



HAL
open science

Stochastic Mirror Descent for Large-Scale Sparse Recovery

Yannis Bekri, Sasila Ilandarideva, Anatoli B. Juditsky, Vianney Perchet

► **To cite this version:**

Yannis Bekri, Sasila Ilandarideva, Anatoli B. Juditsky, Vianney Perchet. Stochastic Mirror Descent for Large-Scale Sparse Recovery. AISTATS - 2023 - 26th International Conference on Artificial Intelligence and Statistics, Apr 2023, Valencia, Spain. pp.1-27. hal-03826176

HAL Id: hal-03826176

<https://hal.science/hal-03826176v1>

Submitted on 24 Oct 2022

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.

Stochastic Mirror Descent for Large-Scale Sparse Recovery

Yannis Bekri¹

yannis.bekri@univ-grenoble-alpes.fr

Sasila Ilandarideva¹

sasila.ilandarideva@univ-grenoble-alpes.fr

Anatoli Juditsky¹

anatoli.juditsky@univ-grenoble-alpes.fr

Vianney Perchet²

vianney.perchet@normalsup.org

¹LJK, Université Grenoble Alpes, Grenoble

²CREST, ENSAE Paris and CRITEO AI Lab, Paris

Abstract

In this paper we discuss an application of Stochastic Approximation to statistical estimation of high-dimensional sparse parameters. The proposed solution reduces to resolving a penalized stochastic optimization problem on each stage of a multistage algorithm; each problem being solved to a prescribed accuracy by the non-Euclidean Composite Stochastic Mirror Descent (CSMD) algorithm. Assuming that the problem objective is smooth and quadratically minorated and stochastic perturbations are sub-Gaussian, our analysis prescribes the method parameters which ensure fast convergence of the estimation error (the radius of a confidence ball of a given norm around the approximate solution). This convergence is linear during the first “preliminary” phase of the routine and is sublinear during the second “asymptotic” phase. We consider an application of the proposed approach to sparse Generalized Linear Regression problem. In this setting, we show that the proposed algorithm attains the optimal convergence of the estimation error under weak assumptions on the regressor distribution. We also present a numerical study illustrating the performance of the algorithm on high-dimensional simulation data.

1 Introduction

Our original motivation is the well known problem of (generalized) linear high-dimensional regression with random design. Formally, consider a dataset of N points $(\phi_i, \eta_i), i \in \{1, \dots, N\}$, where $\phi_i \in \mathbf{R}^n$ are (random) features and $\eta_i \in \mathbf{R}$ are observations, linked by the following equation

$$\eta_i = \mathbf{r}(\phi_i^T x_*) + \sigma \xi_i, \quad i \in [N] := \{1, \dots, N\} \quad (1)$$

where $\xi_i \in \mathbf{R}$ are i.i.d. observation noises. The standard objective is to recover the unknown parameter $x_* \in \mathbf{R}^n$ of the Generalized Linear Regression (1) – which is assumed to belong to a given convex closed set X and to be *s-sparse*, i.e., to have at most $s \ll n$ non-vanishing entries from the data-set.

As mentioned before, we consider random design, where ϕ_i are i.i.d. random variables, so that the estimation problem of x_* can be recast as the following generic Stochastic Optimization problem:

$$g_* = \min_{x \in X} g(x), \quad \text{where} \quad g(x) = \mathbf{E}\{G(x, (\phi, \eta))\}, \quad G(x, (\phi, \eta)) = \mathfrak{s}(\phi^T x) - \phi^T x \eta, \quad (2)$$

with $\mathfrak{s}(\cdot)$ any primitive of $\mathfrak{r}(\cdot)$, i.e., $\mathfrak{r}(t) = \mathfrak{s}'(t)$. The equivalence between the original and the stochastic optimization problems comes from the fact that x_* is a critical point of $g(\cdot)$, i.e., $\nabla g(x_*) = 0$ since, under mild assumptions, $\nabla g(x) = \mathbf{E}\{\mathfrak{r}(\phi^T x) - \mathfrak{r}(\phi^T x_*)\}$. Hence, as soon as g as a unique minimizer (say, g is strongly convex over X), solutions of both problems are identical.

As a consequence, we shall focus on the generic problem (2), that has already been widely tackled. For instance, when given an observation sample (ϕ_i, η_i) , $i \in [N]$, one may build a Sample Average Approximation (SAA) of the objective $g(x)$

$$\widehat{g}_N(x) = \frac{1}{N} \sum_{i=1}^N G(x, (\phi_i, \eta_i)) = \frac{1}{N} \sum_{i=1}^N [\mathfrak{s}(\phi_i^T x) - \phi_i^T x \eta_i] \quad (3)$$

and then solve the resulting problem of minimizing $\widehat{g}_N(x)$ over sparse x 's. The celebrated ℓ_1 -norm minimization approach allows to reduce this problem to convex optimization. We will provide a new algorithm adapted to this high-dimensional case, and instantiating it to the original problem 1.

Existing approaches and related works. Sparse recovery by Lasso and Dantzig Selector has been extensively studied [11, 8, 5, 46, 10, 9]. It computes a solution \widehat{x}_N to the ℓ_1 -penalized problem $\min_x \widehat{g}_N(x) + \lambda \|x\|_1$ where $\lambda \geq 0$ is the algorithm parameter [35]. This delivers “good solutions”, with high probability for sparsity level s as large as $O\left(\frac{N\kappa_\Sigma}{\ln n}\right)$, as soon as the random regressors (the ϕ_i) are drawn independently from a normal distribution with a covariance matrix Σ such that $\kappa_\Sigma I \preceq \Sigma \preceq \rho \kappa_\Sigma I^1$, for some $\kappa_\Sigma > 0, \rho \geq 1$. However, computing this solution may be challenging in a very high-dimensional setting: even popular iterative algorithms, like coordinate descent, loops over a large number of variables. To mitigate this, randomized algorithms [3, 22], screening rules and working sets [19, 30, 34] may be used to diminish the size of the optimization problem at hand, while iterative thresholding [1, 7, 20, 16, 33] is a “direct” approach to enhance sparsity of the solution.

Another approach relies on Stochastic Approximation (SA). As $\nabla G(x, (\phi_i, \eta_i)) = \phi_i(\mathfrak{r}(\phi_i^T x) - \eta_i)$ is an unbiased estimate of $\nabla g(x)$, iterative Stochastic Gradient Descent (SGD) algorithm may be used to build approximate solutions. Unfortunately, unless regressors ϕ are sparse or possess a special structure, standard SA leads to accuracy bounds for sparse recovery proportional to the dimension n which are essentially useless in the high-dimensional setting. This motivates non-Euclidean SA procedures, such as Stochastic Mirror Descent (SMD) [37], its application to sparse recovery enjoys almost dimension free convergence and it has been well studied in the literature. For instance, under bounded regressors and with sub-Gaussian noise, SMD reaches “slow rate” of sparse recovery of the type $g(\widehat{x}_N) - g_* = O\left(\sigma \sqrt{s \ln(n)/N}\right)$ where \widehat{x}_N is the approximate solution after N iterations [44, 45]. Multistage routines may be used to improve the error estimates of SA under strong or uniform convexity assumptions [27, 29, 18]. However, they do not always hold, as in sparse Generalized Linear Regression, where they are replaced by Restricted Strong Convexity conditions. Then multistage procedures [2, 17] based on standard SMD algorithms [24, 38] control the ℓ_2 -error $\|\widehat{x}_N - x_*\|_2$ at the rate $O\left(\frac{\sigma}{\kappa_\Sigma} \sqrt{\frac{s \ln n}{N}}\right)$ with high probability. This is the best “asymptotic” rate attainable when solving (2). However, those algorithms have two major limitations. They both need a number of iterations to reach a given accuracy proportional to the initial error $R = \|x_* - x_0\|_1$ and the sparsity level s must be of order $O\left(\kappa_\Sigma \sqrt{\frac{N}{\ln n}}\right)$ for the sparse linear regression. These limits may be seen as a consequence of dealing with *non-smooth* objective $g(x)$. Although it slightly restricts the scope of corresponding algorithms, we shall consider smooth objectives and algorithm for minimizing composite objectives (cf. [25, 32, 39]) to mitigate the aforementioned drawbacks of the multistage algorithms from [2, 17].

Principal contributions. We provide a refined analysis of *Composite Stochastic Mirror Descent (CSMD)* algorithms for computing sparse solutions to Stochastic Optimization problem leveraging smoothness of the objective. This leads to a new “aggressive” choice of parameters in a multistage algorithm with significantly improved performances compared to those in [2]. We summarize below some properties of the proposed procedure for problem (2).

¹We use $A \preceq B$ for two symmetric matrices A and B if $B - A \succeq 0$, i.e. $B - A$ is positive semidefinite.

