta_t + h_{t+1}e_\zeta, z_\xi) - \ell(\theta_t, z_\xi))$ (zeroth-order) and $\theta_{t+1}^{(\zeta)} = \theta_t^{(\zeta)} - \gamma_{t+1}\partial_{\theta_\zeta}\ell(\theta_t, z_\xi)$ (first order).

Given the past, the *data sampling* and *coordinate sampling* draws should not be related.

Assumption 2 (Conditional Independence) ζ_{t+1} is independent from ξ_{t+1} conditionally on \mathcal{F}_t .

This assumption is natural in the ERM context as in most cases there is no particular link between the sample indexes and the coordinates. Furthermore, the independence property plays an important role in our proofs. The SCGD algorithm defined in (3) is simply written with matrix notation as

$$\theta_{t+1} = \theta_t - \gamma_{t+1}D(\zeta_{t+1})g_t,$$

where $D(k) = e_k e_k^\top \in \mathbb{R}^{p \times p}$ has its entries equal to 0 except the (k, k) which is 1. Observe that the distribution of the random matrix $D(\zeta_{t+1})$ is fully characterized by the matrix

$$D_t = \mathbb{E}[D(\zeta_{t+1})|\mathcal{F}_t] = \text{Diag}(d_t^{(1)}, \dots, d_t^{(p)}).$$

Note that under Assumptions 1 and 2, the average move of SCGD follows a biased gradient direction. For instance, when $c = 0$, the average move of SCGD is given by $\mathbb{E}[\theta_{t+1} - \theta_t | \mathcal{F}_t] = -\gamma_{t+1}D_t \nabla f(\theta_t)$ which bears resemblance to the Conditioned-SGD iteration (Bottou et al., 2018, Section 6.2). Such preprocessing is meant to refine the gradient direction through a matrix multiplication for a better understanding of the underlying structure of the data. A natural question rises on the choice of the matrix D_t among all the possible coordinate sampling distributions.

The SCGD framework is efficient as soon as one can compute each coordinate of the gradient estimate. This is the case for zeroth-order (ZO) optimization with finite differences where the full gradient estimate uses p partial derivatives, each of them requiring two queries of the objective function. SCGD reduces this cost to a single coordinate update.

Remark 1 (Batch coordinates) A natural extension is to consider subsets of coordinates, a.k.a. *block-coordinate descent*. Note that this framework is covered by our approach as the proofs can be extended by summing different matrices $D(\zeta)$. Similarly to mini-batching (Gower et al., 2019), one can consider multiple draws for the coordinates that are to be updated. The selecting random matrix $D(\zeta_{t+1})$ may be replaced by a diagonal matrix with $m (< p)$ non-zero coefficients. For that matter, it is enough to have multiple draws from the categorical distribution $Q(d_t)$.

Remark 2 (Parallelization) Several families of communication-reduction methods such as quantization (Alistarh et al., 2017), gradient sparsification (Wangni et al., 2018; Alistarh et al., 2018) or local-SGD (Patel and Dieuleveut, 2019) have been proposed to reduce the overheads of distribution. The SCGD framework can benefit from such data parallelization techniques. When a fixed number m of machines is available, it is then possible to gain computational acceleration by drawing m times the coordinate distribution $Q(d_t)$ on the different machines and then transmit the batch of selected coordinates to the workers.

2.4 Adaptive and Unbiased Policies

To understand more clearly the differences between SGD and SCGD, we shall rely on a more general iteration scheme. This framework is useful to compare different algorithms in terms of adaptive policies and unbiased estimates. Consider the following general update rule

$$\theta_{t+1} = \theta_t - \gamma_{t+1}h(\theta_t, \omega_{t+1}), \quad t \geq 0 \tag{4}$$

where h is a gradient generator and $(\omega_t)_{t \geq 1}$ is a sequence of random variables which are not necessarily independent nor identically distributed. Observe that both frameworks, SGD and SCGD, are instances of (4). For example, the randomness of SCGD can be expressed through $\omega_t = (\xi_t, \zeta_t)$.

Definition 1 (Policy) *Denote by P_t the distribution of ω_{t+1} given \mathcal{F}_t . The sequence $(P_t)_{t \geq 0}$ is called the policy of the stochastic algorithm.*

The policy of a stochastic algorithm is an important tool as it determines the randomness introduced over time. On the one hand, it provides insights on the expected behavior of the algorithm. On the other hand, it measures the ability to adapt through the iterations.

Definition 2 (Unbiased and Adaptive) *A policy $(P_t)_{t \geq 0}$ is called "unbiased" if: $\forall \theta \in \mathbb{R}^p, t \geq 0, \int h(\theta, \omega)P_t(d\omega) \propto \nabla f(\theta)$. It is called "naive" if P_t does not change with t , otherwise it is adaptive.*

With these definitions in mind, it is clear that the SGD policy (2) under Assumption 1 with $c = 0$ is unbiased and naive, and so does the policy induced by first order gradient in ERM.

Within the framework of SCGD, a policy cannot be unbiased and adaptive as claimed in the next proposition.

Proposition 2 (Unbiased coordinate policy) *Suppose that Assumption 1 is fulfilled with $c = 0$ and that $\text{Span}\{\nabla f(\theta) : \theta \in \mathbb{R}^p\}$ is dense in \mathbb{R}^p , then the only unbiased coordinate sampling policy is $D_t = I_p/p$. It corresponds to uniform coordinate sampling.*

When working under Assumption 1 with $c = 0$, SCGD with uniform coordinate sampling is unbiased and hence similar to SGD. This is confirmed in the numerical experiments (Appendix E and G). However, a uniform sampling does not use any available information to favor coordinates among others. Thus, the approach promoted in the paper is different: past gradient values are used to update the probability weights of D_t . The resulting method is an adaptive algorithm which is biased.

Remark 3 (Importance Coordinate Sampling) *Note that the general framework defined above includes the particular case where the coordinates are selected according to ζ then reweighted as proposed in (Wangni et al., 2018). This corresponds to the choice $h(\theta, \omega_{t+1}) = D_t^{-1}D(\zeta_{t+1})g(\theta, \xi_{t+1})$. Even though such a policy is adaptive and unbiased, it turns out -from our numerical experiments (Appendix F)- that it behaves similarly to the uniform version and is therefore sub-optimal.*

3. Main Theoretical Results

In a general non-convex setting, we investigate the almost sure convergence of SCGD algorithms as well as non-asymptotic bounds. The following assumptions on the objective function f are classical among the SGD literature.

Assumption 3 (Lower bound) *There exists $f^* \in \mathbb{R}$ such that: $\forall \theta \in \mathbb{R}^p, f(\theta) \geq f^*$.*

Assumption 4 (Smoothness) *The objective $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is twice continuously differentiable and L -smooth: $\forall \theta, \eta \in \mathbb{R}^p, \|\nabla f(\theta) - \nabla f(\eta)\|_2 \leq L\|\theta - \eta\|_2$.*

Remark 4 (Coordinate smoothness) *Note that this assumption may be refined using the notion of coordinate smoothness with parameters (L_1, \dots, L_p) where for all $k = 1, \dots, p$, $\partial_k f(\cdot)$ is L_k -Lipschitz, i.e., for all $\theta \in \mathbb{R}^p, \delta \in \mathbb{R}, |\partial_k f(\theta + \delta e_k) - \partial_k f(\theta)| \leq L_k |\delta|$. Within this framework, small values of L_k are associated to a high degree of smoothness in the k -th direction. Conversely, large values of L_k are associated to more difficult minimization problems along that direction. Intuitively, it requires more energy to minimize f along these directions and one should assign more sampling probability on coordinates with larger L_k (see Proposition 9 in the appendix).*

When dealing with stochastic algorithms, the stochastic noise associated to the gradient estimates is the keystone for the theoretical analysis. To treat this term, we consider a weak growth condition, related to the notion of *expected smoothness* as introduced in Gower et al. (2019) (see also Gazagnadou et al. (2019); Gower et al. (2021)).

Assumption 5 (Growth condition) *With probability 1, there exist $0 \leq \mathcal{L}, \sigma^2 < \infty$ such that for all $\theta \in \mathbb{R}^p$ and $h > 0$, we have: $\mathbb{E} [\|g_h(\theta, \xi)\|_2^2] \leq 2\mathcal{L}(f(\theta) - f^*) + \sigma^2$.*

This bound on the stochastic noise $\mathbb{E} [\|g(\theta, \xi)\|_2^2]$ is the key to prove the almost sure convergence of the algorithm. Note that Assumption 5 is weak as it allows the noise to be large when the iterate is far away from the optimal point. In that aspect, it contrasts with uniform bounds of the form $\mathbb{E} [\|g(\theta, \xi)\|_2^2] \leq \sigma^2$ for some deterministic $\sigma^2 > 0$ (Nemirovski and Yudin, 1983; Nemirovski et al., 2009; Shalev-Shwartz et al., 2011). Observe that such uniform bound is recovered by taking $\mathcal{L} = 0$ in Assumption 5 but cannot hold when the objective function f is strongly convex (Nguyen et al., 2018). The standard Robbins-Monro condition, $\sum_{t \geq 1} \gamma_t = +\infty$ and $\sum_{t \geq 1} \gamma_t^2 < +\infty$ is required in the next theorem which serves as a starting point for a comparison between SGD and SCGD methods.

Theorem 1 (Almost sure convergence of biased SGD) *Suppose that Assumptions 1 to 5 are fulfilled and let $(\theta_t)_{t \in \mathbb{N}}$ be the sequence of iterates defined by (2). If the learning rates satisfy the Robbins-Monro condition and $h_t^2 = O(\gamma_t)$ then $\nabla f(\theta_t) \rightarrow 0$ a.s. when $t \rightarrow +\infty$.*

The SCGD framework is very general in the sense that it covers as many algorithms as there are ways to generate both the gradient estimate g_t and the random variables ζ_t that select the coordinates. The next theorem provides the almost sure convergence of particular instances of SCGD algorithms where the true gradient is known and used to define the *coordinate sampling* policy. It recovers the deterministic Gauss-Southwell rule (Nutini et al., 2015) and extends it to the case where the coordinate weights are proportional to any norm of the current gradient $\nabla f(\theta_t)$.

Theorem 2 (Almost sure convergence of particular SCGD) *Suppose that Assumptions 1 to 5 are fulfilled and let $(\theta_t)_{t \in \mathbb{N}}$ be the sequence of iterates defined by (3), i.e., $\theta_{t+1} = \theta_t - \gamma_{t+1} D(\zeta_{t+1}) g_t$. If the learning rates satisfy the standard Robbins-Monro and $h_t^2 = O(\gamma_t)$, then the two following results hold:*

- (a) (maximum gradient) *if the selected coordinate follows the maximum coordinate of the gradient $\zeta_{t+1} = \arg \max_{k=1, \dots, p} |\partial_k f(\theta_t)|$ then $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$.*
- (b) (gradient weights) *if the selection weights are proportional to the gradient norm $D_t \propto (|\nabla_k f(\theta_t)|^q)_{1 \leq k \leq p}$ with $q > 0$ then $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$.*

Remark 5 (Sparse Gradient) *In light of the sparsity assumption used in Wang et al. (2018) (Assumption A5), note that SCGD methods with weights proportional to the gradient coordinates can outperform uniform coordinate sampling as they only select the relevant directions throughout the procedure. Such sparsity framework happens for instance in hyper-parameter tuning problems of learning systems: usually the performance of the system is insensitive to some hyper-parameters which implies the sparsity of the gradients.*

In the general case, one may not have access to the true gradient and can only rely on the estimate g_t . Another assumption is therefore needed on the weights of the *coordinate sampling* policy to ensure that all the coordinates of interest are selected throughout the algorithm. The success of the proposed approach relies on the following restrictions between the *learning rates* sequence $(\gamma_t)_{t \in \mathbb{N}}$ and the weights of the *coordinate policy*. This is formally stated in the following assumption, referred to as the extended Robbins-Monro condition. Denote by β_{t+1} the smallest probability weight at time t , i.e., $\beta_{t+1} = \min_{1 \leq k \leq p} d_t^{(k)}$.

Assumption 6 (Extended Robbins-Monro condition) *$(\gamma_t)_{t \geq 1}, (\beta_t)_{t \geq 1}$ are positive sequences such that $\sum_{t \geq 1} \gamma_t \beta_t = +\infty$ and $\sum_{t \geq 1} \gamma_t^2 < +\infty$.*

From a practical point of view, those are not restrictive as they can always be implemented by the user. In the case $D_t = I_p$, this is simply the standard Robbins-Monro condition.

Theorem 3 (Almost sure convergence of general SCGD) *Suppose that Assumptions 1 to 5 are fulfilled and let $(\theta_t)_{t \in \mathbb{N}}$ be the sequence of iterates defined by (3). Assume moreover that the learning rates satisfy Assumption 6, $h_t^2 = O(\gamma_t)$ and that (β_t) has a positive lower bound, then $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$.*

Remark 6 (Global convergence) *Other convergence results concerning the sequence of iterates towards global minimizers may be obtained by considering stronger assumptions including that f is coercive and the level sets of stationary points $\{\theta, \nabla f(\theta) = 0\} \cap \{\theta, f(\theta) = y\}$ are locally finite for every $y \in \mathbb{R}^d$ (see Gadat et al. (2018) or Appendix B.1).*

For a non-asymptotic analysis, we place ourselves under the Polyak–Łojasiewicz (PL) condition (Polyak, 1963) which does not assume convexity of f but retains many properties of strong convexity, e.g. the fact that every stationary point is a global minimum.

Assumption 7 (PL inequality) *There exists a constant $\mu > 0$ such that:*

$$\forall \theta \in \mathbb{R}^p, \|\nabla f(\theta)\|_2^2 \geq 2\mu (f(\theta) - f^*).$$

Similarly to (Moulines and Bach, 2011), we introduce $\varphi_\alpha : \mathbb{R}_+^* \rightarrow \mathbb{R}, \varphi_\alpha(t) = \alpha^{-1}(t^\alpha - 1)$ if $\alpha \neq 0$ and $\varphi_\alpha(t) = \log(t)$ if $\alpha = 0$. Denoting $\delta_t = \mathbb{E}[f(\theta_t) - f^*]$ and assuming that $\beta_{t+1} \geq \beta > 0$, one can obtain the recursion equation: $\delta_t \leq (1 - 2\mu\beta\gamma_t + L\mathcal{L}\gamma_t^2) \delta_{t-1} + \gamma_t^2(\sigma^2 L + c^2)/2$, leading to the following theorem on non-asymptotic bounds for SCGD methods.

Theorem 4 (Non-asymptotic bounds) *Suppose that Assumptions 1 to 7 are fulfilled and let $(\theta_t)_{t \in \mathbb{N}}$ defined in (3) with $\gamma_t = \gamma t^{-\alpha}$ and $h_t = \sqrt{\gamma_t}$. Denote by $\delta_t = \mathbb{E}[f(\theta_t) - f^*]$ and assume that there exists $\beta > 0$ such that $\beta_{t+1} \geq \beta > 0$. We have for $\alpha \in [0, 1]$:*

- If $0 \leq \alpha < 1$ then

$$\delta_t \leq 2 \exp(2L\mathcal{L}\gamma^2\varphi_{1-2\alpha}(t)) \exp\left(-\frac{\mu\beta\gamma}{4}t^{1-\alpha}\right) \left(\delta_0 + \frac{\sigma^2 + 2c^2}{2\mathcal{L}}\right) + \frac{\gamma(\sigma^2 L + 2c^2)}{\mu\beta}t^{-\alpha}$$

- If $\alpha = 1$ then

$$\delta_t \leq 2 \exp(L\mathcal{L}\gamma^2) \left(\delta_0 + \frac{\sigma^2 + 2c^2}{2\mathcal{L}}\right) t^{-\mu\beta\gamma} + \left(\frac{\sigma^2 L}{2} + c^2\right) \gamma^2 \varphi_{\mu\beta\gamma/2-1}(t) t^{-\mu\beta\gamma/2}$$

Remark 7 (Importance weights) *The conclusion of Theorem 3 remains valid for the update rule $\theta_{t+1} = \theta_t - \gamma_{t+1} W_t D(\zeta_{t+1}) g_t$ where W_t is a diagonal matrix with coefficients $(w_t^{(1)}, \dots, w_t^{(p)})$ such that $\beta_{t+1} = \min_{1 \leq k \leq p} w_t^{(k)} d_t^{(k)}$.*

Remark 8 (Norms and constants) *A quick inspection of the proof reveals that Assumptions 1 and 5 may be replaced respectively by: $\forall \theta \in \mathbb{R}^p, h > 0, \|\mathbb{E}_\xi[g_h(\theta, \xi)] - \nabla f(\theta)\|_\infty \leq ch$ and $\max_{k=1, \dots, p} \mathbb{E}[g_h^{(k)}(\theta, \xi)^2] \leq 2\mathcal{L}(f(\theta) - f(\theta^*)) + \sigma^2$. Since $\|\cdot\|_\infty \leq \|\cdot\|_2 \leq \sqrt{p}\|\cdot\|_\infty$, the above constant scales more efficiently with the dimension.*

Remark 9 (Rates) *The optimal convergence rate in Theorem 4 is of order $O(1/t)$, obtained with $\alpha = 1$ under the condition $\mu\beta\gamma > 2$. Such rate matches optimal asymptotic minimax rate for stochastic approximation (Agarwal et al., 2012) and recovers the rate of (Ajalloeian and Stich, 2020) for SGD with biased gradients.*

4. MUSKETEER Algorithm

This section is dedicated to the algorithm MUSKETEER which performs an adaptive reweighting of the coordinate sampling probabilities to leverage the data structure. Note that this procedure is general and may be applied on top of any stochastic optimization algorithm as soon as one has access to coordinates of a gradient estimate. In view of Theorem 2 and Remark 5, the main idea is to rely on a stochastic version of the Gauss-southwell rule where the coordinates of the gradients are only available through random estimates. The algorithm of interest alternates between two elementary blocks: one for the *exploration* phase and another one for the *exploitation* phase.