Each stage of the algorithm is a specific CSMD recursion; They fall into two phases. During the first (preliminary) phase, the estimation error decreases linearly with the exponent proportional to $\frac{\kappa_\Sigma}{s \ln n}$. When it reaches the value $O\left(\frac{\sigma s}{\sqrt{\kappa_\Sigma}}\right)$, the second (asymptotic) phase begins, and its stages contain exponentially increasing number of iterations per stage, hence the estimation error decreases as $O\left(\frac{\sigma s}{\kappa_\Sigma} \sqrt{\frac{\ln n}{N}}\right)$ where N is the total iteration count.

Organization and notation The remaining of the paper is organized as follows. In Section 2, the general problem is set, and the multistage optimization routine and the study of its basic properties are presented. Then, in Section 3, we discuss the properties of the method and conditions under which it leads to “small error” solutions to sparse GLR estimation problems. Finally, a small simulation study illustrating numerical performance of the proposed routines in high-dimensional GLR estimation problem is presented in Section 3.3.

In the following, E is a Euclidean space and $\|\cdot\|$ is a norm on E ; we denote $\|\cdot\|_*$ the conjugate norm (i.e., $\|x\|_* = \sup_{\|y\| \leq 1} \langle y, x \rangle$). Given a positive semidefinite matrix $\Sigma \in \mathbf{S}_n$, for $x \in \mathbf{R}^n$ we denote $\|x\|_\Sigma = \sqrt{x^T \Sigma x}$ and for any matrix Q , we denote $\|Q\|_\infty = \max_{ij} |[Q]_{ij}|$. We use a generic notation c and C for absolute constants; a shortcut notation $a \lesssim b$ ($a \gtrsim b$) means that the ratio a/b (ratio b/a) is bounded by an absolute constant; the symbols \vee, \wedge and the notation $(\cdot)_+$ respectively refer to “maximum between”, “minimum between” and “positive part”.

2 Multistage Stochastic Mirror Descent for Sparse Stochastic Optimization

This section is dedicated to the formulation of the generic stochastic optimization problem, the description and the analysis of the generic algorithm.

2.1 Problem statement

Let X be a convex closed subset of an Euclidean space E and (Ω, P) a probability space. We consider a mapping $G : X \times \Omega \rightarrow \mathbf{R}$ such that, for all $\omega \in \Omega$, $G(\cdot, \omega)$ is convex on X and smooth, meaning that $\nabla G(\cdot, \omega)$ is Lipschitz continuous on X with a.s. bounded Lipschitz constant,

$$\forall x, x' \in X, \quad \|\nabla G(x, \omega) - \nabla G(x', \omega)\| \leq \mathcal{L}(\omega) \|x - x'\|, \quad \mathcal{L}(\omega) \leq \nu \quad a.s.. \quad (4)$$

We define $g(x) := \mathbf{E}\{G(x, \omega)\}$, where $\mathbf{E}\{\cdot\}$ stands for the expectation with respect to ω , drawn from P . We shall assume that the mapping $g(\cdot)$ is finite, convex and differentiable on X and we aim at solving the following stochastic optimization problem

$$\min_{x \in X} [g(x) = \mathbf{E}\{G(x, \omega)\}], \quad (5)$$

assuming it admits an s -sparse optimal solution x_* for some sparsity structure.

To solve this problem, stochastic oracle can be queried: when given at input a point $x \in X$, generates an $\omega \in \Omega$ from P and outputs $G(x, \omega)$ and $\nabla G(x, \omega) := \nabla_x G(x, \omega)$ (with a slight abuse of notations). We assume that the oracle is *unbiased*, i.e.,

$$\mathbf{E}\{\nabla G(x, \omega)\} = \nabla g(x), \quad \forall x \in X.$$

To streamline presentation, we assume, as it is often the case in applications of stochastic optimization problem (5), that x_* is unconditional, i.e., $\nabla g(x_*) = 0$. or stated otherwise $\mathbf{E}\{\nabla G(x_*, \omega)\} = 0$; we also suppose the sub-Gaussianity of $\nabla G(x_*, \omega)$, namely that, for some $\sigma_* < \infty$

$$\mathbf{E}\left\{\exp\left(\|\nabla G(x_*, \omega)\|_*^2 / \sigma_*^2\right)\right\} \leq \exp(1). \quad (6)$$

2.2 Composite Stochastic Mirror Descent algorithm

As mentioned in the introduction, (stochastic) optimization over the set of sparse solutions can be done through "composite" techniques. We take a similar approach here, by transforming the generic problem 5 into the following *composite Stochastic Optimization problem*, adapted to some norm $\|\cdot\|$, and parameterized by $\kappa \geq 0$,

$$\min_{x \in X} [F_\kappa(x) := \frac{1}{2}g(x) + \kappa\|x\| = \frac{1}{2}\mathbf{E}\{G(x, \omega)\} + \kappa\|x\|]. \quad (7)$$

The purpose of this section is to derive a new (proximal) algorithm. We first provide necessary backgrounds and notations.

Proximal setup, Bregman divergences and Proximal mapping. Let B be the unit ball of the norm $\|\cdot\|$ and $\theta : B \rightarrow \mathbf{R}$ be a *distance-generating function (d.-g.f.)* of B , i.e., a continuously differentiable convex function which is strongly convex with respect to the norm $\|\cdot\|$,

$$\langle \nabla\theta(x) - \nabla\theta(x'), x - x' \rangle \geq \|x - x'\|^2, \quad \forall x, x' \in X.$$

We assume w.l.o.g. that $\theta(x) \geq \theta(0) = 0$ and denote $\Theta = \max_{\|z\| \leq 1} \theta(z)$.

We now introduce a local and renormalized version of the d.-g.f. θ .

Definition 2.1 For any $x_0 \in X$, let $X_R(x_0) := \{z \in X : \|z - x_0\| \leq R\}$ be the ball of radius R around x_0 . It is equipped with the d.-g.f. $\vartheta_{x_0}^R(z) := R^2\theta((z - x_0)/R)$.

Note that $\vartheta_{x_0}^R(z)$ is strongly convex on $X_R(x_0)$ with modulus 1, $\vartheta_{x_0}^R(x_0) = 0$, and $\vartheta_{x_0}^R(z) \leq \Theta R^2$.

Definition 2.2 Given $x_0 \in X$ and $R > 0$, the Bregman divergence V associated to ϑ is defined by

$$V_{x_0}(x, z) = \vartheta_{x_0}^R(z) - \vartheta_{x_0}^R(x) - \langle \nabla\vartheta_{x_0}^R(x), z - x \rangle, \quad x, z \in X.$$

We can now define *composite proximal mapping* on $X_R(x_0)$ [39, 40] with respect to some convex and continuous mapping $h : X \rightarrow \mathbf{R}$.

Definition 2.3 The composite proximal mapping with respect to h and x is defined by

$$\begin{aligned} \text{Prox}_{h, x_0}(\zeta, x) &:= \arg \min_{z \in X_R(x_0)} \{ \langle \zeta, z \rangle + h(z) + V_{x_0}(x, z) \} \\ &= \arg \min_{z \in X_R(x_0)} \{ \langle \zeta - \nabla\vartheta_{x_0}^R(x), z \rangle + h(z) + \vartheta_{x_0}^R(z) \} \end{aligned} \quad (8)$$

If (8) can be efficiently solved to high accuracy and Θ is "not too large" (we refer to [27, 36, 40]); those setups will be called "prox-friendly". We now introduce the main building block of our algorithm, the Composite Stochastic Mirror Descent.

Composite Stochastic Mirror Descent algorithm. Given a sequence of positive *step sizes* $\gamma_i > 0$, the *Composite Stochastic Mirror Descent* (CSMD) is defined by the following recursion

$$x_i = \text{Prox}_{\gamma_i h, x_0}(\gamma_{i-1} \nabla G(x_{i-1}, \omega_i), x_{i-1}), \quad x_0 \in X. \quad (9)$$

After m steps of CSMD, the final output is \hat{x}_m (approximate solution) defined by

$$\hat{x}_m = \frac{\sum_{i=0}^{m-1} \gamma_i x_i}{\sum_{i=0}^{m-1} \gamma_i} \quad (10)$$

For any integer $L \in \mathbf{N}$, we can also define the L -minibatch CSMD. Let $\omega_i^{(L)} = [\omega_i^1, \dots, \omega_i^L]$ be i.i.d. realizations of ω_i . The associated (average) stochastic gradient is then simply defined as

$$H(x_{i-1}, \omega_i^{(L)}) = \frac{1}{L} \sum_{\ell=1}^L \nabla G(x_{i-1}, \omega_i^\ell),$$

which yields the following recursion for the L -minibatch CSMD recursion:

$$x_i^{(L)} = \text{Prox}_{\gamma_i h, x_0} \left(\gamma_{i-1} H \left(x_{i-1}, \omega_i^{(L)} \right), x_{i-1}^{(L)} \right), \quad x_0 \in X, \quad (11)$$

with its approximate solution $\hat{x}_m^{(L)} = \sum_{i=0}^{m-1} \gamma_i x_i^{(L)} / \sum_{i=0}^{m-1} \gamma_i$ after m iterations.

From now on, we set $h(x) = \kappa \|x\|$.

Proposition 2.1 *If step-sizes are constant, i.e., $\gamma_i \equiv \gamma \leq (4\nu)^{-1}$, $i = 0, 1, \dots$, and the initial point $x_0 \in X$ such that $x_* \in X_R(x_0)$ then for any $t \gtrsim \sqrt{1 + \ln m}$, with probability at least $1 - 4e^{-t}$*

$$F_\kappa(\hat{x}_m) - F_\kappa(x_*) \lesssim m^{-1} [\gamma^{-1} R^2 (\Theta + t) + \kappa R + \gamma \sigma_*^2 (m + t)], \quad (12)$$

and the approximate solution $\hat{x}_m^{(L)}$ of the L -minibatch CSMD satisfies

$$F_\kappa(\hat{x}_m^{(L)}) - F_\kappa(x_*) \lesssim m^{-1} [\gamma^{-1} R^2 (\Theta + t) + \kappa R + \gamma \sigma_*^2 \Theta L^{-1} (m + t)]. \quad (13)$$

For the sake of clarity and conciseness, we denote $\text{CSMD}(x_0, \gamma, \kappa, R, m, L)$ the approximate solution $\hat{x}_m^{(L)}$ computed after m iterations of L -minibatch CSMD algorithm with initial point x_0 , step-size γ , and radius R using recursion (11).

2.3 Main contribution: a multistage adaptive algorithm

Our approach to find sparse solution to the original stochastic optimization problem (7) consists in solving a sequence of auxiliary composite problems (7), with their sequence of parameters (κ, x_0, R) defined recursively. For the latter, we need to infer the quality of approximate solution to (5). To this end, we introduce the following *Reduced Strong Convexity* (RSC) assumption, satisfied in the motivating example (it is discussed in the appendix for the sake of fluency):

Assumption [RSC] There exist some $\delta > 0$ and $\rho < \infty$ such that for any feasible solution $\hat{x} \in X$ to the composite problem (7) satisfying, with probability at least $1 - \varepsilon$,

$$F_\kappa(\hat{x}) - F_\kappa(x_*) \leq \nu,$$

it holds, with probability at least $1 - \varepsilon$, that

$$\|\hat{x} - x_*\| \leq \delta [\rho s \kappa + \nu \kappa^{-1}]. \quad (14)$$

Given the different problem parameters $s, \nu, \delta, \rho, \kappa, R$ and some initial point $x_0 \in X$ such that $x_* \in X_R(x_0)$ Algorithm 1 works in stages. Each stage represents a run of CSMD algorithm with properly set penalty parameter κ . More precisely, at stage $k + 1$, given the approximate solution \hat{x}_m^k of stage k , a new instance of CSMD is initialized on $X_{R_{k+1}}(x_0^{k+1})$ with $x_0^{k+1} = \hat{x}_m^k$ and $R_{k+1} = R_k/2$.

Furthermore, those stages are divided into two phases which we refer to as *preliminary* and *asymptotic*:

Preliminary phase: During this phase, the step-sizes γ and the number of CSMD iterations per stage are fixed; the error of approximate solutions converges linearly with the total number of calls to stochastic oracle. This phase terminates when the error of approximate solution becomes independent of the initial error of the algorithm; then the asymptotic phase begins.

Asymptotic phase: In this phase, the step-size decreases and the length of the stage increases linearly; the solution converges sublinearly, with the “standard” rate $O(N^{-1/2})$ where N is the total number of oracle calls. When expensive proximal computation (8) results in high numerical cost of the iterative algorithm, minibatches are used to keep the number of iterations per stage fixed.

Algorithm 1 CSMD-SR

Initialization : Initial point $x_0 \in X$, step-size $\gamma = (4\nu)^{-1}$, initial radius R_0 , confidence level t , total budget N .

Set $m_0 \asymp s\rho\nu\delta^2(\Theta + t)$, $\bar{K}_1 \asymp \ln\left(\frac{R_0^2\nu}{\delta^2\rho\sigma_*^2s}\right) \wedge \frac{N}{m_0}$, $L = 1$

if $R_0 \gtrsim \sigma_*\delta\sqrt{\frac{\rho s}{\nu}}$ **continue** with preliminary stage,

else proceed directly to asymptotic phase

end

for stage $k = 1, \dots, \bar{K}_1$ **do**

▷ Preliminary Phase

Set $\kappa_k \asymp R_k(\delta\rho s)^{-1}$

Compute approximate solution $\hat{x}_{m_0}^k = \text{CSMD}(x_0, \gamma, \kappa_k, R_k, m_0, L)$ at stage k

Reset the prox-center $x_0 = \hat{x}_{m_0}^k$

Set $R_k = R_{k-1}/2$

end for

Set $\hat{x}_N = \hat{x}_{m_0}^{\bar{K}_1}$, $B = N - m_0\bar{K}_1$, $m_1 \asymp m_0$

if $m_1 > B$ **output** : \hat{x}_N and **return**; **endif**

▷ Asymptotic Phase

Set $r_0 = R_{\bar{K}_1}$

Set $k = 1$

while $m_k \leq B$ **do**

Set $\kappa_k \asymp 2^{-k}\sigma_*(\rho\nu s)^{-1/2}$, $\gamma_k \asymp 4^{-k}\nu^{-1}$

Compute approximate solution $\hat{x}_{m_k}^k = \text{CSMD}(x_0, \gamma_k, \kappa_k, r_k, m_k, L)$ at stage k

Reset the prox-center $x_0 = \hat{x}_{m_k}^k$

Set $B = B - m_k$, $k = k + 1$, $r_k = r_{k-1}/2$, $m_k \asymp 4^k m_0$

end while

output : $\hat{x}_N = \hat{x}^k$

In the algorithm description, \bar{K}_1 and $\bar{K}_2 \asymp 1 + \log(\frac{N}{m_0})$ stand for the respective maximal number of stages of the two phases of the method, here, $m_0 \asymp s\rho\nu\delta^2(\Theta + t)$ is the length of stages of the first (preliminary) phase. The pseudo-code for the variant of the asymptotic phase with minibatches is given in Algorithm 2.

The following theorem states the main result of this paper, an upper bound on the precision of the estimator computed by our multistage method.

Theorem 2.1 *Assume that the total sample budget satisfies $N \geq m_0$, so that at least one stage of the preliminary phase of Algorithm 1 is completed, then for $t \gtrsim \sqrt{\ln N}$ the approximate solution \hat{x}_N of Algorithm 1 satisfies, with probability at least $1 - C(\bar{K}_1 + \bar{K}_2)e^{-t}$,*

$$\|\hat{x}_N - x_*\| \lesssim R \exp\left\{-\frac{c}{\delta^2\rho\nu} \frac{N}{s(\Theta + t)}\right\} + \delta^2\rho\sigma_*s\sqrt{\frac{\Theta + t}{N}}.$$

The corresponding solution $\hat{x}_N^{(b)}$ of the minibatch Algorithm 2 satisfies with probability $\geq 1 - C(\bar{K}_1 + \tilde{K}_2)e^{-t}$

$$\|\hat{x}_N^{(b)} - x_*\| \lesssim R \exp\left\{-\frac{c}{\delta^2\rho\nu} \frac{N}{s(\Theta + t)}\right\} + \delta^2\rho\sigma_*s\sqrt{\frac{\Theta(\Theta + t)}{N}}.$$

where $\tilde{K}_2 \asymp 1 + \ln(\frac{N}{\Theta m_0})$ is the bound for the number of stages of the asymptotic phase of the minibatch algorithm.

Remark 2.1 *Along with the oracle computation, proximal computation to be implemented at each iteration of the algorithm is an important part of the computational cost of the method. It becomes even more important during the asymptotic phase when number of iterations per stage increases exponentially fast with the stage count, and may result in poor real-time convergence. The interest of minibatch implementation of the second phase of the algorithm is in reducing drastically the number of iterations per asymptotic stage. The price to paid is an extra factor $\sqrt{\Theta}$ that could also theoretically hinder convergence. However, in the problems of interest (sparse*

Algorithm 2 Asymptotic phase of CSMD-SR with minibatch

Input : The approximate solution $\widehat{x}_{m_0}^{\overline{K}_1}$ at the end of the preliminary stage, step-size parameter γ , radius at the end of the preliminary phase $R_{\overline{K}_1}$, initial batch size $\ell_1 \asymp \Theta$

- 1: Set $r_0 = R_{\overline{K}_1}$, $x_0 = \widehat{x}_{m_0}^{\overline{K}_1}$, $B = N - m_0\overline{K}_1$ ▷ Asymptotic Phase
 - 2: $k = 1$
 - 3: **while** $m_0\ell_k \leq B$ **do**
 - 4: $\kappa_k \asymp 2^{-k}\sigma_*(\rho\nu s)^{-1/2}$
 - 5: Compute approximate solution $\widehat{x}_{m_0}^k = \text{CSMD}(x_0, \gamma_k, \kappa_k, r_k, m_0, L = \ell_k)$ at stage k
 - 6: Reset the prox-center $x_0 = \widehat{x}_{m_0}^k$
 - 7: Set $B = B - m_0\ell_k$, $k = k + 1$, $r_k = r_{k-1}/2$, $\ell_k \asymp 4^k\ell_1$
 - 8: **end while**
- output:** $\widehat{x}_N^{(b)} = \widehat{x}_{m_2}^k$
-

and group-sparse recovery, low rank matrix recovery) Θ is logarithmic in problem dimension. Furthermore, in our numerical experiments we did not observe any accuracy degradation when using the minibatch variant of the method.