Exploration phase. The goal of this phase is twofold: perform stochastic coordinate gradient descent and collect information about the noisy directions of the gradient. The former task is done using the current coordinate sampling distribution $Q(d_n)$ which is fixed during this phase whereas the latter is computed through cumulative gains.

Exploitation phase. This phase is the cornerstone of the probability updates since it exploits the knowledge of the cumulative gains to update the coordinate sampling probability vector d_n in order to sample more often the relevant directions of the optimization problem.

MUSKETEER

Require: $\theta_0 \in \mathbb{R}^p$, $N, T \in \mathbb{N}$, $(\gamma_t)_{t \geq 0}$, $(\lambda_n)_{n \geq 0}$, $\eta > 0$.

1. Initialize probability weights $d_0 = (1/p, \dots, 1/p)$ // start with uniform sampling
 2. Initialize cumulative gains $G_0 = (0, \dots, 0)$
 3. **for** $n = 0, \dots, N - 1$ **do**
 4. Initialize current gain $\tilde{G}_0 = (0, \dots, 0)$
 5. Run **Explore** (T, d_n) // to compute current gain \tilde{G}_T
 6. Run **Exploit** $(G_n, \tilde{G}_T, \lambda_n, \eta)$ // to update weights d_{n+1}
 7. **end for**
 8. Return final point θ_N
-

Consider a fixed iteration $n \in \mathbb{N}$ of MUSKETEER's main loop. The *exploration* phase may be seen as a multi-armed bandit problem (Auer et al., 2002a) where the arms are the gradient coordinates for $k \in \llbracket 1, p \rrbracket$. At each time step $t \in \llbracket 1, T \rrbracket$, a coordinate ζ is drawn according to $Q(d_n)$ and the relative gradient $g_t^{(\zeta)}/d_n^{(\zeta)}$, representing the reward, is observed. Note that an importance sampling strategy is used to produce an unbiased estimate of the gradient when dealing with first order methods. The rewards are then used to build cumulative gains \tilde{G}_T which can be written in a vectorized form as an empirical sum of the visited gradients during the *exploration* phase

$$\forall k \in \llbracket 1, p \rrbracket, \quad \tilde{G}_T^{(k)} = \frac{1}{T} \sum_{t=1}^T \frac{g_t^{(k)}}{d_n^{(k)}} \mathbb{1}_{\{\zeta_{t+1}=k\}}, \quad i.e. \quad \tilde{G}_T = \frac{1}{T} \sum_{t=1}^T D_n^{-1} D(\zeta_{t+1}) g(\theta_t, \xi_{t+1}). \quad (5)$$

This average reduces the noise induced by the gradient estimates but may be sign-dependent. Thus, one may rely on the following cumulative gains which are also considered in the experiments,

$$\tilde{G}_T = \frac{1}{T} \sum_{t=1}^T D_n^{-1} D(\zeta_{t+1}) |g(\theta_t, \xi_{t+1})| \quad \text{or} \quad \tilde{G}_T = \frac{1}{T} \sum_{t=1}^T D_n^{-1} D(\zeta_{t+1}) g(\theta_t, \xi_{t+1})^2. \quad (6)$$

Starting from $G_0 = (0, \dots, 0)$, the total gain G_n is updated in an online manner during the *exploitation* phase using the update rule $G_{n+1} = G_n + (\tilde{G}_T - G_n)/(n+1)$. Once the average cumulative gains are computed, one needs to normalize them to obtain probability weights. Such normalization can be done by a natural ℓ_1 -reweighting or a softmax operator with a parameter $\eta > 0$. To cover both cases, consider the normalizing function $\varphi : \mathbb{R}^p \rightarrow \mathbb{R}^p$ defined by $\varphi(x)^{(k)} = |x^{(k)}| / \sum_{j=1}^p |x^{(j)}|$ or $\varphi(x)^{(k)} = \exp(\eta x^{(k)}) / \sum_{j=1}^p \exp(\eta x^{(j)})$. Following the sequential approach of the EXP3 algorithm (Auer et al., 2002a,b), the probability weights

are updated through a mixture between the normalized average cumulative gains $\varphi(G_n)$ and a uniform distribution. The former term takes into account the knowledge of the gains by exploiting the rewards while the latter ensures exploration. Given a sequence $(\lambda_n) \in [0, 1]^{\mathbb{N}}$, we have for all $k \in \llbracket 1, p \rrbracket$,

$$d_{n+1}^{(k)} = (1 - \lambda_n)\varphi(G_n)^{(k)} + \lambda_n \frac{1}{p}. \quad (7)$$

Explore (T, d_n)	Exploit ($G_n, \tilde{G}_T, \lambda_n, \eta$)
<ol style="list-style-type: none"> 1. for $t = 1, \dots, T$ do 2. Sample coordinate $\zeta \sim Q(d_n)$ and data ξ 3. Move iterate: $\theta_{t+1}^{(\zeta)} = \theta_t^{(\zeta)} - \gamma_{t+1} g_h^{(\zeta)}(\theta_t, \xi)$ 4. Update gain $\tilde{G}_{t+1}^{(\zeta)}$ using (5) or (6) 5. end for 6. Return vector of gains \tilde{G}_T 	<ol style="list-style-type: none"> 1. Update total average gain G_n in an online manner 2. Compute normalized gains $\varphi(G_n)$ with ℓ_1-weights or softmax 3. Update probability weights d_{n+1} with the mixture of Eq.(7)

In view of Theorem 3, the convergence of the sequence of iterates $(\theta_t)_{t \in \mathbb{N}}$ obtained by MUSKETEER relies on the extended Robbins-Monro condition $\sum_{t \geq 1} \beta_t \gamma_t = +\infty$ which is implied by the weaker condition $\sum_{t \geq 1} \lambda_t \gamma_t = +\infty$ for both ℓ_1 and softmax weights. Observe that such a constraint is easily verified with either a fixed value $\lambda_t \equiv \lambda$ in the mixture update or more generally a slowly decreasing sequence, e.g. $\lambda_t = 1/\log(t)$. Since the gradients $\nabla f(\theta_t)$ get smaller through the iterations, the softmax weights get closer to $1/p$. Thus, in the asymptotic regime, there is no favorable directions among all the possible gradient directions. Hence, near the optimum, the *coordinate sampling policy* of MUSKETEER with softmax weights is likely to treat all the coordinates equally.

Theorem 5 (*Weak convergence*) *Suppose that Assumptions 1 to 5 are fulfilled and that the learning rates satisfy the standard Robbins-Monro condition. Then MUSKETEER's coordinate policy $(Q(d_n))_{n \in \mathbb{N}}$ with softmax normalization converges weakly to the uniform distribution, i.e., $Q(d_n) \rightsquigarrow \mathcal{U}(\llbracket 1, p \rrbracket)$ as $n \rightarrow +\infty$.*

Remark 10 (*On the choice of λ_n and η*) *The uniform term in Equation (7) ensures that all coordinates are eventually visited. Taking $\lambda_n \rightarrow 0$ at a specific rate (which can be derived from the proof) gives more importance to the cumulative gains. The parameter η is fixed during the algorithm and may be tuned through an analysis of the regret (Auer et al., 2002a).*

Remark 11 (*Choice of Exploration Size T*) *Choosing the value of T is a central question known as the exploration-exploitation dilemma in reinforcement learning. As T gets large, the exploration phase gathers more information leading to fewer but more accurate updates. Conversely, with a small value of T , the probabilities get updated more often, at the price of less collected information. Setting $T = p$ ensures that, in average, all the coordinates are visited once during the exploration phase. Nevertheless, a smaller value $T = \lfloor \sqrt{p} \rfloor$ is taken in the experiments and lead to great performance.*

Remark 12 (*Asymptotic behavior*) *The previous results highlight two main features of MUSKETEER: the sequence of iterates converges almost surely and the coordinate policy converges weakly. The latter point suggests that, in the long run, MUSKETEER is similar*

to the uniform coordinate version of SCGD. However, the weak convergence of the rescaled process $(\theta_t - \theta^*)/\sqrt{\gamma_t}$ remains an open question. In light of the link between SCGD and Conditioned-SGD, discussed in Section 2.3, we conjecture that the behavior of MUSKETEER with softmax weights is asymptotically equivalent to SCGD with uniform policy. This is in line with the continuity property obtained in Leluc and Portier (2020) within the Conditioned-SGD framework and relates to the convergence of stochastic Newton algorithms (Boyer and Godichon-Baggioni, 2020).

5. Numerical Experiments

In this section, we empirically validate the SCGD framework by running MUSKETEER and competitors on synthetic and real datasets. First, we focus on regularized regression problems adopting the data generation process of (Namkoong et al., 2017) in which the covariates exhibit a certain block structure. Second, MUSKETEER is employed to train different neural networks models on real datasets for multi-label classification task. For ease of reproducibility, the code is available online¹. Technical details and additional results (with different data settings, normalization and hyperparameters) are available in the appendix.

Methods in competition. The set of methods is restricted to zeroth-order methods. This choice leads to an honest comparison based on the number of function queries. MUSKETEER is implemented according to Section 4 with $T = \lfloor \sqrt{p} \rfloor$, softmax and ℓ_1 normalization for the simulated and real data respectively. The different cumulative gains of Eq. (6) are considered, namely AVG, SQR and ABS for the gradients, their squares or their absolute value respectively. The method FULL is the finite difference gradient estimate computed over all coordinates and UNIFORM stands for the uniform coordinate sampling policy. NESTEROV implements the gaussian smoothing of (Nesterov and Spokoiny, 2017). In all cases, the initial parameter is set to $\theta_0 = (0, \dots, 0)^\top \in \mathbb{R}^p$ and the optimal SGD learning rate of the form $\gamma_k = \gamma/(k + k_0)$ is used.

Regularized linear models. We apply the Empirical Risk Minimization paradigm to regularized linear problems. Given a data matrix $X = (x_{i,j}) \in \mathbb{R}^{n \times p}$, labels $y \in \mathbb{R}^n$ or $\{-1, +1\}^n$ and a regularization parameter $\mu > 0$, the *Ridge regression* objective is defined by

$$f(\theta) = \frac{1}{2n} \sum_{i=1}^n (y_i - \sum_{j=1}^p x_{i,j} \theta_j)^2 + \frac{\mu}{2} \|\theta\|_2^2$$

and the ℓ_2 -regularized logistic regression is given by

$$f(\theta) = \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \sum_{j=1}^p x_{i,j} \theta_j)) + \mu \|\theta\|_2^2.$$

Similarly to (Namkoong et al., 2017), we endow the data matrix X with a block structure. The columns are drawn as $X[:, k] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, p \rrbracket$. The parameters are set to $n = 10,000$ samples in dimension $p = 250$ with an exploration size equal to $T = \lfloor \sqrt{p} \rfloor = 15$. The regularization parameter is set to the classical value $\mu = 1/n$.

1. <https://github.com/RemiLELUC/SCGD-Musketeer>

Figure 1 provides the graphs of the optimality gap $t \mapsto f(\theta_t) - f(\theta^*)$ averaged over 20 independent simulations for different values of $\alpha \in \{2; 5; 10\}$. First, note that the uniform sampling strategy shows similar performance to the classical full gradient estimate. Besides, MUSKETEER with average or absolute gains shows the best performance in all configurations. Greater values of α , *i.e.* stronger block structure, improve our relative performance with respect to the other methods as shown by Figures 1(b) and 1(d).

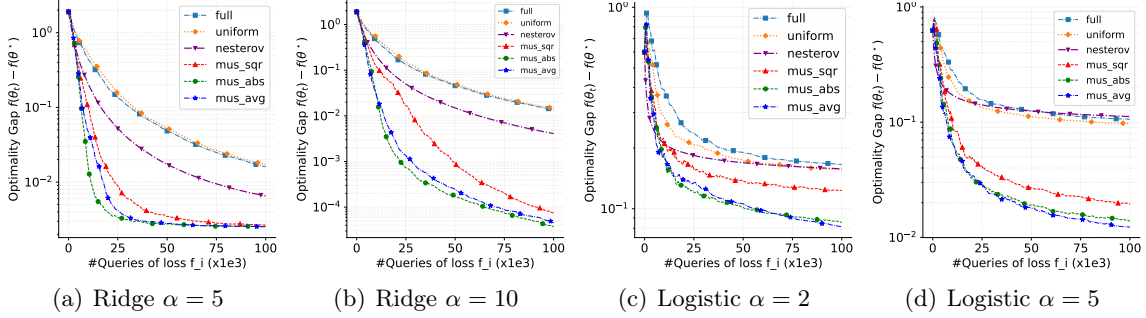


Figure 1: $[f(\theta_t) - f(\theta^*)]$ for Ridge and Logistic on Synthetic data with different block structures.

Neural Networks. We focus on the training of neural networks within the framework of multi-label classification. The datasets in the experiments are popular publicly available deep learning datasets: MNIST (Deng, 2012) and Fashion-MNIST (Xiao et al., 2017). Given an image, the goal is to predict its label among ten classes. The neural architecture is based on linear layers in dimension $p = 55,050$ with $T = 234$. Figure 2 shows the means and standard deviations of the training losses of the different ZO methods averaged over 10 independent runs. Interestingly, the performance of MUSKETEER also benefit from the adaptive structure in terms on accuracy of the test set (see Figures 3(a) and 3(b)). This allows to quantify the statistical gain brought by MUSKETEER over standard ZO methods.

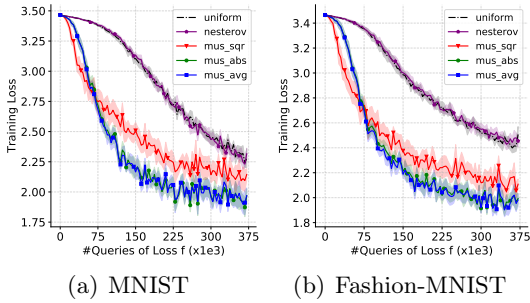


Figure 2: Evolution of training loss.

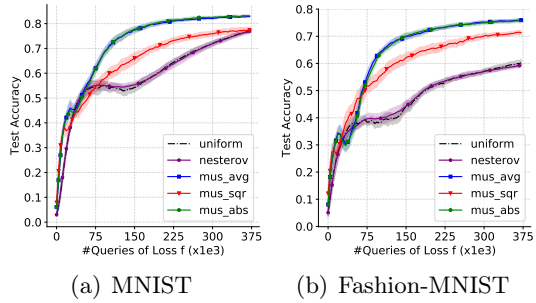


Figure 3: Evolution of test accuracy.

Acknowledgments and Disclosure of Funding

The authors are grateful to the Associate Editor and two anonymous Reviewers for their many valuable comments and interesting suggestions.

Appendices

Table of Contents

A	Technical Proofs	16
A.1	Proof of Proposition 1	16
A.2	Deterministic results for convergence of gradients	16
A.3	Proof of Theorem 1	19
A.4	Proof of Theorem 2	20
A.5	Proof of Theorem 3	22
A.6	Proof of Theorem 4	23
A.7	Proof of Theorem 5	23
B	Additional Results	25
B.1	Almost sure convergence under stronger assumptions	25
B.2	Almost sure convergence of MUSKETEER	25
B.3	Regret analysis in the convex case	25
B.4	Auxiliary Results	27
C	Illustrative Example (stochastic first order)	28
D	Numerical Experiments Details	28
D.1	Regularized linear models	28
D.2	Neural Networks	29
D.3	Hyperparameters and Hardware.	30
E	Numerical Experiments with stochastic first order methods	32
F	Further Numerical Experiments with zeroth-order methods	34
F.1	Ridge Regression (ℓ_1 -reweighting) with different settings of (n, p)	34
F.2	Ridge Regression (softmax reweighting) with different settings of (n, p)	35
F.3	Logistic Regression (ℓ_1 -reweighting) with different settings of (n, p)	36
F.4	Logistic Regression (softmax reweighting) with different settings of (n, p)	37
F.5	Effect of Importance Sampling (IS) on Ridge Regression	38
F.6	Effect of Importance Sampling (IS) on Logistic Regression	39
G	Further Experiments with stochastic first order methods	40
G.1	Comparing learning rates	40
G.2	Ridge Regression with different settings of (n, p)	41
G.3	Logistic Regression with different settings of (n, p)	42

A. Technical Proofs

A.1 Proof of Proposition 1

Under Eq.(1), using Jensen inequality, we find

$$\begin{aligned} \|\mathbb{E}_\xi[g_\mu(\theta, \xi)] - \nabla f(\theta)\|_2^2 &= \left\| \int_{\mathbb{R}^p} x \left(\frac{f(\theta + \mu x) - f(\theta)}{\mu} - x^\top \nabla f(\theta) \right) \nu(\mathrm{d}x) \right\|_2^2 \\ &\leq \int_{\mathbb{R}^p} \|x\|_2^2 \left(\frac{f(\theta + \mu x) - f(\theta)}{\mu} - x^\top \nabla f(\theta) \right)^2 \nu(\mathrm{d}x) \\ &= \mu^{-2} \int_{\mathbb{R}^p} \|x\|_2^2 \left(f(\theta + \mu x) - f(\theta) - \mu x^\top \nabla f(\theta) \right)^2 \nu(\mathrm{d}x) \end{aligned}$$

Using the quadratic bound of L -smooth functions, we obtain

$$\|\mathbb{E}_\xi[g_\mu(\theta, \xi)] - \nabla f(\theta)\|_2^2 \leq \mu^{-2} \frac{L^2}{4} \int \|x\|_2^2 \|\mu x\|_2^4 \nu(\mathrm{d}x) = \mu^2 \frac{L^2}{4} \int \|x\|_2^6 \nu(\mathrm{d}x).$$

A.2 Deterministic results for convergence of gradients

In this section we provide results ensuring the convergence to 0 of several gradient descent algorithms. They are meant to be *high-level* as they may be applied in different situations and *deterministic* because no randomness is measured but only an inclusion of events is considered. The results are key in the proofs.