3 Sparse generalized linear regression by stochastic approximation

3.1 Problem setting

We now consider again the original problem of recovery of a s -sparse signal $x_* \in X \subset \mathbf{R}^n$ from random observations defined by

$$\eta_i = \mathbf{r}(\phi_i^T x_*) + \sigma\xi_i, \quad i = 1, 2, \dots, N, \quad (15)$$

where $\mathbf{r} : \mathbf{R} \rightarrow \mathbf{R}$ is some non-decreasing and continuous “activation function”, and $\phi_i \in \mathbf{R}^n$ and $\xi_i \in \mathbf{R}$ are mutually independent. We assume that ξ_i are sub-Gaussian, i.e., $\mathbf{E}\{e^{\xi_i^2}\} \leq \exp(1)$, while regressors ϕ_i are bounded, i.e., $\|\phi_i\|_\infty \leq \overline{\nu}$. We also denote $\Sigma = \mathbf{E}\{\phi_i\phi_i^T\}$, with $\Sigma \succeq \kappa_\Sigma I$ with some $\kappa_\Sigma > 0$, and $\|\Sigma_j\|_\infty \leq \nu < \infty$.

We will apply the machinery developed in Section 2, with respect to

$$g(x) = \mathbf{E}\{\mathfrak{s}(\phi^T x) - x^T \phi \eta\}$$

where $\mathbf{r}(t) = \nabla \mathfrak{s}(t)$ for some convex and continuously differentiable \mathfrak{s} , applied with the norm $\|\cdot\| = \|\cdot\|_1$ (hence $\|\cdot\|_* = \|\cdot\|_\infty$), from some initial point $x_0 \in X$ such that $\|x_* - x_0\|_1 \leq R$. It remains to prove that the different assumptions of Section 2 are satisfied.

Proposition 3.1 *Assume that \mathbf{r} is \overline{r} -Lipschitz continuous and \underline{r} -strongly monotone (i.e., $|\mathbf{r}(t) - \mathbf{r}(t')| \geq \underline{r}|t - t'|$ which implies that \mathfrak{s} is \underline{r} -strongly convex) then*

1. [Smoothness] $G(\cdot, \omega)$ is $\mathcal{L}(\omega)$ -smooth with $\mathcal{L}(\omega) \leq \overline{r}\overline{\nu}^2$.
2. [Quadratic minoration] g satisfies

$$g(x) - g(x_*) \geq \frac{1}{2}\underline{r}\|x - x_*\|_\Sigma^2. \quad (16)$$

3. [Reduced Strong Convexity] Assumption [RSC] holds with $\delta = 1$ and $\rho = (\kappa_\Sigma \underline{r})^{-1}$.
4. [Sub-Gaussianity] $\nabla G(x_*, \omega_i)$ is $\sigma^2 \overline{\nu}^2$ -sub Gaussian.

The proof is postponed to the appendix. The last point is a consequence of a generalization of the Restricted Eigenvalue property [5], that we detail below (as it gives insight on why Proposition 3.1 holds).

This condition, that we state and call $\mathbf{Q}(\lambda, \psi)$ in the following Lemma 3.1, and is reminiscent of [26] with the corresponding assumptions of [41, 14].

Lemma 3.1 Let $\lambda > 0$ and $0 < \psi \leq 1$, and suppose that for all subsets $I \subset \{1, \dots, n\}$ of cardinality smaller than s the following property is verified:

$$\forall z \in \mathbf{R}^n \quad \|z_I\|_1 \leq \sqrt{\frac{s}{\lambda}} \|z\|_\Sigma + \frac{1}{2}(1 - \psi) \|z\|_1 \quad \mathbf{Q}(\lambda, \psi)$$

where z_I is obtained by zeroing all its components with indices $i \notin I$.

If $g(\cdot)$ satisfies the quadratic minoration condition, i.e., for some $\mu > 0$,

$$g(x) - g(x_*) \geq \frac{1}{2}\mu \|x - x_*\|_\Sigma^2, \quad (17)$$

and that \hat{x} is an admissible solution to (7) satisfying, with probability at least $1 - \varepsilon$,

$$F_\kappa(\hat{x}) \leq F_\kappa(x_*) + v.$$

Then, with probability at least $1 - \varepsilon$,

$$\|\hat{x} - x_*\|_1 \leq \frac{s\kappa}{\lambda\mu\psi} + \frac{v}{\kappa\psi}. \quad (18)$$

Remark 3.1 Condition $\mathbf{Q}(\lambda, \psi)$ generalizes the classical Restricted Eigenvalue (RE) property [5] and Compatibility Condition [46], and is the most relaxed condition under which classical bounds for the error of ℓ_1 -recovery routines were established. Validity of $\mathbf{Q}(\lambda, \psi)$ with some $\lambda > 0$ is necessary for Σ to possess the celebrated null-space property [13]

$$\exists \psi > 0 : \max_{I, |I| \leq s} \|z_I\|_1 \leq \frac{1}{2}(1 - \psi) \|z\|_1 \quad \forall z \in \text{Ker}(\Sigma)$$

which is necessary and sufficient for the s -goodness of Σ (i.e., $\hat{x} \in \text{Argmin}_u \{\|u\| : \Sigma u = \Sigma x_*\}$ reproduces exactly every s -sparse signal x_* in the noiseless case).

When Σ possesses the nullspace property, $\mathbf{Q}(\lambda, \psi)$ may hold for Σ with nontrivial kernel; this is typically the case for random matrices [41, 42] such as rank deficient Wishart matrices, etc. When Σ is a regular matrix, condition $\mathbf{Q}(\lambda, \psi)$ may also hold with constant λ which is much higher than the minimal eigenvalue of Σ when the eigenspace corresponding to small eigenvalues of Σ does not contain vectors z with $\|z_I\|_1 > \frac{1}{2}(1 - \psi) \|z\|_1$.

Remarks. In the case of linear regression where $\tau(t) = t$, it holds

$$\begin{aligned} g(x) &= \mathbf{E}\left\{\frac{1}{2}(\phi^T x)^2 - x^T \phi \eta\right\} = \frac{1}{2} \mathbf{E}\left\{(\phi^T (x_* - x))^2 - (\phi^T x_*)^2\right\} \\ &= \frac{1}{2}(x - x_*)^T \Sigma (x - x_*) - \frac{1}{2} x_*^T \Sigma x_* = \frac{1}{2} \|x - x_*\|_\Sigma^2 - \frac{1}{2} \|x_*\|_\Sigma^2 \end{aligned}$$

and $\nabla G(x, \omega) = \phi \phi^T (x - x_*) - \sigma \xi \phi$. In this case $\mathcal{L}(\omega) \leq \|\phi \phi^T\|_\infty \leq \bar{v}^2$.

Note that quadratic minoration bound (16) for $g(x) - g(x_*)$ is often overly pessimistic. Indeed, consider for instance, Gaussian regressor $\phi \sim \mathcal{N}(0, \Sigma)$ (such regressors are not a.s. bounded, we consider this example only for illustration purposes) and activation τ , define for some $0 \leq \alpha \leq 1$ (with the convention, $0/0 = 0$)

$$\tau(t) = \begin{cases} t, & |t| \leq 1, \\ \text{sign}(t)[\alpha^{-1}(|t|^\alpha - 1) + 1], & |t| > 1. \end{cases} \quad (19)$$

When passing from ϕ to $\varphi = \Sigma^{-1/2} \phi$ and from x to $z = \Sigma^{1/2} x$ and using the fact that

$$\varphi = \frac{z z^T}{\|z\|_2^2} \varphi + \underbrace{\left(I - \frac{z z^T}{\|z\|_2^2}\right)}_{=: \chi} \varphi$$

with independent $\frac{z z^T}{\|z\|_2^2} \varphi$ and χ , we obtain

$$H(x) = \mathbf{E}\{\phi[\tau(\phi^T x)]\} = \mathbf{E}\left\{\frac{z z^T}{\|z\|_2^2} \varphi \tau(\varphi^T z)\right\} = \frac{z}{\|z\|_2} \mathbf{E}\{\tau(\varsigma \|z\|_2)\} = \frac{\Sigma^{1/2} x}{\|x\|_\Sigma} \mathbf{E}\{\tau(\varsigma \|x\|_\Sigma)\}$$

where $\varsigma \sim \mathcal{N}(0, 1)$. Thus, $H(x)$ is proportional to $\frac{\Sigma^{1/2}x}{\|x\|_\Sigma}$ with coefficient

$$h(\|x\|_\Sigma) = \mathbf{E} \{ \varsigma \mathbf{r}(\varsigma \|x\|_\Sigma) \}.$$

Figure 1 represents the mapping h for different values of α (on the left), along with the corresponding mapping H on a $\|\cdot\|_\Sigma$ -ball centered at the origin of radius r (on the right).