Lemma 6 (Deterministic result 1) *Let $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be a L -smooth function, $(\gamma_t)_{t \geq 1}$ a positive sequence of learning rates such that $\sum_t \gamma_t = \infty$. Let (θ_t) a random sequence obtained by the SGD update rule $\theta_{t+1} = \theta_t - \gamma_{t+1} g_t$. Let $\omega \in \Omega$ such that the following limits exist:*

$$(i) \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty \quad (ii) \sum_{t \geq 1} \gamma_t (g_{t-1}(\omega) - \nabla f(\theta_{t-1}(\omega))) < \infty$$

then $\nabla f(\theta_t(\omega)) \rightarrow 0$ as $t \rightarrow \infty$.

The next Lemma is the equivalent of Lemma 6 for a specific procedure which, at each iteration, moves only one well-chosen coordinate: the one with highest gradient value.

Lemma 7 (Deterministic result 2) *Let $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be a L -smooth function (with respect to $|\cdot|_\infty$), $(\gamma_t)_{t \geq 1}$ a positive sequence of learning rates such that $\sum_t \gamma_t = \infty$. Let (θ_t) a random sequence obtained by the SCGD update rule $\theta_{t+1} = \theta_t - \gamma_{t+1} D(\zeta_{t+1}) g_t$ with $\zeta_{t+1} = \arg \max_{k=1, \dots, p} |\partial_k f(\theta_t)|$. Let $\omega \in \Omega$ such that the following limits exist:*

$$(i) \sum_{t \geq 0} \gamma_{t+1} |\nabla f(\theta_t(\omega))|_\infty^2 < \infty \quad (ii) \sum_{t \geq 1} \gamma_t D(\zeta_t)(g_{t-1}(\omega) - \nabla f(\theta_{t-1}(\omega))) < \infty,$$

then $\nabla f(\theta_t(\omega)) \rightarrow 0$ as $t \rightarrow \infty$.

We conclude with one last result which is valid for procedure where only one coordinate (chosen randomly) is moved at each iteration.

Lemma 8 (Deterministic result 3) *Let $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be a L -smooth function, $(\gamma_t)_{t \geq 1}$ a positive sequence of learning rates such that $\sum_t \gamma_t = \infty$. Let (θ_t) a random sequence obtained by the SCGD update rule $\theta_{t+1} = \theta_t - \gamma_{t+1} D(\zeta_{t+1}) g_t$ where $\zeta_{t+1} \sim Q(d_t)$. Let $\omega \in \Omega$ such that the following limits exist:*

$$(i) \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty \quad (ii) \sum_{t \geq 1} \gamma_t (D(\zeta_t(\omega)) g_{t-1}(\omega) - D_{t-1} \nabla f(\theta_{t-1}(\omega))) < \infty$$

then $\nabla f(\theta_t(\omega)) \rightarrow 0$ as $t \rightarrow \infty$.

Proof of Lemma 6. The proof (and in particular the reasoning by contradiction) is inspired from the proof of Proposition 1 in Bertsekas and Tsitsiklis (2000). For ease of notation we omit the ω in the proof. Note that condition (i) along with $\sum_t \gamma_t = \infty$ imply that $\liminf_t \|\nabla f(\theta_t)\| = 0$. Now, by contradiction, let $\varepsilon > 0$ and assume that

$$\limsup_t \|\nabla f(\theta_t)\| > \varepsilon$$

We have that there is infinitely many t such that $\|\nabla f(\theta_t)\| < \varepsilon/2$ and also infinitely many t such that $\|\nabla f(\theta_t)\| > \varepsilon$. It follows that there is infinitely many *crossings* between the sets $\{t \in \mathbb{N} : \|\nabla f(\theta_t)\| < \varepsilon/2\}$ and $\{t \in \mathbb{N} : \|\nabla f(\theta_t)\| > \varepsilon\}$. A *crossing* is a collection of indexes $I_k = \{L_k, L_k + 1, \dots, U_k - 1\}$ with $L_k \leq U_k$ ($I_k = \emptyset$ when $L_k = U_k$) such that for all $t \in I_k$,

$$\|\nabla f(\theta_{L_k-1})\| < \varepsilon/2 \leq \|\nabla f(\theta_t)\| \leq \varepsilon < \|\nabla f(\theta_{U_k})\|.$$

Define the following partial Cauchy sequence $R_k = \sum_{t=L_k}^{U_k} \gamma_t (g_{t-1} - \nabla f(\theta_{t-1}))$ and note that condition (ii) implies that $R_k \rightarrow 0$ as $k \rightarrow \infty$. For all $k \geq 1$,

$$\begin{aligned} \varepsilon/2 &\leq \|\nabla f(\theta_{U_k})\|_2 - \|\nabla f(\theta_{L_k-1})\|_2 \\ &\leq \|\nabla f(\theta_{U_k}) - \nabla f(\theta_{L_k-1})\|_2 \\ &\leq L \|\theta_{U_k} - \theta_{L_k-1}\|_2, \end{aligned}$$

where we use that ∇f is L -Lipschitz. Then using the update rule $\theta_t - \theta_{t-1} = -\gamma_t g_{t-1}$, we have by sum

$$\begin{aligned} \varepsilon/2 &\leq L \left\| \sum_{t=L_k}^{U_k} \theta_t - \theta_{t-1} \right\|_2 = L \left\| \sum_{t=L_k}^{U_k} \gamma_t g_{t-1} \right\|_2 \\ &\leq L \left\| \sum_{t=L_k}^{U_k} \gamma_t \nabla f(\theta_{t-1}) \right\|_2 + L \left\| \sum_{t=L_k}^{U_k} \gamma_t (g_{t-1} - \nabla f(\theta_{t-1})) \right\|_2 \\ &\leq L \sum_{t=L_k}^{U_k} \gamma_t \|\nabla f(\theta_{t-1})\|_2 + L \|R_k\|_2 \end{aligned}$$

Since in the previous equation $\|\nabla f(\theta_{t-1})\|_2 > \varepsilon/2$, we get

$$(\varepsilon/2)^2 \leq L \sum_{t=L_k}^{U_k} \gamma_t \|\nabla f(\theta_{t-1})\|_2^2 + (\varepsilon/2)L \|R_k\|_2$$

But since $\sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t)\|^2$ is finite and $\lim_k R_k = 0$, the previous upper bound goes to 0 and implies a contradiction. ■

Proof of Lemma 7. For ease of readability, the variable ω is removed during the proof. By assumption, $|\nabla_{\zeta_{t+1}} f(\theta_t)| = |\nabla f(\theta_t)|_\infty$. Hence, (i) yields that $\liminf_t |\nabla f(\theta_t)|_\infty = 0$. The proof is by contradiction. Suppose that $\limsup_t |\nabla f(\theta_t)|_\infty > \epsilon$. There exists a sequence of *crossings* between the sets $\{t \in \mathbb{N} : |\nabla f(\theta_t)|_\infty < \epsilon/2\}$ and $\{t \in \mathbb{N} : |\nabla f(\theta_t)|_\infty > \epsilon\}$. Formally, there is a collection of indexes $I_k = \{L_k, L_k + 1, \dots, U_k - 1\}$ with $L_k \leq U_k$ ($I_k = \emptyset$ when $L_k = U_k$) such that for all $t \in I_k$,

$$|\nabla f(\theta_{L_k-1})|_\infty < \epsilon/2 \leq |\nabla f(\theta_t)|_\infty \leq \epsilon < |\nabla f(\theta_{U_k})|_\infty.$$

Define

$$R_k = \sum_{t=L_k}^{U_k} \gamma_t D_{\zeta_t} (g_{t-1} - \nabla f(\theta_{t-1}))$$

and use that ∇f is L -smooth to get

$$\begin{aligned} (\epsilon/2) &\leq |\nabla f(\theta_{U_k})|_\infty - |\nabla f(\theta_{L_k-1})|_\infty \\ &\leq L |\theta_{U_k} - \theta_{L_k-1}|_\infty \\ &\leq L \left| \sum_{t=L_k}^{U_k} \gamma_t D_{\zeta_t} \nabla f(\theta_{t-1}) \right|_\infty + L \left| \sum_{t=L_k}^{U_k} \gamma_t D_{\zeta_t} (g_{t-1} - \nabla f(\theta_{t-1})) \right|_\infty \\ &= L \left| \sum_{t=L_k}^{U_k} \gamma_t D_{\zeta_t} \nabla f(\theta_{t-1}) \right|_\infty + L |R_k|_\infty \\ &\leq L \sum_{t=L_k}^{U_k} \gamma_t |D_{\zeta_t} \nabla f(\theta_{t-1})|_\infty + L |R_k|_\infty \end{aligned}$$

Noting that $|D_{\zeta_t} \nabla f(\theta_{t-1})|_\infty = |\nabla f(\theta_{t-1})|_\infty > \epsilon/2$, we get

$$(\epsilon/2)^2 \leq L \sum_{t=L_k}^{U_k} \gamma_t |\nabla f(\theta_{t-1})|_\infty^2 + (\epsilon/2)L |R_k|_\infty.$$

As the previous upper bound converges to 0 by assumption we reach a contradiction. ■

Proof of Lemma 8. Following the proof of Lemma 6, we assume that $\limsup_t \|\nabla f(\theta_t)\|_2 > \epsilon$ and consider the same collection of crossing indexes (L_k, U_k) to obtain that

$$\epsilon/2 \leq L \sum_{t=L_k}^{U_k} \gamma_t \|D_{t-1} \nabla f(\theta_{t-1})\|_2 + L \|R_k\|_2$$

where $R_k = \sum_{t=L_k}^{U_k} \gamma_t (D(\zeta_t)g_{t-1} - D_{t-1} \nabla f(\theta_{t-1}))$ is a sequence that goes to 0. Since in the previous equation $D_{t-1} \preceq I_d$ and $\|\nabla f(\theta_{t-1})\|_2 > \epsilon/2$, we get

$$(\epsilon/2)^2 \leq L \sum_{t=L_k}^{U_k} \gamma_t \|\nabla f(\theta_{t-1})\|_2^2 + (\epsilon/2)L \|R_k\|_2$$

and a contradiction follows as the above term goes to 0. ■

A.3 Proof of Theorem 1

The proof follows from applying Lemma 6 in which two conditions are required:

$$(i) \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty \quad (ii) \sum_{t \geq 1} \gamma_t (g_{t-1}(\omega) - \nabla f(\theta_{t-1}(\omega))) < \infty.$$

Proof of condition (i). We classically rely on the Robbins-Siegmund Theorem (Theorem 10 in Section B.4). Since $\theta \mapsto f(\theta)$ is L -smooth, we have the quadratic bound $f(\eta) \leq f(\theta) + \langle \nabla f(\theta), \eta - \theta \rangle + \frac{L}{2} \|\eta - \theta\|_2^2$. Using the update rule $\theta_{t+1} = \theta_t - \gamma_{t+1} g_t$, we get

$$\begin{aligned} f(\theta_{t+1}) &\leq f(\theta_t) + \langle \nabla f(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L}{2} \|\theta_{t+1} - \theta_t\|_2^2 \\ &= f(\theta_t) - \gamma_{t+1} \langle \nabla f(\theta_t), g_t \rangle + \frac{L}{2} \gamma_{t+1}^2 \|g_t\|_2^2. \end{aligned}$$

Using that

$$2\langle a, b \rangle = \|a\|_2^2 + \|b\|_2^2 - \|a - b\|_2^2 \geq \|a\|_2^2 - \|a - b\|_2^2$$

and taking the conditional expectation, we get

$$\begin{aligned} \mathbb{E}_t[f(\theta_{t+1})] &\leq f(\theta_t) - \gamma_{t+1} \langle \nabla f(\theta_t), \mathbb{E}_t[g_t] \rangle + \frac{L}{2} \gamma_{t+1}^2 \mathbb{E}_t[\|g_t\|_2^2] \\ &\leq f(\theta_t) - \frac{\gamma_{t+1}}{2} \|\nabla f(\theta_t)\|_2^2 + \frac{\gamma_{t+1}}{2} \|\nabla f(\theta_t) - \mathbb{E}_t[g_t]\|_2^2 + \frac{L}{2} \gamma_{t+1}^2 \mathbb{E}_t[\|g_t\|_2^2] \end{aligned}$$

On the one hand, using Assumption 1, we obtain

$$\|\nabla f(\theta_t) - \mathbb{E}_t[g_t]\|_2^2 \leq h_{t+1}^2 c^2$$

On the other hand, using Assumption 5, there exist $0 \leq \mathcal{L}, \sigma^2 < \infty$ such that almost surely

$$\forall t \in \mathbb{N}, \quad \mathbb{E}_t[\|g_t\|_2^2] = \mathbb{E}_\xi[\|g(\theta_t, \xi)\|_2^2] \leq 2\mathcal{L}(f(\theta_t) - f^*) + \sigma^2.$$

Injecting $-f^*$ on both sides, it follows that

$$\mathbb{E}_t[f(\theta_{t+1}) - f^*] \leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \frac{\gamma_{t+1}}{2} \|\nabla f(\theta_t)\|_2^2 + \gamma_{t+1} h_{t+1}^2 c^2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2$$

Introduce $V_t = f(\theta_t) - f^*$, $W_t = \gamma_{t+1} \|\nabla f(\theta_t)\|_2^2 / 2$, $a_t = L\mathcal{L}\gamma_{t+1}^2$ and $b_t = c^2 h_{t+1}^2 \gamma_{t+1} + (L/2)\gamma_{t+1}^2 \sigma^2$. These four random sequences are non-negative \mathcal{F}_t -measurable sequences with $\sum_t a_t < \infty$ and $\sum_t b_t < \infty$ almost surely. We have: $\forall t \in \mathbb{N}, \mathbb{E}[V_{t+1} | \mathcal{F}_t] \leq (1 + a_t)V_t - W_t + b_t$. We can apply Robbins-Siegmund Theorem to have

$$(a) \sum_{t \geq 0} W_t < \infty \text{ a.s.} \quad (b) V_t \xrightarrow{\text{a.s.}} V_\infty, \mathbb{E}[V_\infty] < \infty. \quad (c) \sup_{t \geq 0} \mathbb{E}[V_t] < \infty.$$

Therefore we have a.s. that $(f(\theta_t))$ converges to a finite value $f_\infty \in L^1$ and $\sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t)\|_2^2 < +\infty$. There exists an event $\Omega_0 \subset \Omega$ such that, $\mathbb{P}(\Omega_0) = 1$ and for every $\omega \in \Omega_0$, $\lim_t f(\theta_t(\omega)) < \infty$ and $\sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty$.

Proof of condition (ii). We place ourselves on the event Ω_0 and omit the ω in notation for ease of clarity. First, since $\limsup_t f(\theta_t) < \infty$, we have that $(f(\theta_t))$ is bounded almost surely. It yields, in virtue of Assumption 5 that $\mathbb{E}_t[\|g_t\|_2^2] \leq 2\mathcal{L}(f(\theta_t) - f^*) + \sigma^2 \leq C$ where C is a some finite random variable and the latter holds almost surely. It then follows that, almost surely $\sum_{t \geq 1} \gamma_t^2 \mathbb{E}_t[\|g_t\|^2] \leq C \sum_{t \geq 1} \gamma_t^2 < \infty$. Now, observe that condition (ii) is satisfied as soon as

$$(a) \left\| \sum_{t \geq 0} \gamma_{t+1} (g_t - \mathbb{E}_t[g_t]) \right\|_2 < \infty \quad \text{and} \quad (b) \left\| \sum_{t \geq 0} \gamma_{t+1} (\mathbb{E}_t[g_t] - \nabla f(\theta_t)) \right\|_2 < \infty.$$

Equation (a) involves martingale increments whose quadratic variation satisfies

$$\sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|g_t - \mathbb{E}_t[g_t]\|^2] \leq \sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|g_t\|^2] < \infty,$$

which ensures that $\sum_{t \geq 0} \gamma_{t+1} (g_t - \mathbb{E}_t[g_t]) < \infty$ a.s. in virtue of Theorem 2.17 in Hall and Heyde (1980). The term in equation (b) is bounded using assumption 1 and we have

$$\sum_{t \geq 0} \gamma_{t+1}^2 \|\mathbb{E}_t[g_t] - \nabla f(\theta_t)\|_2^2 \leq c^2 \sum_{t \geq 0} \gamma_{t+1}^2 h_t^2 < \infty,$$

which finally proves

$$(ii) \sum_{t \geq 0} \gamma_{t+1} (g_t(\omega) - \nabla f(\theta_t(\omega))) < \infty$$

and gives, in virtue of Lemma 6 the conclusion $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$.