Figure 1: Given the activation function \mathbf{r} in (19) and $\alpha = (0, 0.01, 0.1, 0.25, 1)$; left plot: mappings h ; right plot: moduli of strong monotonicity of mappings H on $\{x : \|x\|_\Sigma \leq r\}$ as function of r .

3.2 Stochastic Mirror Descent algorithm

In this section, we describe the statistical properties of approximate solutions of Algorithm 1 when applied to the sparse recovery problem. We shall use the following distance-generating function of the ℓ_1 -ball of \mathbf{R}^n (cf. [27, Section 5.7.1])

$$\theta(x) = \frac{c}{p} \|x\|_p^p, \quad p = \begin{cases} 2, & n = 2 \\ 1 + \frac{1}{\ln(n)}, & n \geq 3, \end{cases} \quad c = \begin{cases} 2, & n = 2, \\ e \ln n, & n \geq 3. \end{cases} \quad (20)$$

It immediately follows that θ is strongly convex with modulus 1 w.r.t. the norm $\|\cdot\|_1$ on its unit ball, and that $\Theta \leq e \ln n$. In particular, Theorem 2.1 entails the following statement.

Proposition 3.2 *For $t \gtrsim \sqrt{\ln N}$, assuming the samples budget is large enough, i.e., $N \geq m_0$ (so that at least one stage of the preliminary phase of Algorithm 1 is completed), the approximate solution \hat{x}_N output satisfies with probability at least $1 - Ce^{-t} \ln N$,*

$$\|\hat{x}_N - x_*\|_1 \lesssim R \exp \left\{ -c \frac{r\kappa_\Sigma}{r\nu^2} \frac{N}{s(\ln n + t)} \right\} + \frac{\sigma \bar{\nu} s}{r\kappa_\Sigma} \sqrt{\frac{\ln n + t}{N}} \quad (21)$$

The corresponding solution $\hat{x}_N^{(b)}$ of the minibatch variant of the algorithm satisfies with probability $\geq 1 - Ce^{-t} \ln N$,

$$\|\hat{x}_N^{(b)} - x_*\|_1 \lesssim R \exp \left\{ -c \frac{r\kappa_\Sigma}{r\nu^2} \frac{N}{s(\ln n + t)} \right\} + \frac{\sigma \bar{\nu} s}{r\kappa_\Sigma} \sqrt{\frac{\ln n (\ln n + t)}{N}}$$

Remark 3.2 *Bounds for the ℓ_1 -norm of the error $\hat{x}_N - x_*$ (or $\hat{x}_N^{(b)} - x_*$) established in Proposition 3.2 allows us to quantify prediction error $g(\hat{x}_N) - g(x_*)$ (and $g(\hat{x}_N^{(b)}) - g(x_*)$), and also lead to bounds for $\|\hat{x}_N - x_*\|_\Sigma$ and $\|\hat{x}_N - x_*\|_2$ (respectively, for $\|\hat{x}_N^{(b)} - x_*\|_\Sigma$ and $\|\hat{x}_N^{(b)} - x_*\|_2$). For instance, Proposition 2.1 in the present setting implies the bound on the prediction error after N steps of the algorithm that reads*

$$g(\hat{x}_N) - g(x_*) \lesssim \frac{R^2 \kappa_\Sigma r}{s} \exp \left\{ -\frac{c\kappa_\Sigma r}{\delta^2 r\nu^2} \frac{N}{s(\Theta + t)} \right\} + \frac{\sigma^2 \bar{\nu}^2 s(\Theta + t)}{\kappa_\Sigma r N}$$

with probability $\geq 1 - C \ln N e^{-t}$. We conclude by (16) that

$$\begin{aligned} \|\widehat{x}_N - x_*\|_2^2 &\leq \kappa_\Sigma^{-1} \|\widehat{x}_N - x_*\|_\Sigma^2 \leq 2\kappa_\Sigma^{-1} r^{-1} [g(\widehat{x}_N) - g(x_*)] \\ &\lesssim \frac{R^2}{s} \exp \left\{ -\frac{c\kappa_\Sigma r}{\delta^2 \bar{\nu}^2} \frac{N}{s(\Theta + t)} \right\} + \frac{\sigma^2 \bar{\nu}^2 s(\Theta + t)}{\kappa_\Sigma^2 r^2 N}. \end{aligned}$$

In other words, the error $\|\widehat{x}_N - x_*\|_2$ converges geometrically to the “asymptotic rate” $\frac{\sigma \bar{\nu}}{\kappa_\Sigma r} \sqrt{\frac{s(\Theta + t)}{N}}$ which is the “standard” rate established in the setting (cf. [1, 5, 35], etc).

Remark 3.3 The proposed approach allows also to address the situation in which regressors are not a.s. bounded. For instance, consider the case of random regressors with i.i.d sub-Gaussian entries such that

$$\forall j \leq n, \quad \mathbf{E} \left[\exp \left(\frac{[\phi_i]_j^2}{\varkappa^2} \right) \right] \leq 1.$$

Using the fact that the maximum of uniform norms $\|\phi_i\|_\infty$, $1 \leq i \leq m$, concentrates around $\varkappa \sqrt{\ln mn}$ along with independence of noises ξ_i of ϕ_i , the “smoothness” and “sub-Gaussianity” assumptions of Proposition 3.2 can be stated “conditionally” to the event $\{\omega : \max_{i \leq m} \|\phi_i\|_\infty^2 \lesssim \varkappa^2 (\ln[mn] + t)\}$ of probability greater than $1 - e^{-t}$. For instance, when replacing the bound for the uniform norm of regressors with $\varkappa^2 (\ln[mn] + t)$ in the definition of algorithm parameters and combining with appropriate deviation inequality for martingales (cf., e.g., [4]), one arrives at the bound for the error $\|\widehat{x}_N - x_*\|_1$ of Algorithm 1 which is similar to (21) of Proposition 3.2 in which $\bar{\nu}$ is replaced with $\varkappa \sqrt{\ln[mn] + t}$.

3.3 Numerical experiments

In this section, we present results of a small simulation study illustrating the theoretical part of the previous section.² We consider the GLR model (15) with activation function (19) where $\alpha = 1/2$. In our simulations, x_* is an s -sparse vector with s nonvanishing components sampled independently from the standard s -dimensional Gaussian distribution; regressors ϕ_i are sampled from a multivariate Gaussian distribution $\phi \sim \mathcal{N}(0, \Sigma)$, where Σ is a diagonal covariance matrix with diagonal entries $\sigma_1 \leq \dots \leq \sigma_n$. In Figure 2 we report on the experiment in which we compare the performance of the CSMD-SR algorithm from Section 2.3 to that of four other methods. The contenders are (1) “vanilla” non-Euclidean SMD algorithm constrained to the ℓ_1 -ball equipped with the distance generating function (20), (2) composite non-Euclidean dual averaging algorithm (p -Norm RDA) from [47], (3) multistage SMD-SR of [23], and (4) “vanilla” Euclidean SGD. The regularization parameter of the ℓ_1 penalty in (2) is set to the theoretically optimal value $\lambda = 2\sigma \sqrt{2 \log(n)/T}$. The corresponding dimension of the parameter space is $n = 500000$, the sparsity level of the optimal point x_* is $s = 200$, and the “total budget” of oracle calls is $N = 250000$; we use the identity regressor covariance matrix ($\Sigma = I_n$) and $\sigma \in \{0.001, 0.1\}$. To reduce computation time we use the minibatch versions of the multi-stage algorithms—CSMD-SR and algorithm (3), the data to compute stochastic gradient realizations $\nabla G(x_i, \omega) = \phi(\mathbf{r}(\phi^T x_i) - \eta)$ at the current search point x_i being generated “on the fly.” We repeat simulations 20 times and plot the median value along with the first and the last deciles of the error $\|\widehat{x}_i - x_*\|_1$ at each iteration of the algorithm against the number of oracle calls.

²The reader is invited to check Section C of the supplementary material for more experimental results.

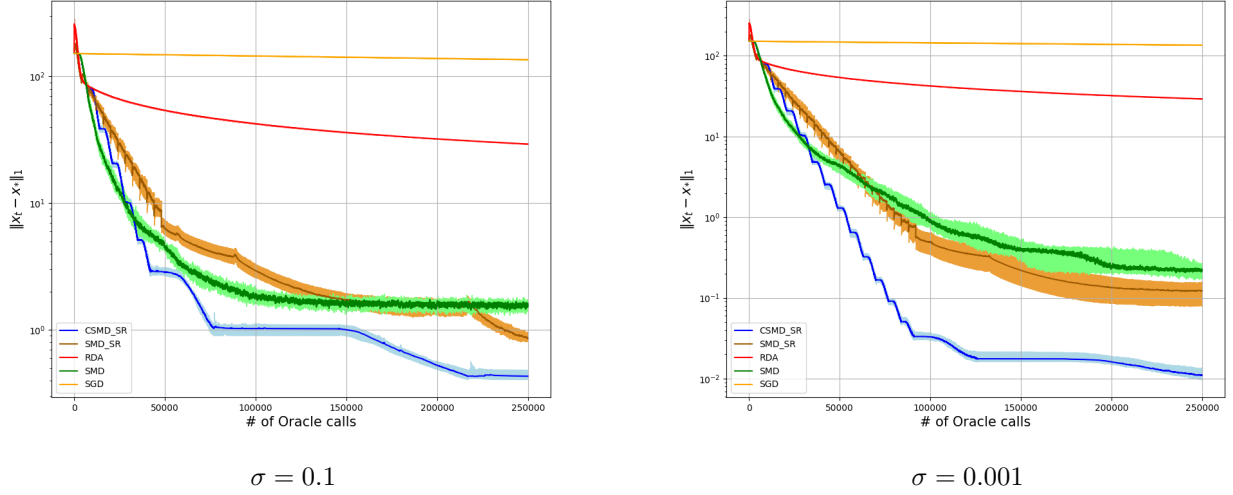


Figure 2: CSMD-SR and “vanilla” SMD in Generalized Linear Regression problem: ℓ_1 error as a function of the number of oracle calls

The proposed method outperforms other algorithms which struggle to reach the regime where the stochastic noise is dominant.




Figure 3: Preliminary stages of the CSMD-SR and its variant with data recycling: linear regression experiment (left pane), GLR with activation $\tau_{1/10}(t)$ (right pane).

In the second experiment we report on here, we study the behavior of the multistage algorithm derived from Algorithm 2 in which, instead of using independent data samples, we reuse the same data at each stage of the method. In Figure 3 we present results of comparison of the CSMD-SR algorithm with its variant with data recycle. This version is of interest as it attains fast the noise regime while using limited amount of samples. In our first experiment, we consider linear regression problem with parameter dimension $n = 100\,000$ and sparsity

level $s = 75$ of the optimal solution; we consider the GLR model (15) with activation function $\tau_{1/10}(t)$ in the second experiment. We choose $\Sigma = I_n$ and $\sigma = 0.001$; we run 14 (preliminary) stages of the algorithm with $m_0 = 3500$ in the first simulation and $m_0 = 4500$ in the second. We believe that the results speak for themselves.

Acknowledgements

This work was supported by Multidisciplinary Institute in Artificial intelligence MIAI @ Grenoble Alpes (ANR-19-P3IA-0003), “Investissements d’avenir” program (ANR20-CE23-0007-01), FAIRPLAY project, LabEx Ecodec (ANR11-LABX-0047), and ANR-19-CE23-0026. The authors would also like to acknowledge CRITEO AI Lab for supporting this work.

References

- [1] A. Agarwal, S. Negahban, and M. J. Wainwright. Fast global convergence of gradient methods for high-dimensional statistical recovery. *The Annals of Statistics*, 40(5):2452 – 2482, 2012.
- [2] A. Agarwal, S. Negahban, and M. J. Wainwright. Stochastic optimization and sparse statistical recovery: Optimal algorithms for high dimensions. In *Advances in Neural Information Processing Systems*, pages 1538–1546, 2012.
- [3] M. Baes, M. Burgisser, and A. Nemirovski. A randomized mirror-prox method for solving structured large-scale matrix saddle-point problems. *SIAM Journal on Optimization*, 23(2):934–962, 2013.
- [4] B. Bercu, B. Delyon, and E. Rio. *Concentration inequalities for sums and martingales*. Springer, 2015.
- [5] P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- [6] L. Birgé, P. Massart, et al. Minimum contrast estimators on sieves: exponential bounds and rates of convergence. *Bernoulli*, 4(3):329–375, 1998.
- [7] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.
- [8] E. Candes, T. Tao, et al. The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics*, 35(6):2313–2351, 2007.
- [9] E. J. Candes and Y. Plan. A probabilistic and riplless theory of compressed sensing. *IEEE transactions on information theory*, 57(11):7235–7254, 2011.
- [10] E. J. Candès, Y. Plan, et al. Near-ideal model selection by ℓ_1 minimization. *The Annals of Statistics*, 37(5A):2145–2177, 2009.
- [11] E. J. Candes, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- [12] G. Chen and M. Teboulle. Convergence analysis of a proximal-like minimization algorithm using bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, 1993.
- [13] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. *Journal of the American mathematical society*, 22(1):211–231, 2009.
- [14] A. Dalalyan and P. Thompson. Outlier-robust estimation of a sparse linear model using ℓ_1 -penalized Huber’s m -estimator. In *Advances in Neural Information Processing Systems*, pages 13188–13198, 2019.

- [15] X. Fan, I. Grama, and Q. Liu. Hoeffding’s inequality for supermartingales. *Stochastic Processes and their Applications*, 122(10):3545–3559, 2012.
- [16] R. Foygel Barber and W. Ha. Gradient descent with non-convex constraints: local concavity determines convergence. *Information and Inference: A Journal of the IMA*, 7(4):755–806, 03 2018.
- [17] P. Gaillard and O. Wintenberger. Sparse accelerated exponential weights. In *20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. arXiv preprint arXiv:1610.05022.
- [18] S. Ghadimi and G. Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization, ii: shrinking procedures and optimal algorithms. *SIAM Journal on Optimization*, 23(4):2061–2089, 2013.
- [19] L. E. Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination for the lasso and sparse supervised learning problems. *arXiv preprint arXiv:1009.4219*, 2010.
- [20] P. Jain, A. Tewari, and P. Kar. On iterative hard thresholding methods for high-dimensional m-estimation. In *Advances in Neural Information Processing Systems*, pages 685–693, 2014.
- [21] A. Juditsky, F. K. Karzan, and A. Nemirovski. On a unified view of nullspace-type conditions for recoveries associated with general sparsity structures. *Linear Algebra and its Applications*, 441:124–151, 2014.
- [22] A. Juditsky, F. Kılınç Karzan, and A. Nemirovski. Randomized first order algorithms with applications to ℓ_1 -minimization. *Mathematical Programming*, 142(1):269–310, 2013.
- [23] A. Juditsky, A. Kulunchakov, and H. Tsytseus. Sparse recovery by reduced variance stochastic approximation. *arXiv preprint arXiv:2006.06365*, 2020.
- [24] A. Juditsky, A. Nazin, A. Tsybakov, and N. Vayatis. Generalization error bounds for aggregation by mirror descent with averaging. In *Advances in neural information processing systems*, pages 603–610, 2006.
- [25] A. Juditsky and A. Nemirovski. First order methods for nonsmooth convex large-scale optimization, II: utilizing problems structure. *Optimization for Machine Learning*, pages 149–183.
- [26] A. Juditsky and A. Nemirovski. Accuracy guarantees for ℓ_1 -recovery. *IEEE Transactions on Information Theory*, 57(12):7818–7839, 2011.
- [27] A. Juditsky and A. Nemirovski. First order methods for nonsmooth convex large-scale optimization, I: general purpose methods. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*. MIT Press Cambridge, 2011.
- [28] A. Juditsky and A. S. Nemirovski. Large deviations of vector-valued martingales in 2-smooth normed spaces. *arXiv preprint arXiv:0809.0813*, 2008.
- [29] A. Juditsky and Y. Nesterov. Deterministic and stochastic primal-dual subgradient algorithms for uniformly convex minimization. *Stochastic Systems*, 4(1):44–80, 2014.
- [30] M. Kowalski, P. Weiss, A. Gramfort, and S. Anthoine. Accelerating ista with an active set strategy. In *OPT 2011: 4th International Workshop on Optimization for Machine Learning*, page 7, 2011.
- [31] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- [32] Y. Lei and K. Tang. Stochastic composite mirror descent: Optimal bounds with high probabilities. *Advances in Neural Information Processing Systems*, 31, 2018.
- [33] H. Liu and R. Foygel Barber. Between hard and soft thresholding: optimal iterative thresholding algorithms. *Information and Inference: A Journal of the IMA*, 9(4):899–933, 2020.

- [34] J. Mairal. *Sparse coding for machine learning, image processing and computer vision*. PhD thesis, Cachan, Ecole normale supérieure, 2010.
- [35] L. Meier, S. Van De Geer, and P. Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- [36] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [37] A. S. Nemirovski and D. B. Yudin. *Complexity of problems and effectiveness of methods of optimization (Russian book)*. Nauka, Moscow, 1979. Translated as *Problem complexity and method efficiency in optimization*, J. Wiley & Sons, New York 1983.
- [38] Y. Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- [39] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical programming*, 140(1):125–161, 2013.
- [40] Y. Nesterov and A. Nemirovski. On first-order algorithms for ℓ_1 /nuclear norm minimization. *Acta Numerica*, 22:509–575, 2013.
- [41] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11(Aug):2241–2259, 2010.
- [42] H. Rauhut. Compressive sensing and structured random matrices. *Theoretical foundations and numerical methods for sparse recovery*, 9(1):92, 2010.
- [43] B. Recht, W. Xu, and B. Hassibi. Null space conditions and thresholds for rank minimization. *Mathematical programming*, 127(1):175–202, 2011.
- [44] S. Shalev-Shwartz and A. Tewari. Stochastic methods for ℓ_1 -regularized loss minimization. *Journal of Machine Learning Research*, 12(Jun):1865–1892, 2011.
- [45] N. Srebro, K. Sridharan, and A. Tewari. Smoothness, low noise and fast rates. In *Advances in neural information processing systems*, pages 2199–2207, 2010.
- [46] S. Van De Geer and P. Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [47] L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(88):2543–2596, 2010.