A.4 Proof of Theorem 2

Part (a) Maximum gradient. The proof follows from applying Lemma 7 in which two conditions are required:

$$(i) \sum_{t \geq 0} \gamma_{t+1} |\nabla f(\theta_t(\omega))|_\infty^2 < \infty \quad (ii) \sum_{t \geq 0} \gamma_{t+1} D(\zeta_{t+1})(g_t(\omega) - \nabla f(\theta_t(\omega))) < \infty.$$

Proof of condition (i). Again, we rely on the quadratic bound

$$\begin{aligned} f(\theta_{t+1}) &\leq f(\theta_t) - \gamma_{t+1} \langle \nabla f(\theta_t), D(\zeta_{t+1})g_t \rangle + \frac{L}{2} \gamma_{t+1}^2 \|D(\zeta_{t+1})g_t\|_2^2 \\ &= f(\theta_t) - \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t) g_t^{(\zeta_{t+1})} + \frac{L}{2} \gamma_{t+1}^2 g_t^{(\zeta_{t+1})2} \end{aligned}$$

Taking the expectation with respect to ξ_{t+1} and using Assumption 2, we find

$$\mathbb{E}_{\xi_{t+1}}[f(\theta_{t+1}) - f^*] \leq f(\theta_t) - f^* - \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t) \tilde{g}_t^{(\zeta_{t+1})} + \frac{L}{2} \gamma_{t+1}^2 \mathbb{E}_{\xi_{t+1}}[g_t^{(\zeta_{t+1})2}]$$

where $\tilde{g}_t = E_\xi[g_{h_{t+1}}(\theta_t, \xi)]$. We use the inequality $2ab \geq a^2 - (a-b)^2$ and Assumption 1 to get

$$\begin{aligned} 2\nabla_{\zeta_{t+1}} f(\theta_t) \tilde{g}_t^{(\zeta_{t+1})} &\geq \nabla_{\zeta_{t+1}} f(\theta_t)^2 - (\nabla_{\zeta_{t+1}} f(\theta_t) - \tilde{g}_t^{(\zeta_{t+1})})^2 \\ &\geq \nabla_{\zeta_{t+1}} f(\theta_t)^2 - \max_{k=1, \dots, p} (\partial_k f(\theta_t) - \tilde{g}_t^{(k)})^2 \\ &\geq \nabla_{\zeta_{t+1}} f(\theta_t)^2 - c^2 h_{t+1}^2 \end{aligned}$$

We also have, invoking Assumption 5, that

$$\mathbb{E}_{\xi_{t+1}}[g_t^{(\zeta_{t+1})^2}] \leq \max_{k=1, \dots, p} \mathbb{E}_{\xi_{t+1}}[g_t^{(k)^2}] \leq 2\mathcal{L}(f(\theta_t) - f^*) + \sigma^2.$$

We finally obtain that

$$\begin{aligned} &\mathbb{E}_{\xi_{t+1}}[f(\theta_{t+1}) - f^*] \\ &\leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t)^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2. \end{aligned}$$

Apply Robbins-Siegmund Theorem to obtain that almost surely

$$\sum_{t \geq 0} \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t)^2 = \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t)\|_\infty^2 < \infty.$$

Proof of condition (ii). Note that from the proof of Theorem 1, we already have $\sum_{t \geq 0} \gamma_{t+1} (g_t(\omega) - \nabla f(\theta_t(\omega))) < \infty$, so using that

$$\|D(\zeta_{t+1})(g_t(\omega) - \nabla f(\theta_t(\omega)))\|_2 \leq \|g_t(\omega) - \nabla f(\theta_t(\omega))\|_2,$$

we deduce the convergence $\sum_{t \geq 0} \gamma_{t+1} D(\zeta_{t+1})(g_t(\omega) - \nabla f(\theta_t(\omega))) < \infty$ which gives, in virtue of Lemma 7 the result $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$. ■

Part (b) gradient weights. Here we assume that the weights of the coordinate sampling policy are proportional to any norm of the current gradient: $D_t \propto (|\partial_k f(\theta_t)|^q)_{1 \leq k \leq p}$ with $q > 0$. As before, the proof follows from applying Lemma 7. The proof of condition (i) relies on the equivalence of the norms in finite dimension.

Proof of condition (i). From the proof of Theorem 2, we get

$$\begin{aligned} &\mathbb{E}_{\xi_{t+1}}[f(\theta_{t+1}) - f^*] \\ &\leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t)^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2. \end{aligned}$$

Taking the expectation with respect to ζ_{t+1} , we get

$$\begin{aligned} &\mathbb{E}_t[f(\theta_{t+1}) - f^*] \\ &\leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \sum_{k=1}^p d_{t,k} \partial_k f(\theta_t)^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2. \end{aligned}$$

Apply Robbins-Siegmund Theorem to obtain $\sum_{t \geq 0} \gamma_{t+1} \nabla f(\theta_t)^\top D_t \nabla f(\theta_t) < \infty$ almost surely. Now observe that since $D_t \propto (|\partial_k f(\theta_t)|^q)_{1 \leq k \leq p}$, it means that for all $k = 1, \dots, p$ we have $d_{t,k} \propto |\partial_k f(\theta_t)|^q / \|\nabla f(\theta_t)\|_q^q$ and

$$\nabla f(\theta_t)^\top D_t \nabla f(\theta_t) = \sum_{k=1}^p d_{t,k} \partial_k f(\theta_t)^2 \propto \sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^q}{\|\nabla f(\theta_t)\|_q^q} \partial_k f(\theta_t)^2 \propto \frac{\|\nabla f(\theta_t)\|_{q+2}^{q+2}}{\|\nabla f(\theta_t)\|_q^q}.$$

All norms are equivalent on \mathbb{R}^p and using Hölder's inequality we have for $0 < l < q$ that $\|\cdot\|_l \leq p^{1/l-1/q} \|\cdot\|_q$ so the last term is lower bounded as

$$\frac{\|\nabla f(\theta_t)\|_{q+2}^{q+2}}{\|\nabla f(\theta_t)\|_q^q} \geq C \|\nabla f(\theta_t)\|_{q+2}^2 \quad \text{with } C = p^{-2/(q+2)},$$

and again using the equivalence of the norms we get the square of the infinity norm $\nabla f(\theta_t)^\top D_t \nabla f(\theta_t) \propto \|\nabla f(\theta_t)\|_\infty^2$ which finally proves

$$(i) \quad \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t)\|_\infty^2 < \infty.$$

Proof of condition (ii). It is the same as for *Part (a) maximum gradient*. We deduce the convergence $\sum_{t \geq 0} \gamma_{t+1} D(\zeta_{t+1})(g_t(\omega) - \nabla f(\theta_t(\omega))) < \infty$ which gives, in virtue of Lemma 7 the result $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$. ■

A.5 Proof of Theorem 3

Similarly to the proof of Theorem 1, we rely on Lemma 8 where g_{t-1} is replaced by $D(\zeta_t)g_{t-1}$. Therefore we need to check that, with probability 1, it holds that

$$(i) \quad \sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty \quad (ii) \quad \sum_{t \geq 0} \gamma_{t+1} (D(\zeta_t)g_t(\omega) - D_t \nabla f(\theta_t(\omega))) < \infty.$$

Proof of condition (i). From the proof of Theorem 2, we get

$$\begin{aligned} & \mathbb{E}_{\xi_{t+1}}[f(\theta_{t+1}) - f^*] \\ & \leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \nabla_{\zeta_{t+1}} f(\theta_t)^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2. \end{aligned}$$

Taking the expectation with respect to ζ_{t+1} and using that $\min_{k=1, \dots, d} d_{t,k} \geq \beta$ gives

$$\begin{aligned} & \mathbb{E}_t[f(\theta_{t+1}) - f^*] \\ & \leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \sum_{k=1}^p d_{t,k} \partial_k f(\theta_t)^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2 \\ & \leq (1 + L\mathcal{L}\gamma_{t+1}^2)(f(\theta_t) - f^*) - \gamma_{t+1} \beta \|\nabla f(\theta_t)\|_2^2 / 2 + c^2 \gamma_{t+1} h_{t+1}^2 / 2 + \frac{L}{2} \gamma_{t+1}^2 \sigma^2, \end{aligned}$$

and Robbins-Siegmund Theorem allows to conclude $\sum_{t \geq 0} \gamma_{t+1} \|\nabla f(\theta_t)\|_2^2 < +\infty$.

Proof of condition (ii). Again, we place ourselves on the event Ω_0 and omit the ω in notation for ease of clarity. First, note that $\|g_t\|_2^2 \leq g_t^{(\zeta_{t+1})^2} \leq \|g_t\|_2^2$. As a consequence, $\sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|D(\zeta_{t+1})g_t\|_2^2] \leq \sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|g_t\|_2^2]$ and this last series converges as shown in the proof of Theorem 1. Now observe that condition (ii) is satisfied as soon as

$$(a) \quad \left\| \sum_{t \geq 0} \gamma_{t+1} (D(\zeta_{t+1})g_t - \mathbb{E}_t[D(\zeta_{t+1})g_t]) \right\|_2 < \infty$$

$$(b) \quad \left\| \sum_{t \geq 0} \gamma_{t+1} (\mathbb{E}_t[D(\zeta_{t+1})g_t] - D_t \nabla f(\theta_t)) \right\|_2 < \infty$$

Note that equation (a) involves martingale increments whose quadratic variation satisfies

$$\sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|D(\zeta_{t+1})g_t - \mathbb{E}_t[D(\zeta_{t+1})g_t]\|_2^2] \leq \sum_{t \geq 0} \gamma_{t+1}^2 \mathbb{E}_t[\|D(\zeta_{t+1})g_t\|_2^2] < \infty,$$

which proves Equation (a). Finally the term in equation (b) is bounded using assumption 1 and $\|D_t\|_2 \leq 1$. We have $\sum_{t \geq 0} \gamma_{t+1}^2 \|\mathbb{E}_t[D(\zeta_{t+1})g_t] - D_t \nabla f(\theta_t)\|_2^2 \leq c^2 \sum_{t \geq 0} \gamma_{t+1}^2 h_t^2 < \infty$ which finally proves condition (ii) and gives, in virtue of Lemma 8 that $\nabla f(\theta_t) \rightarrow 0$ almost surely as $t \rightarrow +\infty$.

A.6 Proof of Theorem 4

From the proof of Theorem 3 and using β as a uniform lower bound on β_{t+1} , we have

$$\mathbb{E}_t[f(\theta_{t+1}) - f^*] \leq (1 + L\mathcal{L}\gamma_{t+1}^2) [f(\theta_t) - f^*] - \gamma_{t+1}\beta \|\nabla f(\theta_t)\|_2^2 + \frac{\sigma^2 L + c^2}{2} \gamma_{t+1}^2.$$

Inject the PL inequality $\|\nabla f(\theta_t)\|_2^2 \geq 2\mu(f(\theta_t) - f(\theta^*))$ from Assumption 7 to have

$$\mathbb{E}_t[f(\theta_{t+1}) - f^*] \leq (1 - 2\mu\beta\gamma_{t+1} + L\mathcal{L}\gamma_{t+1}^2) [f(\theta_t) - f^*] + \frac{\sigma^2 L + c^2}{2} \gamma_{t+1}^2.$$

Define $\delta_t = \mathbb{E}[f(\theta_t) - f^*]$ to finally obtain the recursion equation

$$\delta_t \leq (1 - 2\mu\beta\gamma_t + L\mathcal{L}\gamma_t^2) \delta_{t-1} + \frac{\sigma^2 L + c^2}{2} \gamma_t^2$$

Applying the same result from (Moulines and Bach, 2011) with the family of functions φ_α defined by $\varphi_\alpha(t) = \alpha^{-1}(t^\alpha - 1)$ if $\alpha \neq 0$ and $\varphi_\alpha(t) = \log(t)$ if $\alpha = 0$ along with the learning rates $\gamma_t = \gamma t^{-\alpha}$.

$$\delta_t \leq \begin{cases} 2 \exp(2L\mathcal{L}\gamma^2 \varphi_{1-2\alpha}(t)) \exp\left(-\frac{\mu\beta\gamma}{4} t^{1-\alpha}\right) \left(\delta_0 + \frac{\sigma^2 + 2c^2}{2\mathcal{L}}\right) + \frac{\gamma(\sigma^2 L + 2c^2)}{\mu\beta} t^{-\alpha} & \text{if } \alpha < 1 \\ 2 \exp(L\mathcal{L}\gamma^2) \left(\delta_0 + \frac{\sigma^2 + 2c^2}{2\mathcal{L}}\right) t^{-\mu\beta\gamma} + \left(\frac{\sigma^2 L}{2} + c^2\right) \gamma^2 \varphi_{\mu\beta\gamma/2-1}(t) t^{-\mu\beta\gamma/2} & \text{if } \alpha = 1 \end{cases}$$

A.7 Proof of Theorem 5

Starting from $G_0 = (0, \dots, 0)$, the total average gain G_n is updated in an online manner during the exploitation phase and collects all the empirical sums of the gradient gradient estimates as

$$G_n = \frac{1}{nT} \sum_{t=1}^{nT} D_t^{-1} D(\zeta_{t+1}) g(\theta_t, \xi_{t+1}), \quad \mathbb{E}[G_n] = \frac{1}{nT} \sum_{t=1}^{nT} \nabla f(\theta_t).$$

The goal is to show that $G_n \rightarrow 0$ using martingale properties. Thanks to Theorem 3, we have the almost sure convergence $\theta_t \rightarrow \theta^*$ which gives, since $\theta \mapsto \nabla f(\theta)$ is continuous, that $\nabla f(\theta_t) \rightarrow 0$ almost surely. Applying Cesaro's Lemma, it holds that $\mathbb{E}[G_n] \rightarrow 0$. It is enough to consider the difference $\left(G_n^{(k)} - \mathbb{E}\left[G_n^{(k)}\right]\right)$ for each $k \in \llbracket 1, p \rrbracket$. Introducing the martingale increments

$$\Delta_{t+1}^{(k)} = \frac{g(\theta_t, \xi_{t+1})^{(k)}}{d_t^{(k)}} \mathbb{1}_{\{\zeta_{t+1}=k\}} - \partial_k f(\theta_t), \quad \mathbb{E}\left[\Delta_{t+1}^{(k)} | \mathcal{F}_t\right] = 0.$$

It remains to show that, with probability 1,

$$G_n^{(k)} - \mathbb{E}\left[G_n^{(k)}\right] = \frac{1}{nT} \sum_{t=1}^{nT} \Delta_{t+1}^{(k)} \rightarrow 0.$$

Or equivalently, that, for each coordinate $k \in \llbracket 1, p \rrbracket$

$$\sum_{t=1}^{nT} \Delta_{t+1}^{(k)} = o(n). \quad (8)$$

The latter being a sum of martingale increments, we are in position to apply the strong law of large numbers for martingales which can be found as Assertion 2 of Theorem 1.18 in (Bercu et al., 2015). Using Assumption 5, there exist $0 \leq \mathcal{L}, \sigma^2 < \infty$ such that almost surely

$$\forall t \in \mathbb{N}, \quad \mathbb{E}\left[(g(\theta_t, \xi_{t+1})^{(k)})^2 | \mathcal{F}_t\right] \leq 2\mathcal{L}(f(\theta_t) - f^*) + \sigma^2.$$

Using the almost sure convergence $\theta_t \rightarrow \theta^*$, we deduce that there exists a compact set K which contains the sequence of iterates $(\theta_t)_{t \in \mathbb{N}}$ and using that f is continuous gives the upper bound

$$\forall k \in \llbracket 1, p \rrbracket \quad \mathbb{E}\left[(g(\theta_t, \xi_{t+1})^{(k)})^2 | \mathcal{F}_t\right] \leq M = 2\mathcal{L} \sup_{\theta \in K} (f(\theta) - f(\theta^*)) + \sigma^2.$$

Hence, the quadratic variation is bounded as follows

$$\begin{aligned} \sum_{t=1}^{nT} \mathbb{E}\left[(\Delta_{t+1}^{(k)})^2 | \mathcal{F}_t\right] &\leq \sum_{t=1}^{nT} \mathbb{E}\left[\left(\frac{g(\theta_t, \xi_{t+1})^{(k)}}{d_t^{(k)}}\right)^2 | \mathcal{F}_t\right] \\ &\leq (p/\lambda)^2 \sum_{t=1}^{nT} \mathbb{E}[(g(\theta_t, \xi_{t+1})^{(k)})^2 | \mathcal{F}_t] \\ &\leq (p/\lambda)^2 nTM. \end{aligned}$$

Equation (8) follows from applying the previously mentioned law of large numbers.

B. Additional Results

B.1 Almost sure convergence under stronger assumptions

Similar to Gadat et al. (2018), we consider some stronger assumptions where the function f is coercive and there exists a unique stationary point θ^* . In such framework, the sequences of iterates $(\theta_t)_{t \geq 0}$ obtained by both SGD and SCGD satisfy $\theta_t \rightarrow \theta^*$ almost surely as $t \rightarrow +\infty$.

f is coercive and $\{\theta \in \mathbb{R}^p : \nabla f(\theta) = 0\} = \{\theta^*\}$. Following the proofs of Theorems 1 and 3, we may apply Robbins-Siegmund Theorem. There exists an event $\Omega_0 \subset \Omega$ such that, $\mathbb{P}(\Omega_0) = 1$ and for every $\omega \in \Omega_0$, $\limsup_t f(\theta_t(\omega)) < \infty$ and the series $\sum_t \eta_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2$ converges (where $\eta_t = \gamma_t$ for SGD and $\eta_t = \gamma_t \beta_t$ for SCGD). Since $\lim_{\|\theta\| \rightarrow \infty} f(\theta) = \infty$, we deduce that for every $\omega \in \Omega_0$, the sequence $(\theta_t(\omega))_{t \geq 0}$ is bounded in \mathbb{R}^p . Therefore the limit set $\chi_\infty(\omega)$ (set of accumulation points) of the sequence $(\theta_t(\omega))$ is non-empty. The convergence of the series $\sum_t \eta_{t+1} \|\nabla f(\theta_t(\omega))\|_2^2 < \infty$ along with the condition $\sum_t \eta_{t+1} = +\infty$ only imply that : $\liminf_{t \rightarrow \infty} \|\nabla f(\theta_t(\omega))\|_2^2 = 0$, $\mathbb{P} - a.s.$

Hence, since $\theta \mapsto \nabla f(\theta)$ is continuous, there exists a limit point $\theta_\infty(\omega) \in \chi_\infty(\omega)$ such that $\|\nabla f(\theta_\infty(\omega))\|_2^2 = 0$, *i.e.*, $\nabla f(\theta_\infty(\omega)) = 0$. Because the set of solutions $\{\theta \in \mathbb{R}^p, \nabla f(\theta) = 0\}$ is reduced to the singleton $\{\theta^*\}$, we have $\theta_\infty(\omega) = \theta^*$. Since $(f(\theta_t(\omega)))$ converges, it implies that $\lim_t f(\theta_t(\omega)) = f^*$ and for every limit point $\theta \in \chi_\infty(\omega)$, we have $f(\theta) = f^*$. Since the set $\{\theta \in \mathbb{R}^p, f(\theta) = f^*\}$ is equal to $\{\theta^*\}$, the limit set $\chi_\infty(\omega)$ is also reduced to $\{\theta^*\}$.