A Proofs

We use notation \mathbf{E}_i for conditional expectation given x_0 and $\omega_1, \dots, \omega_i$.

A.1 Proof of Proposition 2.1

The result of Proposition 2.1 is an immediate consequence of the following statement.

Proposition A.1 *Let*

$$f(x) = \frac{1}{2}g(x) + h(x), \quad x \in X.$$

In the situation of Section 2.2, let $\gamma_i \leq (4\nu)^{-1}$ for all $i = 0, 1, \dots$, and let \hat{x}_m be defined in (10), where x_i are iterations (9). Then for any $t \geq 2\sqrt{2} + \ln m$ there is $\bar{\Omega}_m \subset \Omega$ such that $\text{Prob}(\bar{\Omega}_m) \geq 1 - 4e^{-t}$ and for all $\omega^m = [\omega_1, \dots, \omega_m] \in \bar{\Omega}_m$,

$$\begin{aligned} \left(\sum_{i=0}^{m-1} \gamma_i \right) [f(\hat{x}_m) - f(x_*)] &\leq \sum_{i=0}^{m-1} \left[\frac{1}{2} \gamma_i \langle \nabla g(x_i), x_i - x_* \rangle + \gamma_{i+1} (h(x_{i+1}) - h(x_*)) \right] \\ &\leq V(x_0, x_*) + \gamma_0 [h(x_0) - h(x_*)] - \gamma_m [h(x_m) - h(x_*)] \\ &\quad + V(x_0, x_*) + 15tR^2 + \sigma_*^2 \left[7 \sum_{i=0}^{m-1} \gamma_i^2 + 24t\bar{\gamma}^2 \right]. \end{aligned} \quad (22)$$

In particular, when using the constant stepsize strategy with $\gamma_i \equiv \gamma$, $0 < \gamma \leq (4\nu)^{-1}$, one has

$$\begin{aligned} &\frac{1}{2} [g(\hat{x}_m) - g(x_*)] + [h(\hat{x}_m) - h(x_*)] \\ &\leq \frac{V(x_0, x_*) + 15tR^2}{\gamma m} + \frac{h(x_0) - h(x_m)}{m} + \gamma \sigma_*^2 \left(7 + \frac{24t}{m} \right). \end{aligned} \quad (23)$$

Proof. Denote $H_i = \nabla G(x_{i-1}, \omega_i)$. In the sequel, we use the shortcut notation $\vartheta(z)$ and $V(x, z)$ for $\vartheta_{x_0}^R(z)$ and $V_{x_0}(x, z)$ when exact values x_0 and R are clear from the context.

1^o. From the definition of x_i and of the composite prox-mapping (8) (cf. Lemma A.1 of [40]), we conclude that there is $\eta_i \in \partial h(x_i)$ such that

$$\langle \gamma_{i-1} H_i + \gamma_i \eta_i + \nabla \vartheta(x_i) - \nabla \vartheta(x_{i-1}), z - x_i \rangle \geq 0, \quad \forall z \in \mathcal{X},$$

implying, as usual [12], that $\forall z \in \mathcal{X}$

$$\langle \gamma_{i-1} H_i + \gamma_i \eta_i, x_i - z \rangle \leq V(x_{i-1}, z) - V(x_i, z) - V(x_{i-1}, x_i).$$

In particular,

$$\begin{aligned} &\gamma_{i-1} \langle H_i, x_{i-1} - x_* \rangle + \gamma_i \langle \eta_i, x_i - x_* \rangle \\ &\leq V(x_{i-1}, x_*) - V(x_i, x_*) - V(x_{i-1}, x_i) + \gamma_{i-1} \langle H_i, x_{i-1} - x_i \rangle \\ &\leq V(x_{i-1}, x_*) - V(x_i, x_*) + \frac{1}{2} \gamma_{i-1}^2 \|H_i\|_*^2. \end{aligned}$$

Observe that due to the Lipschitz continuity of $\nabla G(\cdot, \omega)$ one has

$$\nu \langle \nabla G(x, \omega) - \nabla G(x', \omega), x - x' \rangle \geq \|\nabla G(x, \omega) - \nabla G(x', \omega)\|_*^2, \quad \forall x, x' \in \mathcal{X}, \quad (24)$$

so that

$$\begin{aligned} \|\nabla G(x, \omega)\|_*^2 &\leq 2\|\nabla G(x, \omega) - \nabla G(x_*, \omega)\|_*^2 + 2\|\nabla G(x_*, \omega)\|_*^2 \\ &\leq 2\nu \langle \nabla G(x, \omega) - \nabla G(x_*, \omega), x - x_* \rangle + 2\|\nabla G(x_*, \omega)\|_*^2 \\ &= 2\nu \langle \nabla G(x, \omega), x - x_* \rangle - 2\nu \langle \nabla G(x_*, \omega), x - x_* \rangle + 2\|\nabla G(x_*, \omega)\|_*^2 \end{aligned}$$

so that

$$\begin{aligned} & \gamma_{i-1} \langle H_i, x_{i-1} - x_* \rangle + \gamma_i \langle \eta_i, x_i - x_* \rangle \\ & \leq V(x_{i-1}, x_*) - V(x_i, x_*) + \gamma_{i-1}^2 [\nu \langle H_i, x_{i-1} - x_* \rangle - \nu \zeta_i + \tau_i] \end{aligned}$$

where $\zeta_i = \langle \nabla G(x_*, \omega_i), x_{i-1} - x_* \rangle$ and $\tau_i = \|\nabla G(x_*, \omega)\|_*^2$. As a result, by convexity of h we have for $\gamma_i \leq (4\nu)^{-1}$

$$\begin{aligned} & \frac{3}{4} \gamma_{i-1} \langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle + \gamma_i [h(x_i) - h(x_*)] \\ & \leq (\gamma_{i-1} - \gamma_{i-1}^2 \nu) \langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle + \gamma_i \langle \eta_i, x_i - x_* \rangle \\ & \leq V(x_{i-1}, x_*) - V(x_i, x_*) + (\gamma_{i-1} - \gamma_{i-1}^2 \nu) \langle \xi_i, x_{i-1} - x_* \rangle + \gamma_{i-1}^2 [\tau_i - \nu \zeta_i] \end{aligned}$$

where we put $\xi_i = H_i - \nabla g(x_{i-1})$. When summing from $i = 1$ to m we obtain

$$\begin{aligned} & \sum_{i=1}^m \gamma_{i-1} \left(\frac{3}{4} \langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle + [h(x_{i-1}) - h(x_*)] \right) \\ & \leq V(x_0, x_*) + \underbrace{\sum_{i=1}^m [\gamma_{i-1}^2 (\tau_i - \nu \zeta_i) + \gamma_{i-1} (1 - \gamma_{i-1} \nu) \langle \xi_i, x_{i-1} - x_* \rangle]}_{=: R_m} \\ & \quad + \gamma_0 [h(x_0) - h(x_*)] - \gamma_m [h(x_m) - h(x_*)]. \end{aligned} \tag{25}$$

2°. We have

$$\begin{aligned} \gamma_{i-1} \langle \xi_i, x_{i-1} - x_* \rangle &= \gamma_{i-1} \overbrace{\langle [\nabla G(x_{i-1}, \omega_i) - \nabla G(x_*, \omega_i)] - \nabla g(x_{i-1}), x_{i-1} - x_* \rangle}^{v_i} \\ & \quad + \gamma_{i-1} \langle \nabla G(x_*, \omega_i), x_{i-1} - x_* \rangle \\ &= \gamma_{i-1} [v_i + \zeta_i], \end{aligned}$$

so that

$$R_m = \sum_{i=1}^m \gamma_{i-1}^2 \tau_i + \sum_{i=1}^m (\gamma_{i-1} - \gamma_{i-1}^2 \nu) v_i + \sum_{i=1}^m (\gamma_{i-1} - 2\nu \gamma_{i-1}^2) \zeta_i =: r_m^{(1)} + r_m^{(2)} + r_m^{(3)}. \tag{26}$$

Note that $r_m^{(3)}$ is a sub-Gaussian martingale. Indeed, one has $\mathbf{E}_{i-1} \{\zeta_i\} = 0$ a.s.,³ and

$$|\zeta_i| \leq \|x_{i-1} - x_*\| \|\nabla G(x_*, \omega)\|_*,$$

so that by the sub-Gaussian hypothesis (6), $\mathbf{E}_{i-1} \left\{ \exp \left(\underbrace{\frac{\zeta_i^2}{4R^2 \sigma_*^2}}_{\nu_*^2} \right) \right\} \leq \exp(1)$. As a result (cf. the proof of