B.2 Almost sure convergence of MUSKETEER

By definition, we have for all $k \in \llbracket 1, p \rrbracket$,

$$d_{t+1}^{(k)} = (1 - \lambda_t) \varphi(G_t)^{(k)} + \lambda_t \frac{1}{p}$$

implying that $\beta_{t+1} = \min_{k \in \llbracket 1, p \rrbracket} d_t^{(k)} \geq \lambda_t/p$. As a consequence, as soon as $\sum_{t \geq 1} \lambda_t \gamma_t = +\infty$, the assumption $\sum_{t \geq 1} \beta_t \gamma_t = +\infty$ is satisfied. Applying Theorem 3 we obtain the almost sure convergence of MUSKETEER. The condition $\sum_{t \geq 1} \lambda_t \gamma_t = +\infty$ is easily satisfied with a fixed value $\lambda_t \equiv \lambda$ in the mixture update and one can also use a slowly decreasing sequence, *e.g.* $\lambda_t = 1/\log(t)$.

B.3 Regret analysis in the convex case

In order to better understand the benefits of the adaptive sampling strategies over standard uniform sampling, let us consider a particular setting where the objective function f is convex. The following proposition available in Namkoong et al. (2017) presents a regret analysis which is useful for interpretability.

Proposition 9 (Regret analysis for convex f and unbiased estimates) *Assume that f is convex and consider the sequence of iterates obtained by $\theta_{t+1} = \theta_t - \gamma D_t^{-1} D(\zeta_{t+1}) g_t$ with constant step size $\gamma > 0$. We have*

$$\mathbb{E} \left[f \left(\frac{1}{T} \sum_{t=1}^T \theta_t \right) - f(\theta^*) \right] \leq \frac{\|\theta^*\|^2}{2\gamma T} + \frac{\gamma}{2T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^2}{d_t^{(k)}} \right].$$

Proof Assume that the objective f is convex and consider the average estimate $\bar{\theta}_T = \frac{1}{T} \sum_{t=1}^T \theta_t$, along with the following quantity: $S(f, \hat{\theta}) = \mathbb{E}[f(\hat{\theta})] - f^*$. Using convexity we have on the one hand $f(\theta_t) - f^* \leq \langle \theta_t - \theta^*, \nabla f(\theta_t) \rangle$ and on the other hand

$$f(\bar{\theta}_T) - f^* \leq \frac{1}{T} \sum_{t=1}^T (f(\theta_t) - f^*)$$

which give together the following upper bound

$$f(\bar{\theta}_T) - f^* \leq \frac{1}{T} \sum_{t=1}^T \langle \theta_t - \theta^*, \nabla f(\theta_t) \rangle.$$

Using an unbiased gradient estimate v_t , *i.e.* $\mathbb{E}_t[v_t] = \nabla f(\theta_t)$, we can write

$$\mathbb{E}[f(\bar{\theta}_T)] - f^* \leq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \langle \theta_t - \theta^*, \mathbb{E}_t[v_t] \rangle \right].$$

The term in the expectation is bounded using Lemma 11 with $v_t = D_t^{-1} D(\zeta_{t+1}) g_t$ as

$$\frac{1}{T} \sum_{t=1}^T \langle \theta_t - \theta^*, v_t \rangle \leq \frac{\|\theta^*\|^2}{2\gamma T} + \frac{\gamma}{2T} \sum_{t=1}^T \|D_t^{-1} D(\zeta_{t+1}) g_t\|^2.$$

Take the expectation on both side to control the regret as

$$S(f, \bar{\theta}_T) \leq \frac{\|\theta^*\|^2}{2\gamma T} + \frac{\gamma}{2T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^2}{d_t^{(k)}} \right].$$

■

The term in expectation should be minimized with respect to the probability weights $d_t^{(k)}$. Intuitively, in order to maintain the overall sum as small as possible, the large gradient coordinates should be sampled more often, *i.e.* we would like to have $d_t^{(k)}$ large whenever $|\partial_k f(\theta_t)|^2$ is large. This is in line with the framework of coordinate smoothness discussed in Remark 4 and the work of Allen-Zhu et al. (2016).

(Uniform Coordinate Sampling) For all $k \in \llbracket 1, p \rrbracket$, we have $d_t^{(k)} = 1/p$ so that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^2}{d_t^{(k)}} \right] = \frac{p}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p |\partial_k f(\theta_t)|^2 \right] = \frac{p}{T} \sum_{t=1}^T \mathbb{E} [\|\nabla f(\theta_t)\|^2].$$

(MUSKETEER) For all $k \in \llbracket 1, p \rrbracket$, we have $d_t^{(k)} = (1 - \lambda_{t-1})\varphi(G_{t-1})^{(k)} + \lambda_{t-1}/p$ so that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^2}{d_t^{(k)}} \right] = \frac{p}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^p \frac{|\partial_k f(\theta_t)|^2}{(1 - \lambda_{t-1})p\varphi(G_{t-1})^{(k)} + \lambda_{t-1}} \right],$$

where the denominator is strictly larger than 1 for all the coordinates associated to large gains. Indeed, let $k \in \llbracket 1, p \rrbracket$ the index of such coordinate. Since it is a rewarding coordinate, the normalizing step implies that $\varphi(G_{t-1})^{(k)} > 1/p$ and $(1 - \lambda_{t-1})p\varphi(G_{t-1})^{(k)} + \lambda_{t-1} > 1$. This property translates the adaptive nature of the probability weights used in the MUSKETEER strategy.

B.4 Auxiliary Results

Theorem 10 (Robbins and Siegmund, 1971) Consider a filtration $(\mathcal{F}_n)_{n \geq 0}$ and four sequences of random variables $(V_n)_{n \geq 0}$, $(W_n)_{n \geq 0}$, $(a_n)_{n \geq 0}$ and $(b_n)_{n \geq 0}$ that are adapted and non-negative. Assume that almost surely $\sum_k a_k < \infty$ and $\sum_k b_k < \infty$. Assume moreover that $\mathbb{E}[V_0] < \infty$ and $\forall n \in \mathbb{N} : \mathbb{E}[V_{n+1} | \mathcal{F}_n] \leq (1 + a_n)V_n - W_n + b_n$. Then it holds

$$(a) \sum_k W_k < \infty \text{ a.s.} \quad (b) V_n \xrightarrow{\text{a.s.}} V_\infty, \mathbb{E}[V_\infty] < \infty. \quad (c) \sup_{n \geq 0} \mathbb{E}[V_n] < \infty.$$

Lemma 11 Let $\theta_1, \dots, \theta_T$ be an arbitrary sequence of vectors. Any algorithm with initialization $\theta_1 = 0$ and update rule $\theta_{t+1} = \theta_t - \gamma v_t$ satisfies

$$\sum_{t=1}^T \langle \theta_t - \theta^*, v_t \rangle \leq \frac{\|\theta^*\|^2}{2\gamma} + \frac{\gamma}{2} \sum_{t=1}^T \|v_t\|^2.$$

In particular, for $B, \rho > 0$, if we have $\|v_t\| \leq \rho$ and we set $\gamma = \sqrt{B^2/(\rho^2 T)}$ then for every θ^* with $\|\theta^*\| \leq B$, we have $T^{-1} \sum_{t=1}^T \langle \theta_t - \theta^*, v_t \rangle \leq B\rho/\sqrt{T}$.

Proof Using algebraic manipulations (completing the square), we obtain:

$$\begin{aligned} \langle \theta_t - \theta^*, v_t \rangle &= \frac{1}{\gamma} \langle \theta_t - \theta^*, \gamma v_t \rangle \\ &= \frac{1}{2\gamma} \left(-\|\theta_t - \theta^* - \gamma v_t\|^2 + \|\theta_t - \theta^*\|^2 + \gamma^2 \|v_t\|^2 \right) \\ &= \frac{1}{2\gamma} \left(-\|\theta_{t+1} - \theta^*\|^2 + \|\theta_t - \theta^*\|^2 \right) + \frac{\gamma}{2} \|v_t\|^2 \end{aligned}$$

where the last equality follows from the definition of the update rule. Summing the equality over t , we have

$$\sum_{t=1}^T \langle \theta_t - \theta^*, v_t \rangle = \frac{1}{2\gamma} \sum_{t=1}^T \left(-\|\theta_{t+1} - \theta^*\|^2 + \|\theta_t - \theta^*\|^2 \right) + \frac{\gamma}{2} \sum_{t=1}^T \|v_t\|^2$$

The first sum on the right-hand side is a telescopic sum that collapses to $\|\theta_1 - \theta^*\|^2 - \|\theta_{T+1} - \theta^*\|^2$. Then we have

$$\begin{aligned} \sum_{t=1}^T \langle \theta_t - \theta^*, v_t \rangle &= \frac{1}{2\gamma} \left(\|\theta_1 - \theta^*\|^2 - \|\theta_{T+1} - \theta^*\|^2 \right) + \frac{\gamma}{2} \sum_{t=1}^T \|v_t\|^2 \\ &\leq \frac{1}{2\gamma} \|\theta_1 - \theta^*\|^2 + \frac{\gamma}{2} \sum_{t=1}^T \|v_t\|^2 \\ &= \frac{1}{2\gamma} \|\theta^*\|^2 + \frac{\gamma}{2} \sum_{t=1}^T \|v_t\|^2 \end{aligned}$$

where the last equality is due to the definition $\theta_1 = 0$. This proves the first part of the lemma. The second part follows by upper bounding $\|\theta^*\|$ by B , $\|v_t\|$ by ρ , dividing by T , and plugging in the value of γ ■

C. Illustrative Example (stochastic first order)

We perform a comparison on a simple example in dimension $p = 2$ with the functions $f(x, y) = (x^2 + y^2)/2$ and $h(x, y) = x^2/2$. Note that the function h only depends on the first coordinate and an adaptive coordinate descent method should favor this direction. Figure 4 presents the optimization paths of the different methods: SGD, Uniform and MUKSTEER. With the function f which does not present any particular design or favorable descent direction, the Uniform and Musketeer policies perform as good as classical SGD. More interestingly, when dealing with the function h , our method MUSKETEER (red) finds that the horizontal direction associated to axis (Ox) is the relevant one for optimization. After collecting some information during the exploration phase, the probability weights got updated to favor the horizontal direction, leading to a faster convergence. For a visual demonstration of these optimization paths, please refer to the mp4-files *optimize_f.mp4* and *optimize_h.mp4* available in the supplementary material².

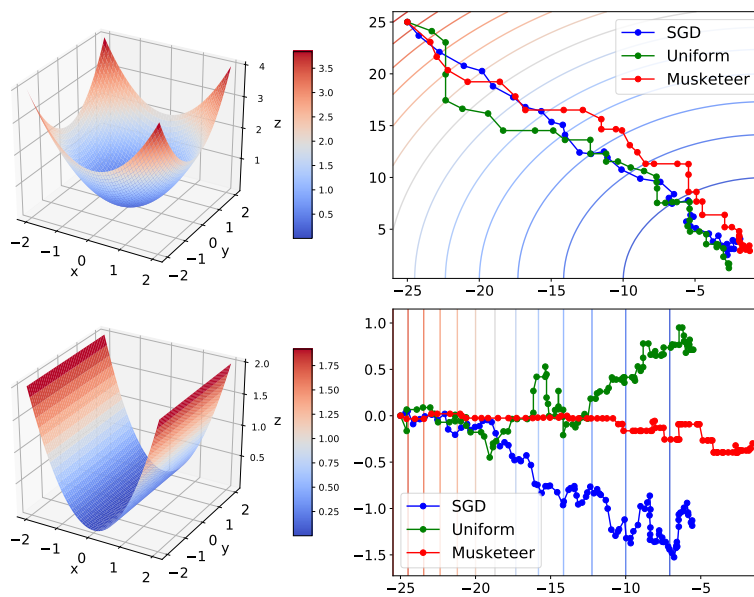


Figure 4: Comparison of SGD/Uniform/Musketeer on simple 2D-examples

D. Numerical Experiments Details

D.1 Regularized linear models

We consider the ERM paradigm with linear models, namely regularized regression problems with objectives of the form $f(\theta) = (1/n) \sum_{i=1}^n f_i(\theta) + \mu \|\theta\|^2$. Similarly to (Namkoong et al., 2017), we endow the data matrix X with a block structure. The columns are drawn as $X[:, k] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, p \rrbracket$. The parameters are set to $n = 10,000$ samples in dimension $p = 250$ with an exploration size equal to $T = \lfloor \sqrt{p} \rfloor = 15$. The regularization parameter is set to the classical value $\mu = 1/n$. We update the parameter

². <https://github.com/RemiLELUC/SCGD-Musketeer>

vector with the optimal learning rate $\gamma_k = \gamma/(k + k_0)$ in the experiments. Other learning rates in the framework of stochastic first order methods are considered in Appendix G.

- (zeroth-order) For the Ridge regression, we set $\gamma = 3, k_0 = 10$ and for the logistic regression $\gamma = 10, k_0 = 5$. The gradient estimate g is computed using queries of a function f_i where $i \sim \mathcal{U}(\llbracket 1, n \rrbracket)$. We use the ℓ_1 -reweighting with $\lambda_t = 1/\log(t)$ or softmax with $\lambda_n \equiv 0.5$, which both satisfy Assumption 6.
- (first order) The learning rate is equal to $\gamma_k = 1/k$ ($\gamma = 1, k_0 = 0$). The gradient estimate g is computed using mini-batches of size 8. The weighting parameter $\eta > 0$ in the softmax part of the probability weights is set to $\eta = 1$ and the parameter λ in Equation (7) is chosen as $\lambda_t = 1/\log(t)$ which satisfies the extended Robbins-Monro condition 6.

D.2 Neural Networks

Dataset description and parameter configuration. The three datasets in the experiments are popular publicly available deep learning datasets. The underlying machine learning task is the one of multi-label classification.

- **MNIST** (Deng, 2012): a database of handwritten digits with a training set of 60,000 examples and a test set of 10,000 examples. The digits have been size-normalized and centered in a fixed-size image. The original black and white (bilevel) images from NIST were size normalized to fit in a 20x20 pixel box while preserving their aspect ratio. The resulting images contain grey levels as a result of the anti-aliasing technique used by the normalization algorithm. The images were centered in a 28x28 image by computing the center of mass of the pixels, and translating the image so as to position this point at the center of the 28x28 field. Each training and test example is assigned to the corresponding handwritten digit between 0 and 9.

- **Fashion-MNIST** (Xiao et al., 2017): a dataset of Zalando’s article images, composed of a training set of 60,000 examples and a test set of 10,000 examples. Each example is a 28x28 grayscale image, associated with a label from 10 classes. It shares the same image size and structure of training and testing splits as the MNIST database. Each training and test example is assigned to one of the following labels: T-shirt/top (0); Trouser (1); Pullover (2); Dress (3); Coat (4); Sandal (5); Shirt (6); Sneaker (7); Bag (8); Ankle boot (9).

- **Kuzushiji-MNIST**: This dataset is a drop-in replacement for the MNIST dataset (28x28 grayscale, 70,000 images), provided in the original MNIST format as well as a NumPy format. Since MNIST is restricted to 10 classes, one character here represents each of the 10 rows of Hiragana when creating Kuzushiji-MNIST.

- **CIFAR10** (Krizhevsky et al., 2009): The CIFAR-10 dataset consists of 60,000 32×32 colour images in 10 classes, with 6,000 images per class. There are 50,000 training images and 10,000 test images. The dataset is divided into five training batches and one test batch, each with 10,000 images. The test batch contains exactly 1,000 randomly-selected images from each class. The training batches contain the remaining images in random order, but some training batches may contain more images from one class than another. Between them, the training batches contain exactly 5,000 images from each class. Each training and test example is assigned to one of the following labels: airplane (0); automobile (1); bird (2); cat (3); deer (4); dog (5); frog (6); horse (7); ship (8); truck (9).

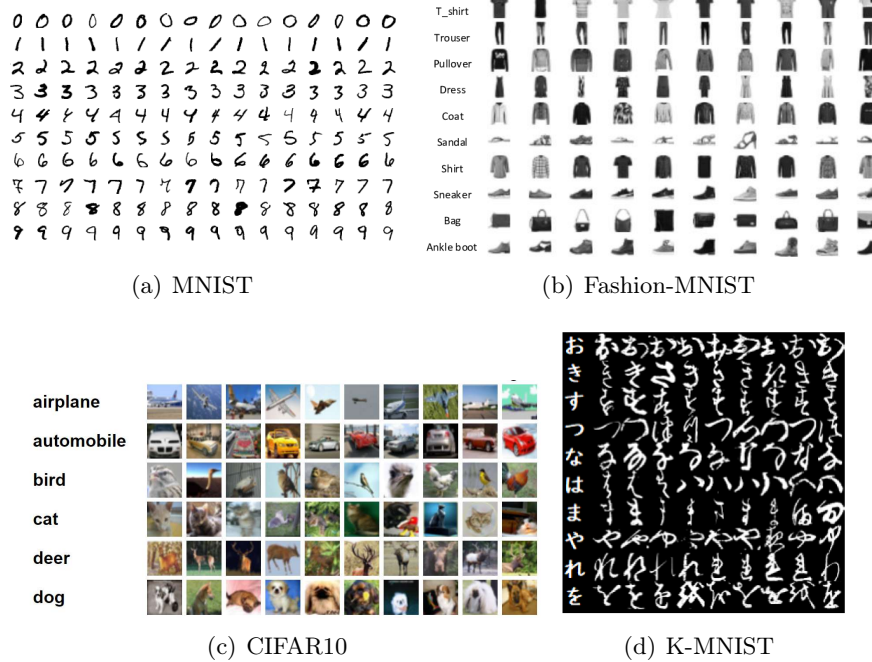


Figure 5: Samples for Mnist, Fashion-Mnist, K-Mnist and CIFAR-10.

Two different neural networks are used in the experiments: one with linear layers for MNIST, Fashion-MNIST, K-MNIST another one with convolutional layers for CIFAR10. For the first network, the total number of parameters is $p = 55,050$. For the second network, the dimension is $p = 64,862$. In both cases, the exploration size is $T = \lfloor \sqrt{p} \rfloor$. In the experiments with stochastic first order methods, we use batches of coordinates with $m = p/10$.

D.3 Hyperparameters and Hardware.

Hyperparameters. When training neural networks with linear layers, we use: `batch_size` = 32; `input_size` = 28*28; `hidden_size` = 32; `output_size` = 64, along with the parameters

- (zeroth-order) $\gamma = 10$ (Mnist and Fashion-Mnist) $\gamma=15$ (Kmnist); $h = 0.01$; ℓ_1 normalization with $\lambda_n = 1/\log(n)$; softmax normalization with $\lambda_n \equiv 0.2$ and $\eta = 5$.

- (first order) $\gamma = 0.01$ (Mnist,Fashion-Mnist,Cifar10); normalization = softmax with $\eta \in \{1, 2, 10\}$; $\lambda_t = 0$ (only exponential weights).

Hardware. The experiments of linear models are run using a processor Intel Core i7-10510U CPU 1.80GHz \times 8; the neural networks are trained using GPU from Google Colab (GPU: Nvidia K80 / T4; GPU Memory: 12GB/16GB; GPU Memory Clock: 0.82GHz/1.59GHz; Performance: 4.1 TFLOPS / 8.1 TFLOPS)

ZO Neural Networks with ℓ_1 normalization.

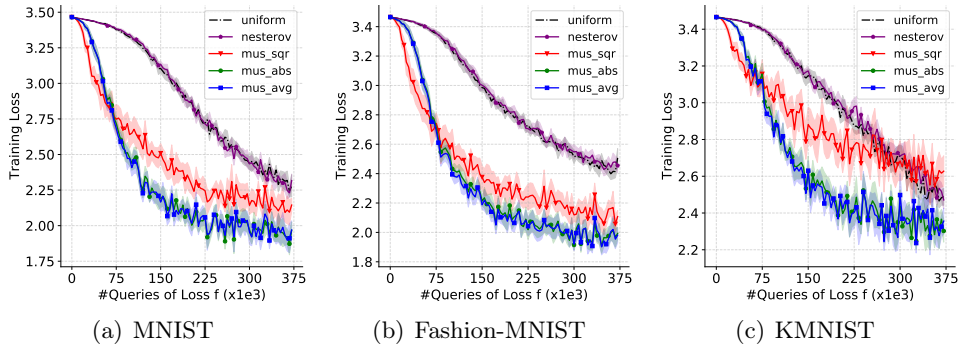


Figure 6: Training Loss ZO Neural Networks with ℓ_1 normalization.

ZO Neural Networks, Comparison of ℓ_1 and Softmax normalizations.

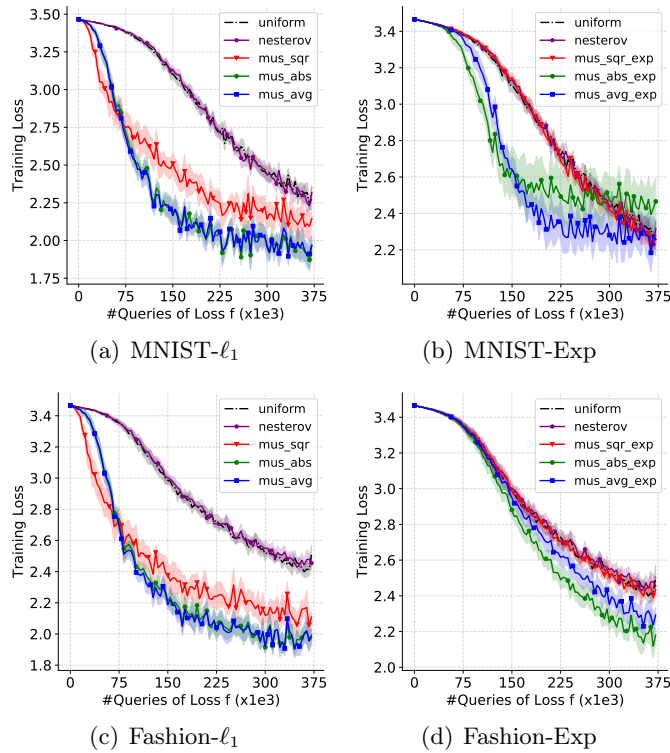


Figure 7: Training Loss ZO Neural Networks with ℓ_1 and Softmax normalizations.

E. Numerical Experiments with stochastic first order methods

In this section, we empirically validate the SCGD framework by running MUSKETEER and competitors on synthetic and real datasets problems with stochastic first order methods. First, we focus on ridge regression and regularized logistic regression problems adopting the data generation process of (Namkoong et al., 2017) in which the covariates exhibit a certain block structure. Second, MUSKETEER is employed to train different neural networks models on real datasets for multi-label classification task. From a practical point of view, the optimization procedure is implemented through a PyTorch optimizer which allows an easy deployment and integration.

Methods in competition. The set of methods in competition is restricted to stochastic coordinate-based methods along with standard SGD playing the role of the baseline. This choice allows an honest comparison as the parameter tuning can be the same for all methods. MUSKETEER is implemented according to Section 4 with an exploration size $T = \lfloor \sqrt{p} \rfloor$ and different values of η are used to feed the discussion on the adaptiveness. The method UNIFORM stands for the uniform coordinate sampling policy in SCGD. The method ADAPTIVE is the importance sampling based method described in Remark 3. This method is no longer part of the SCGD framework and corresponds to the one developed in (Wangni et al., 2018). Among the different methods, MUSKETEER is the only one exhibiting a bias when generating gradients. In all cases, $\theta_0 = (0, \dots, 0)^\top \in \mathbb{R}^p$ and the optimal SGD learning rate $\gamma_k = 1/k$ is used. For a fair comparison of SGD against SCGD, we normalize the number of passes over the coordinates: one SGD step updates the p coordinates of θ so we allow to take p steps for the coordinate-based methods in the mean time.

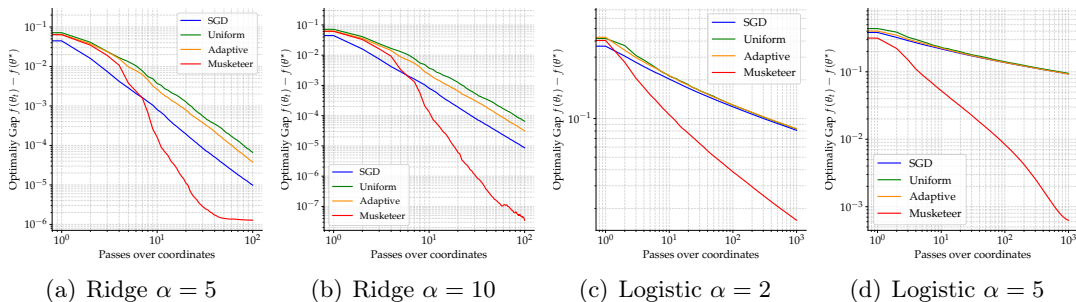


Figure 8: $[f(\theta_t) - f(\theta^*)]$ for Linear Models on Synthetic data with different block structures.

Linear models. We apply ERM to regularized regression and classification problems. Similarly to (Namkoong et al., 2017), we endow the data matrix X with a block structure. The columns are drawn as $X[:, k] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for $k \in \llbracket 1, p \rrbracket$. The parameters are set to $n = 10,000$ samples in dimension $p = 250$ and $T = 15$. Figure 8 provides the graphs of the optimality gap $t \mapsto f(\theta_t) - f^*$ averaged over 20 independent simulations for different values of $\alpha \in \{2; 5; 10\}$. First, note that the uniform sampling strategy shows similar performance to the classical SGD and that the (unbiased) importance sampling version ADAPTIVE is also of the same order. Besides, the clear winner is MUSKETEER as it offers the best performance in all configurations. Greater values of α (stronger block structure)

improve our relative performance with respect to the other methods as shown by Figures 8(b) and 8(d).

Neural Networks. To assess the practical performance of MUSKETEER, we focus on the training of neural networks within the framework of multi-label classification. The datasets in the experiments are popular publicly available deep learning datasets: MNIST (Deng, 2012), Fashion-MNIST (Xiao et al., 2017) and CIFAR10 (Krizhevsky et al., 2009). Given an image, the goal is to predict its label among ten classes. Two different neural networks are used in the experiments: one with linear layers for MNIST and Fashion-MNIST ($p = 55,050$ and $T = 234$), another one with convolutional layers for CIFAR10 ($p = 64,862$ and $T = 254$).

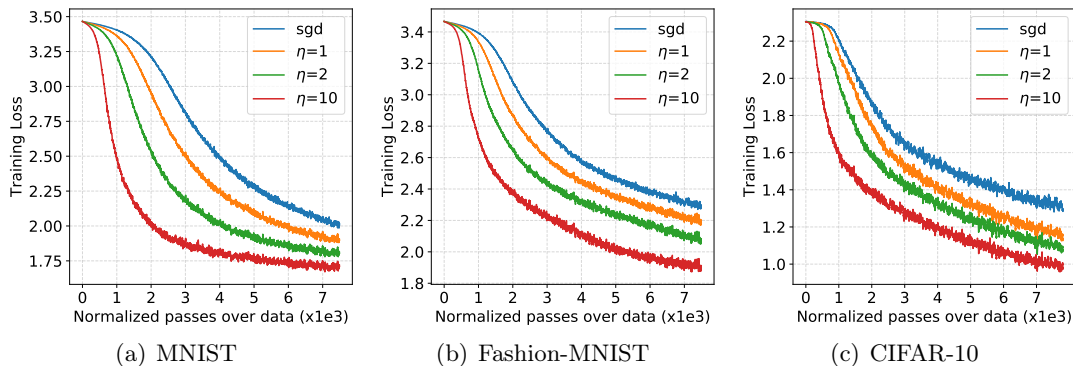


Figure 9: Training Loss of SGD vs. MUSKETEER on real-world datasets.

Figure 9 compares the evolution of the training loss of SGD against MUSKETEER averaged over 10 independent simulations with different values of η . A great value of this parameter strengthens the adaptive scheme as it gives more importance to the weights in Equation (7), leading to stronger decrease of the objective function. Interestingly, the performance of MUSKETEER also benefit from such adaptive structure in terms on accuracy of the test set (see Table 1). This allows to quantify the statistical gain brought by MUSKETEER over SGD.

	SGD	$\eta = 1$	$\eta = 2$	$\eta = 10$
MNIST	84.7 \pm 1.0	86.7 \pm 0.5	88.9 \pm 0.4	91.3\pm0.2
FASHION	64.7 \pm 1.2	68.5 \pm 1.0	71.2 \pm 0.7	77.1\pm0.8
CIFAR10	51.4 \pm 1.4	57.7 \pm 0.8	59.7 \pm 1.0	62.7\pm0.8

Table 1: Test Accuracy (in %).

F. Further Numerical Experiments with zeroth-order methods

F.1 Ridge Regression (ℓ_1 -reweighting) with different settings of (n, p)

We consider the Ridge regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB+1 : kB+B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 0, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 10$ for the Ridge regression. The parameter α represents the block structure and is set to $\alpha = 5$. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs. The learning rates is the same for all methods, fixed to $\gamma_k = 1/(k + 10)$. The different settings are: number of samples $n \in \{1,000; 2,000; 5,000\}$ and dimension $p \in \{20; 50; 100; 200\}$. We use the ℓ_1 normalization in Equation (7) with $\lambda_n = 1/\log(n)$.

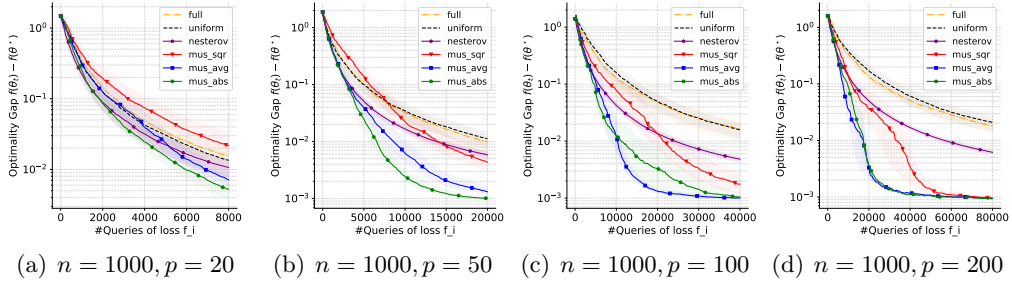


Figure 10: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 1000$ and $p = 20, 50, 100, 200$

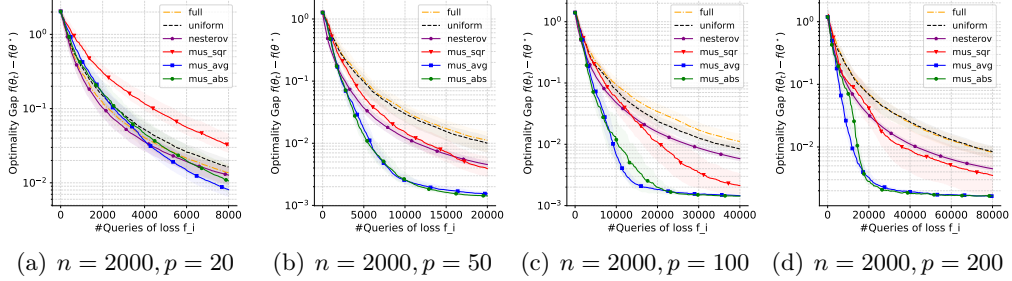


Figure 11: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 2000$ and $p = 20, 50, 100, 200$

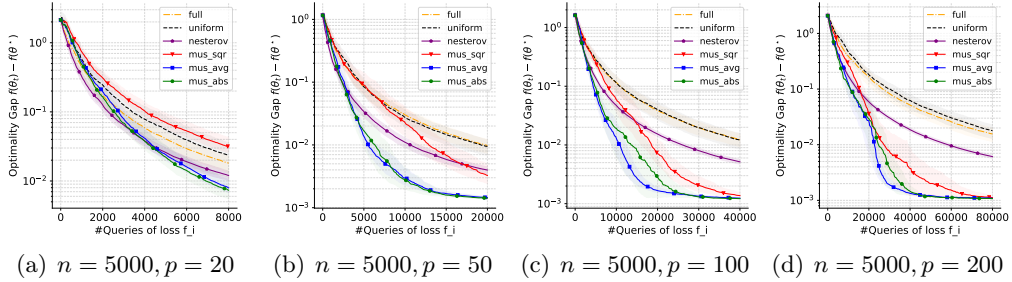
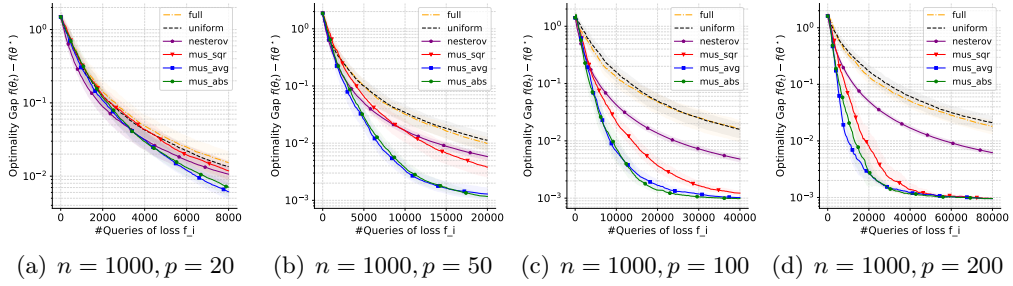
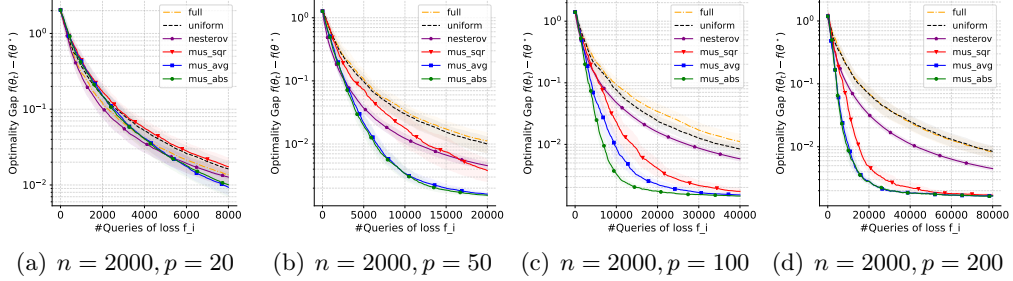
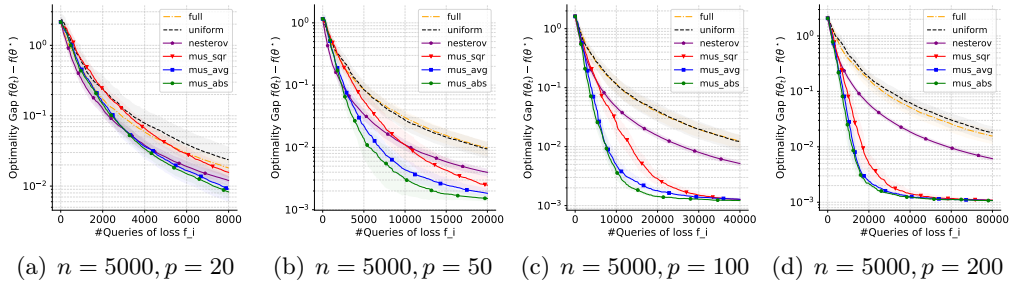


Figure 12: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 5000$ and $p = 20, 50, 100, 200$

F.2 Ridge Regression (softmax reweighting) with different settings of (n, p)

We consider the Ridge regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB+1 : kB+B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 0, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 10$ for the Ridge regression. The parameter α represents the block structure and is set to $\alpha = 5$. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs. The learning rates is the same for all methods, fixed to $\gamma_k = 1/(k + 10)$. The different settings are: number of samples $n \in \{1, 000; 2, 000; 5, 000\}$ and dimension $p \in \{20; 50; 100; 200\}$. We use the softmax normalization in Equation (7) with $\lambda_n \equiv 0.5$ and $\eta = 1$.


 Figure 13: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 1000$ and $p = 20, 50, 100, 200$

 Figure 14: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 2000$ and $p = 20, 50, 100, 200$

 Figure 15: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 5000$ and $p = 20, 50, 100, 200$

F.3 Logistic Regression (ℓ_1 -reweighting) with different settings of (n, p)

We consider the ℓ_2 -Logistic regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB+1 : kB+B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 5$ for the Logistic regression. The parameter α represents the block structure and is set to $\alpha = 5$. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs. The learning rates is the same for all methods, fixed to $\gamma_k = 10/(k + 5)$. The different settings are: number of samples $n \in \{1,000; 2,000; 5,000\}$ and dimension $p \in \{20; 50; 100; 200\}$. We use the ℓ_1 normalization in Equation (7) with $\lambda_n = 1/\log(n)$.

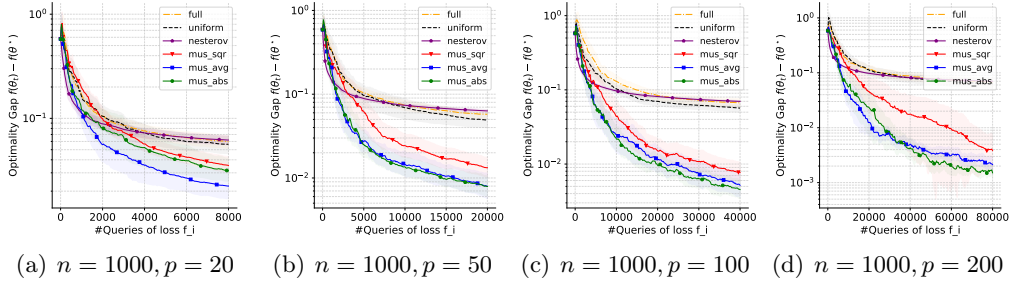


Figure 16: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 1000$ and $p = 20, 50, 100, 200$

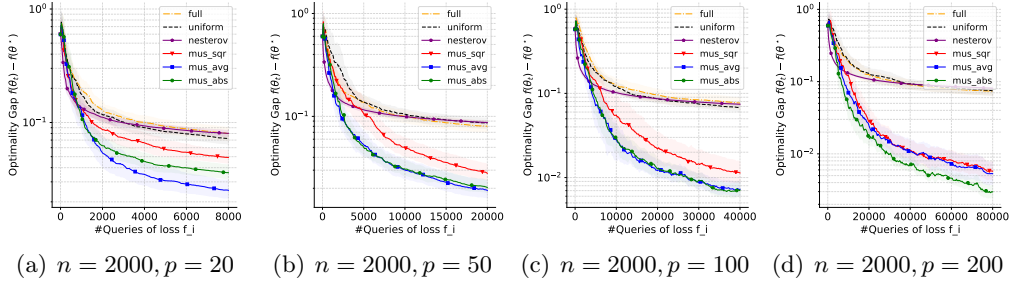


Figure 17: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 2000$ and $p = 20, 50, 100, 200$


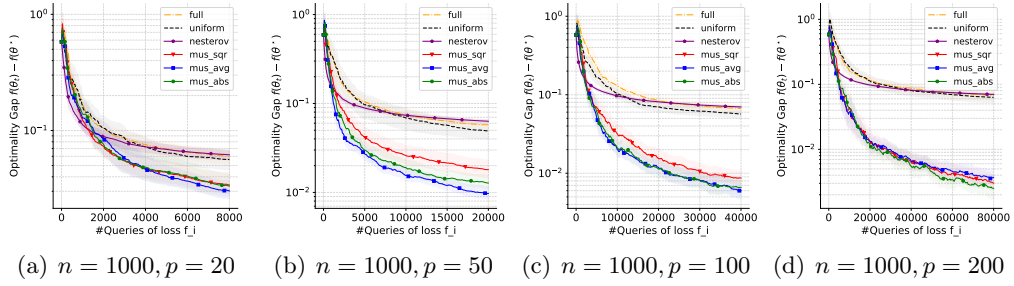
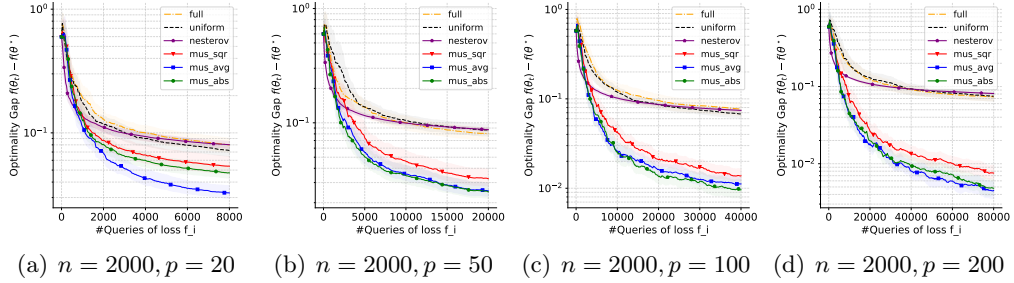



Figure 18: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 5000$ and $p = 20, 50, 100, 200$

F.4 Logistic Regression (softmax reweighting) with different settings of (n, p)

We consider the ℓ_2 -Logistic regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB+1 : kB+B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 5$ for the Logistic regression. The parameter α represents the block structure and is set to $\alpha = 5$. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs. The learning rates is the same for all methods, fixed to $\gamma_k = 10/(k + 5)$. The different settings are: number of samples $n \in \{1, 000; 2, 000; 5, 000\}$ and dimension $p \in \{20; 50; 100; 200\}$. We use the softmax normalization in Equation (7) with $\lambda_n \equiv 0.5$ and $\eta = 1$.


 Figure 19: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 1000$ and $p = 20, 50, 100, 200$

 Figure 20: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 2000$ and $p = 20, 50, 100, 200$

 Figure 21: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 5000$ and $p = 20, 50, 100, 200$

F.5 Effect of Importance Sampling (IS) on Ridge Regression

We consider the same setting as in Subsection F.1 and study the effect of using importance sampling weights in the update rule of MUSKETEER. Indeed, MUSKETEER update rule is defined with the following biased gradient estimate $\theta_{t+1} = \theta_t - \gamma_{t+1}D(\zeta_{t+1})g_t$ and the importance sampling (IS) strategy consists in adding D_t^{-1} to reach an unbiased estimate

$$\theta_{t+1} = \theta_t - \gamma_{t+1}D_t^{-1}D(\zeta_{t+1})g_t.$$

For the different configurations, we compare the MUSKETEER methods with their importance sampling counterparts. The Figures below show that the importance sampling methods perform similarly to the uniform coordinate sampling strategy and are therefore sub-optimal.

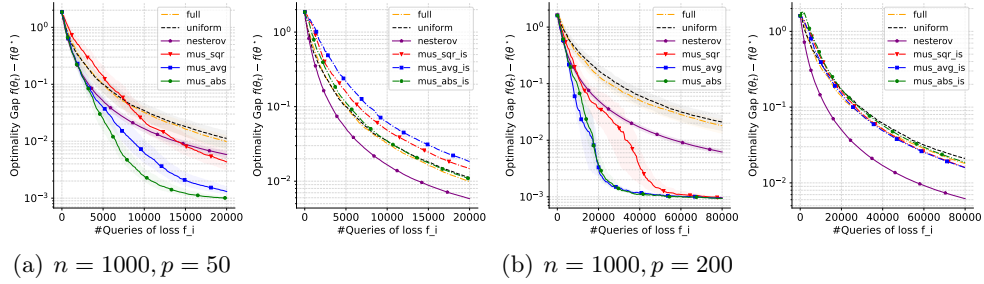


Figure 22: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 1000$ and $p = 50, 200$

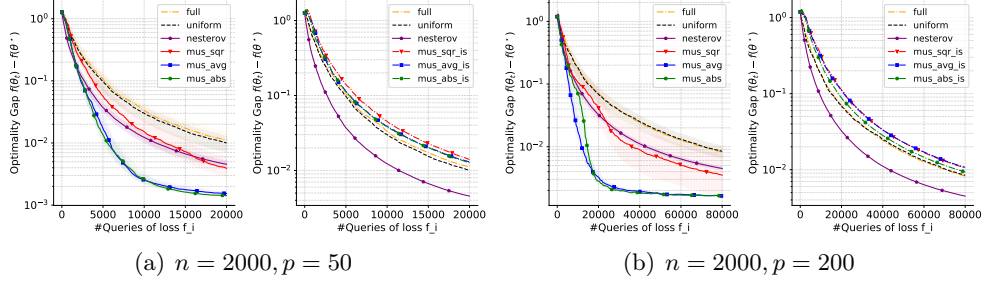


Figure 23: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 2000$ and $p = 50, 200$

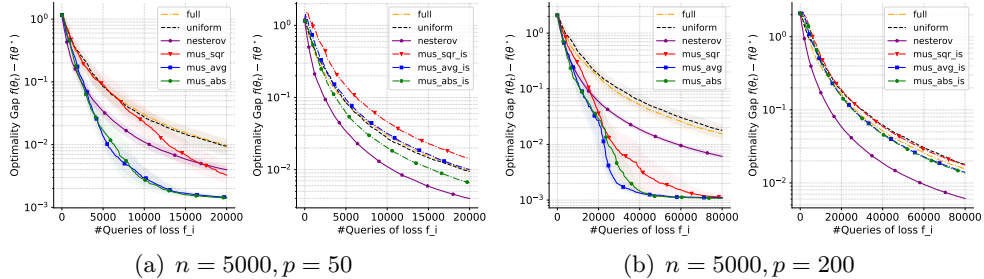


Figure 24: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 5000$ and $p = 50, 200$

F.6 Effect of Importance Sampling (IS) on Logistic Regression

We consider the same setting as in Subsection F.3 and study the effect of using importance sampling weights in the update rule of MUSKETEER. Indeed, MUSKETEER update rule is defined with the following biased gradient estimate $\theta_{t+1} = \theta_t - \gamma_{t+1}D(\zeta_{t+1})g_t$ and the importance sampling (IS) strategy consists in adding D_t^{-1} to reach an unbiased estimate

$$\theta_{t+1} = \theta_t - \gamma_{t+1}D_t^{-1}D(\zeta_{t+1})g_t.$$

For the different configurations, we compare the MUSKETEER methods with their importance sampling counterparts. The Figures below show that the importance sampling methods perform similarly to the uniform coordinate sampling strategy and are therefore sub-optimal.

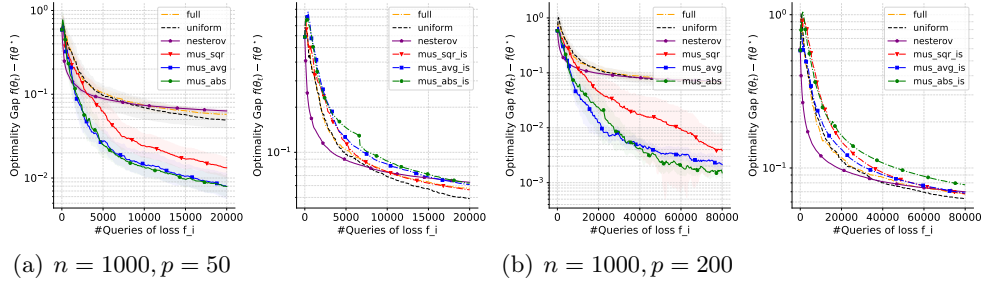


Figure 25: $[f(\theta_t) - f^*]$ for Logistic Regression with $n = 1000$ and $p = 50, 200$

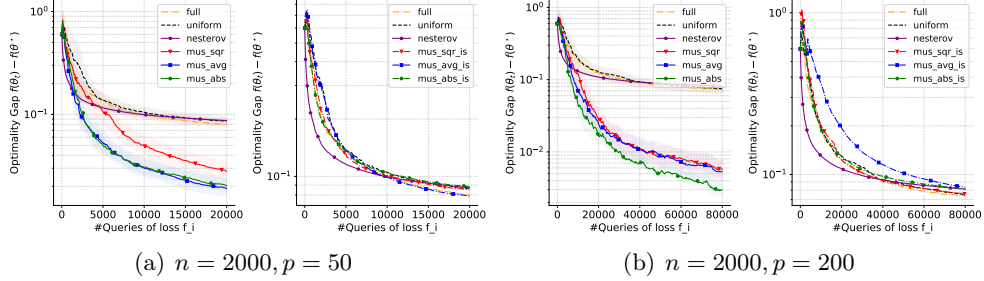


Figure 26: $[f(\theta_t) - f^*]$ for Logistic Regression with $n = 2000$ and $p = 50, 200$

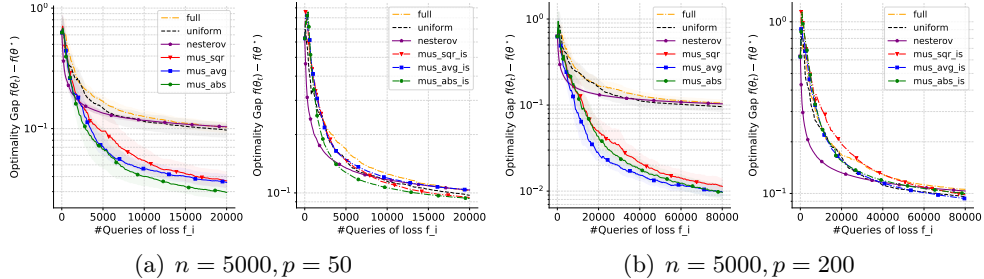


Figure 27: $[f(\theta_t) - f^*]$ for Logistic Regression with $n = 5000$ and $p = 50, 200$

G. Further Experiments with stochastic first order methods

G.1 Comparing learning rates

This section investigates the effect of different learning rates $\gamma_k = \gamma/k$ with $\gamma \in \{0.5; 1; 1.5; 2\}$. It reveals a safe behavior of MUSKETEER as it performs better than the other methods in all configurations with a stronger difference when dealing with small values of γ . We consider the Ridge regression problem with regularization parameter $\mu = 1/n$ and run several experiments in the setting $n = 5,000$ samples and dimension $p \in \{20; 100; 200\}$. We endow the data matrix X with a block structure. The columns are drawn as $X[:, k] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, p \rrbracket$. The parameter α of block structure is $\alpha = 8$. The gradient estimate g is computed using mini-batches of size 4. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs for $N = 100$ iterations with normalized passes over coordinates.




Figure 28: $[f(\theta_t) - f^*]$ for Ridge Regression with $p = 20$ and $\gamma \in \{0.5; 1; 1.5; 2\}$

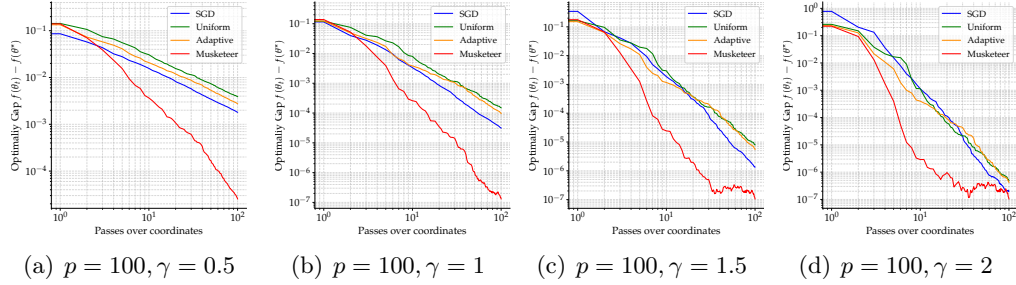


Figure 29: $[f(\theta_t) - f^*]$ for Ridge Regression with $p = 100$ and $\gamma \in \{0.5; 1; 1.5; 2\}$

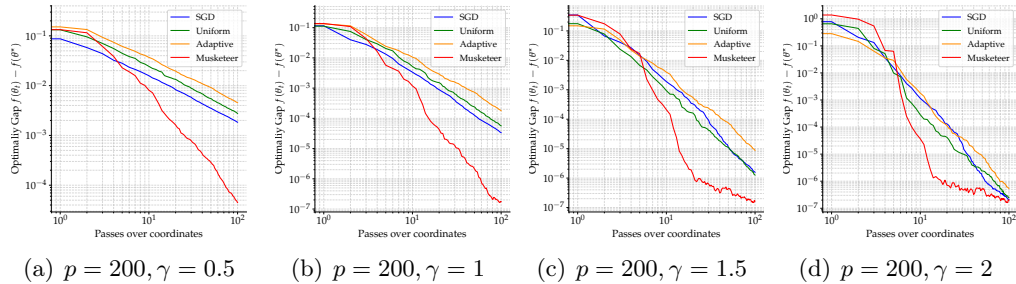


Figure 30: $[f(\theta_t) - f^*]$ for Ridge Regression with $p = 200$ and $\gamma \in \{0.5; 1; 1.5; 2\}$

G.2 Ridge Regression with different settings of (n, p)

We consider the Ridge regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB + 1 : kB + B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 0, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 5$ for the Ridge regression. The parameter α represents the block structure and is set to $\alpha = 10$. The data sampling process ξ of gradient estimate g is computed using mini-batches of size 8. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs for $N = 1000$ iterations with normalized passes over coordinates. The learning rates is the same for all methods, fixed to $\gamma_k = 1/k$. The different settings are: number of samples $n \in \{1,000; 2,000; 5,000\}$ and dimension $p \in \{20; 50; 100; 200\}$.

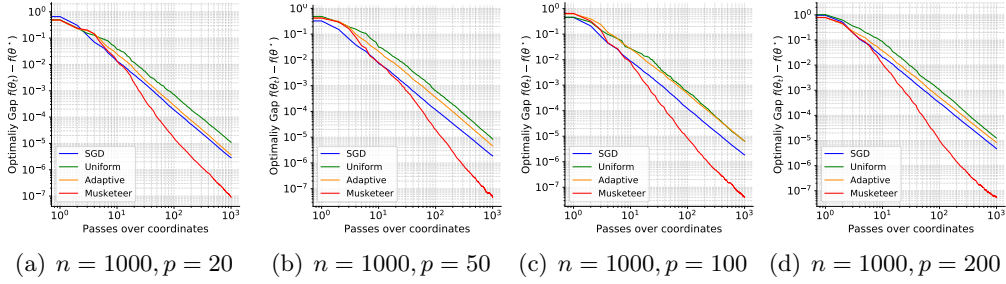


Figure 31: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 1000$ and $p = 20, 50, 100, 200$

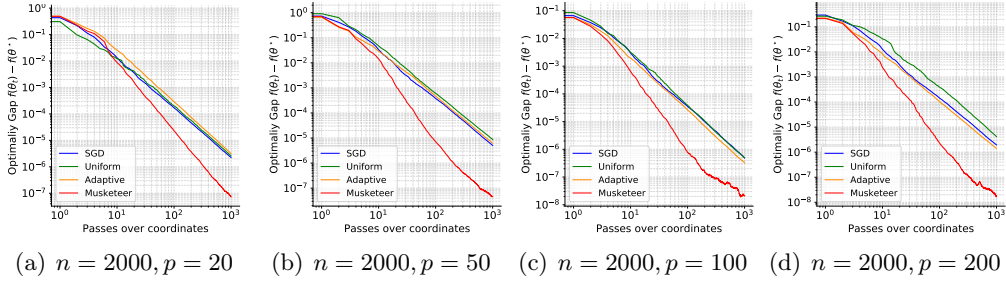


Figure 32: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 2000$ and $p = 20, 50, 100, 200$

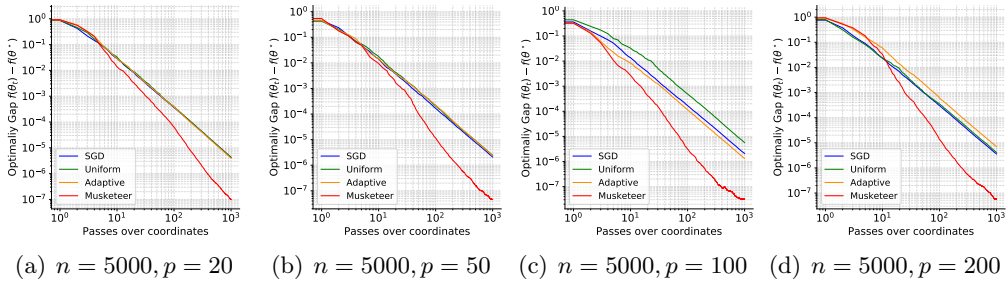


Figure 33: $[f(\theta_t) - f^*]$ for Ridge Regression with $n = 5000$ and $p = 20, 50, 100, 200$

G.3 Logistic Regression with different settings of (n, p)

We consider the ℓ_2 -Logistic regression problem with the classical regularization parameter value $\mu = 1/n$ and run several experiments in various settings of (n, p) . We endow the data matrix X with a block structure. The columns are drawn as $X[:, kB+1 : kB+B] \sim \mathcal{N}(0, \sigma_k^2 I_n)$ with $\sigma_k^2 = k^{-\alpha}$ for all $k \in \llbracket 1, (p/B) - 1 \rrbracket$. The parameter B is the block-size and is set to $B = 2$ for the Logistic regression. The parameter α represents the block structure and is set to $\alpha = 5$. The data sampling process ξ of gradient estimate g is computed using mini-batches of size 32. The different Figures below present the evolution of the optimality gap $t \mapsto [f(\theta_t) - f^*]$ averaged over 20 independent runs for $N = 1000$ iterations with normalized passes over coordinates. The learning rates is the same for all methods, fixed to $\gamma_k = 1/k$. The different settings are: number of samples $n \in \{1,000; 2,000; 5,000\}$ and dimension $p \in \{20; 50; 100; 200\}$.

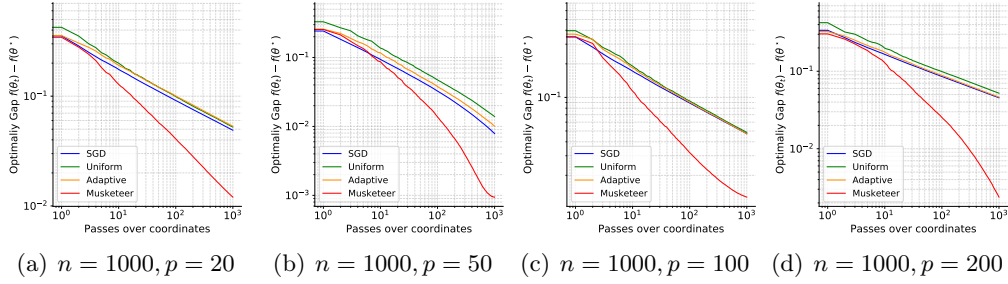


Figure 34: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 1000$ and $p = 20, 50, 100, 200$

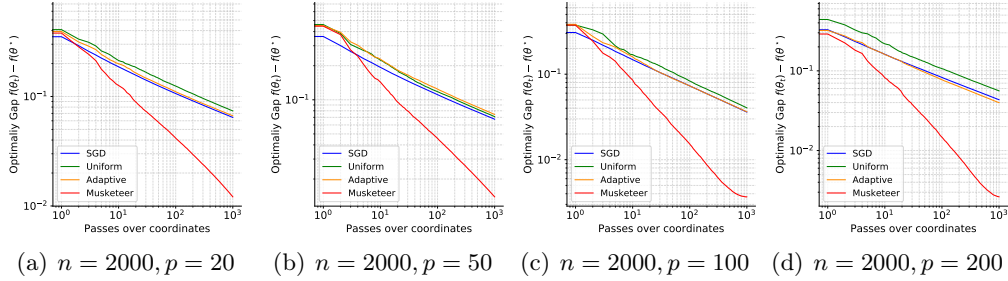


Figure 35: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 2000$ and $p = 20, 50, 100, 200$

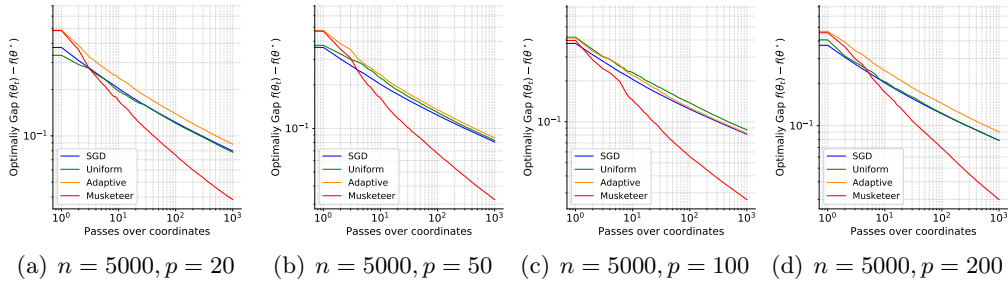


Figure 36: $[f(\theta_t) - f^*]$ for logistic Regression with $n = 5000$ and $p = 20, 50, 100, 200$

References

- Alekh Agarwal, Peter L Bartlett, Pradeep Ravikumar, and Martin J Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.
- Ahmad Ajalloeian and Sebastian U Stich. Analysis of sgd with biased gradient estimators. *arXiv preprint arXiv:2008.00051*, 2020.
- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, pages 1709–1720, 2017.
- Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cédric Renggli. The convergence of sparsified gradient methods. In *Advances in Neural Information Processing Systems*, pages 5973–5983, 2018.
- Zeyuan Allen-Zhu, Zheng Qu, Peter Richtárik, and Yang Yuan. Even faster accelerated coordinate descent using non-uniform sampling. In *International Conference on Machine Learning*, pages 1110–1119, 2016.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002a.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002b.
- Amir Beck and Luba Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM journal on Optimization*, 23(4):2037–2060, 2013.
- Bernard Bercu, Bernard Delyon, and Emmanuel Rio. *Concentration inequalities for sums and martingales*. Springer, 2015.
- Dimitri P Bertsekas and John N Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Claire Boyer and Antoine Godichon-Baggioni. On the asymptotic rate of convergence of stochastic newton algorithms and their weighted averaged versions. *arXiv preprint arXiv:2011.09706*, 2020.
- Dominik Csiba, Zheng Qu, and Peter Richtárik. Stochastic dual coordinate ascent with adaptive probabilities. In *International Conference on Machine Learning*, pages 674–683, 2015.
- Bernard Delyon and François Portier. Asymptotic optimality of adaptive importance sampling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3138–3148. Curran Associates Inc., 2018.

- Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- John C Duchi, Peter L Bartlett, and Martin J Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- Olivier Fercoq and Peter Richtárik. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- Olivier Fercoq, Zheng Qu, Peter Richtárik, and Martin Takáč. Fast distributed coordinate descent for non-strongly convex losses. In *2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2014.
- Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, page 385–394, USA, 2005. Society for Industrial and Applied Mathematics. ISBN 0898715857.
- Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.
- Sébastien Gadat, Fabien Panloup, Sofiane Saadane, et al. Stochastic heavy ball. *Electronic Journal of Statistics*, 12(1):461–529, 2018.
- Nidham Gazagnadou, Robert Gower, and Joseph Salmon. Optimal mini-batch and step sizes for saga. In *International conference on machine learning*, pages 2142–2150. PMLR, 2019.
- Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. Stochastic optimization for large-scale optimal transport. *Advances in neural information processing systems*, 29, 2016.
- Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Tobias Glasmachers and Urun Dogan. Accelerated coordinate descent with adaptive coordinate frequencies. In *Asian Conference on Machine Learning*, pages 72–86, 2013.
- Robert M Gower, Peter Richtárik, and Francis Bach. Stochastic quasi-gradient methods: Variance reduction via jacobian sketching. *Mathematical Programming*, 188(1):135–192, 2021.
- Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. In *International Conference on Machine Learning*, pages 5200–5209. PMLR, 2019.
- P. Hall and C.C. Heyde. *Martingale Limit Theory and Its Application*. Probability and mathematical statistics. Academic Press, 1980. ISBN 9781483240244. URL <https://books.google.fr/books?id=wdLajgEACAAJ>.

- Josiah Hanna, Scott Niekum, and Peter Stone. Importance sampling policy evaluation with an estimated behavior policy. In *International Conference on Machine Learning*, pages 2605–2613. PMLR, 2019.
- Jack Kiefer, Jacob Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yin Tat Lee and Aaron Sidford. Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 147–156. IEEE, 2013.
- Rémi Leluc and François Portier. Towards asymptotic optimality with conditioned stochastic gradient descent. *arXiv preprint arXiv:2006.02745*, 2020.
- Xiangru Lian, Huan Zhang, Cho-Jui Hsieh, Yijun Huang, and Ji Liu. A comprehensive linear speedup analysis for asynchronous stochastic parallel optimization from zeroth-order to first-order. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Sijia Liu, Pin-Yu Chen, Bhavya Kailkhura, Gaoyuan Zhang, Alfred O Hero III, and Pramod K Varshney. A primer on zeroth-order optimization in signal processing and machine learning: Principals, recent advances, and applications. *IEEE Signal Processing Magazine*, 37(5): 43–54, 2020.
- Ilya Loshchilov, Marc Schoenauer, and Michele Sebag. Adaptive coordinate descent. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 885–892, 2011.
- Zhaosong Lu and Lin Xiao. On the complexity analysis of randomized block-coordinate descent methods. *Mathematical Programming*, 152(1-2):615–642, 2015.
- Eric Moulines and Francis R Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459, 2011.
- Hongseok Namkoong, Aman Sinha, Steve Yadowlowsky, and John C Duchi. Adaptive sampling probabilities for non-smooth optimization. In *International Conference on Machine Learning*, pages 2574–2583, 2017.
- I Necoara, Y Nesterov, and F Glineur. A random coordinate descent method on large-scale optimization problems with linear constraints. Technical report, Technical Report, 2014.
- Deanna Needell, Rachel Ward, and Nati Srebro. Stochastic gradient descent, weighted sampling, and the randomized kaczmarz algorithm. In *Advances in neural information processing systems*, pages 1017–1025, 2014.

- 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.
- Arkadi Semenovich Nemirovski and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.
- Lam Nguyen, Phuong Ha Nguyen, Marten Dijk, Peter Richtárik, Katya Scheinberg, and Martin Takáč. Sgd and hogwild! convergence without the bounded gradients assumption. In *International Conference on Machine Learning*, pages 3750–3758. PMLR, 2018.
- Julie Nutini, Mark Schmidt, Issam Laradji, Michael Friedlander, and Hoyt Koepke. Coordinate descent converges faster with the gauss-southwell rule than random selection. In *International Conference on Machine Learning*, pages 1632–1641, 2015.
- Guillaume Papa, Pascal Bianchi, and Stéphan Cléménçon. Adaptive sampling for incremental optimization using stochastic gradient descent. In *International Conference on Algorithmic Learning Theory*, pages 317–331. Springer, 2015.
- Kumar Kshitij Patel and Aymeric Dieuleveut. Communication trade-offs for local-sgd with large step size. *Advances In Neural Information Processing Systems 32 (Nips 2019)*, 32 (CONF), 2019.
- Dmytro Perekrestenko, Volkan Cevher, and Martin Jaggi. Faster coordinate descent via adaptive importance sampling. In *Artificial Intelligence and Statistics*, pages 869–877. PMLR, 2017.
- Boris Teodorovich Polyak. Gradient methods for minimizing functionals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 3(4):643–653, 1963.
- Zheng Qu and Peter Richtárik. Coordinate descent with arbitrary sampling i: Algorithms and complexity. *Optimization Methods and Software*, 31(5):829–857, 2016.
- Zheng Qu, Peter Richtárik, and Tong Zhang. Quartz: Randomized dual coordinate ascent with arbitrary sampling. In *Advances in neural information processing systems*, pages 865–873, 2015.
- Peter Richtárik and Martin Takáč. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming*, 144 (1-2):1–38, 2014.
- Peter Richtárik and Martin Takáč. On optimal probabilities in stochastic coordinate descent methods. *Optimization Letters*, 10(6):1233–1243, 2016a.

- Peter Richtárik and Martin Takáč. Parallel coordinate descent methods for big data optimization. *Mathematical Programming*, 156(1-2):433–484, 2016b.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- Herbert Robbins and David Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pages 233–257. Elsevier, 1971.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.
- 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.
- Ohad Shamir. An optimal algorithm for bandit and zero-order convex optimization with two-point feedback. *The Journal of Machine Learning Research*, 18(1):1703–1713, 2017.
- Yining Wang, Simon Du, Sivaraman Balakrishnan, and Aarti Singh. Stochastic zeroth-order optimization in high dimensions. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1356–1365. PMLR, 09–11 Apr 2018. URL <https://proceedings.mlr.press/v84/wang18e.html>.
- Jianqiao Wangni, Jialei Wang, Ji Liu, and Tong Zhang. Gradient sparsification for communication-efficient distributed optimization. In *Advances in Neural Information Processing Systems*, pages 1299–1309, 2018.
- Tong Tong Wu, Kenneth Lange, et al. Coordinate descent algorithms for lasso penalized regression. *Annals of Applied Statistics*, 2(1):224–244, 2008.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.