Proposition 4.2 in [28]),

$$\forall t \quad \mathbf{E}_{i-1} \{e^{t\zeta_i}\} \leq \exp \left(t \mathbf{E}_{i-1} \{\zeta_i\} + \frac{3}{4} t^2 \nu_*^2 \right) = \exp \left(3t^2 R^2 \sigma_*^2 \right),$$

and applying (31a) to $S_m = r_m^{(3)}$ with

$$r_m = 6R^2 \sigma_*^2 \sum_{i=0}^{m-1} (\gamma_i - 2\nu \gamma_i^2)^2 \leq 6R^2 \sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2$$

³We use notation \mathbf{E}_{i-1} for the conditional expectation given $x_0, \omega_1, \dots, \omega_{i-1}$.

we conclude that for some $\Omega_m^{(3)}$ such that $\text{Prob}(\Omega_m^{(3)}) \geq 1 - e^{-t}$ and all $\omega^m \in \Omega_m^{(3)}$

$$r_m^{(3)} \leq 2 \sqrt{3tR^2\sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2} \leq 3tR^2 + 3\sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2. \quad (27)$$

Next, again by (6), due to the Jensen inequality, $\mathbf{E}_{i-1}\{\tau_i\} \leq \sigma_*^2$, and

$$\mathbf{E}_{i-1}\{\exp(t\|\nabla G(x_*, \omega_i)\|_*)\} \leq \exp(t\mathbf{E}_{i-1}\{\|\nabla G(x_*, \omega_i)\|_*\} + \frac{3}{4}t^2\sigma_*^2) \leq \exp(t\sigma_* + \frac{3}{4}t^2\sigma_*^2).$$

Thus, when setting

$$\mu_i = \gamma_{i-1}\sigma_*, \quad s_i^2 = \frac{3}{2}\gamma_{i-1}\sigma_*^2, \quad \bar{s} = \max_i \gamma_i s_i,$$

$M_m = r_m^{(1)}$, $v_m + h_m = \frac{21}{4}\sigma_*^4 \sum_{i=0}^{m-1} \gamma_i^4$, and applying the bound (31b) of Lemma A.1 we obtain

$$r_m^{(1)} \leq 3\sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2 + \underbrace{\sqrt{21t\sigma_*^4 \sum_{i=0}^{m-1} \gamma_i^4 + 3t\bar{\gamma}^2\sigma_*^2}}_{=:\Delta_m^{(1)}}$$

for $\bar{\gamma} = \max_i \gamma_i$ and $\omega^m \in \Omega_m^{(1)}$ where $\Omega_m^{(1)}$ is of probability at least $1 - e^{-x}$. Because

$$\bar{\gamma}^2 \sum_{i=0}^{m-1} \gamma_i^2 \geq \sum_{i=0}^{m-1} \gamma_i^4,$$

whenever $\sqrt{21t\sigma_*^4 \sum_{i=0}^{m-1} \gamma_i^4} \geq \sum_{i=0}^{m-1} \gamma_i^2$, one has $21t\bar{\gamma}^2 \geq \sum_{i=0}^{m-1} \gamma_i^2$ and

$$21t \sum_{i=0}^{m-1} \gamma_i^4 \leq 21t\bar{\gamma}^2 \sum_{i=0}^{m-1} \gamma_i^2 \leq (21t\bar{\gamma}^2)^2$$

Thus,

$$\Delta_m^{(1)} \leq \min \left[21t\sigma_*^2\bar{\gamma}^2, \sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2 \right] \leq 21t\sigma_*^2\bar{\gamma}^2 + \sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2,$$

and

$$r_m^{(1)} \leq \sigma_*^2 \left[4 \sum_{i=0}^{m-1} \gamma_i^2 + 24t\bar{\gamma}^2 \right] \quad (28)$$

for $\omega^m \in \Omega_m^{(1)}$.

Finally, by the Lipschitz continuity of ∇G (cf. (24)), when taking expectation w.r.t. the distribution of ω_i , we get

$$\begin{aligned} \mathbf{E}_{i-1}\{v_i^2\} &\leq 4R^2\mathbf{E}_{i-1}\{\|\nabla G(x_{i-1}, \omega_i) - \nabla G(x_*, \omega_i)\|_*^2\} \\ &\leq 4R^2\nu\mathbf{E}_{i-1}\{\langle \nabla G(x_{i-1}, \omega_i) - \nabla G(x_*, \omega_i), x_{i-1} - x_* \rangle\} = 4R^2\nu\langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle. \end{aligned}$$

On the other hand, one also has $|v_i| \leq 2\nu\|x_{i-1} - x_i\|^2 \leq 8\nu R^2$. We can now apply Lemma A.2 with $\sigma_i^2 = 4\gamma_{i-1}^2 R^2\nu\langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle$ to conclude that for $t \geq 2\sqrt{2} + \ln m$

$$r_m^{(2)} \leq 4 \underbrace{\sqrt{tR^2\nu \sum_{i=0}^{m-1} \gamma_i^2 \langle \nabla g(x_i), x_i - x_* \rangle + 16t\nu R^2\bar{\gamma}}}_{=:\Delta_m^{(2)}}$$

for all $\omega^m \in \Omega_m^{(2)}$ such that $\text{Prob}(\Omega_m^{(2)}) \geq 1 - 2e^{-t}$. Note that

$$\Delta_m^{(2)} \leq 2tR^2 + \frac{1}{4}\nu \sum_{i=0}^{m-1} \gamma_i^2 \langle \nabla g(x_i), x_i - x_* \rangle,$$

and $\gamma_i \leq (4\nu)^{-1}$, so that

$$r_m^{(2)} \leq \nu \sum_{i=0}^{m-1} \gamma_i^2 \langle \nabla g(x_i), x_i - x_* \rangle + 12tR^2 \leq \frac{1}{4} \sum_{i=0}^{m-1} \gamma_i \langle \nabla g(x_i), x_i - x_* \rangle + 12tR^2 \quad (29)$$

for $\omega^m \in \Omega_m^{(2)}$.

3°. When substituting bounds (27)–(29) into (26) we obtain

$$\begin{aligned} R_m &\leq \frac{1}{4} \sum_{i=0}^{m-1} \gamma_i \langle \nabla g(x_i), x_i - x_* \rangle + 12tR^2 + \sigma_*^2 \left[4 \sum_{i=0}^{m-1} \gamma_i^2 + 24t\bar{\gamma}^2 \right] + 2\sqrt{3tR^2\sigma_*^2 \sum_{i=0}^{m-1} \gamma_i^2} \\ &\leq \frac{1}{4} \sum_{i=0}^{m-1} \gamma_i \langle \nabla g(x_i), x_i - x_* \rangle + 15tR^2 + \sigma_*^2 \left[7 \sum_{i=0}^{m-1} \gamma_i^2 + 24t\bar{\gamma}^2 \right] \end{aligned}$$

for all $\omega^m \in \bar{\Omega}_m = \bigcap_{i=1}^3 \Omega_m^{(i)}$ with $\text{Prob}(\bar{\Omega}_m) \geq 1 - 4e^{-t}$ and $t \geq 2\sqrt{2 + \ln m}$.

When substituting the latter bound into (25) and utilizing the convexity of g and h we arrive at

$$\begin{aligned} &\left(\sum_{i=0}^{m-1} \gamma_i \right) \left(\frac{1}{2} [g(\hat{x}_m) - g(x_*)] + [h(\hat{x}_m) - h(x_*)] \right) \leq \sum_{i=0}^{m-1} \gamma_i \left(\frac{1}{2} [g(x_i) - g(x_*)] + [h(x_i) - h(x_*)] \right) \\ &\leq \sum_{i=1}^m \gamma_{i-1} \left(\frac{1}{2} \langle \nabla g(x_{i-1}), x_{i-1} - x_* \rangle + [h(x_{i-1}) - h(x_*)] \right) \\ &\leq V(x_0, x_*) + 15tR^2 + \sigma_*^2 \left[7 \sum_{i=0}^{m-1} \gamma_i^2 + 24t\bar{\gamma}^2 \right] + \gamma_0 [h(x_0) - h(x_*)] - \gamma_m [h(x_m) - h(x_*)]. \end{aligned}$$

In particular, for constant stepsizes $\gamma_i \equiv \gamma$ we get

$$\begin{aligned} &\frac{1}{2} [g(\hat{x}_m) - g(x_*)] + [h(\hat{x}_m) - h(x_*)] \\ &\leq \frac{V(x_0, x_*) + 15tR^2}{\gamma m} + \frac{h(x_0) - h(x_m)}{m} + \gamma \sigma_*^2 \left(7 + \frac{24t}{m} \right). \end{aligned}$$

This implies the first statement of the proposition.

5°. To prove the bound for the minibatch solution $\hat{x}_m^{(L)} = \left(\sum_{i=0}^{m-1} \gamma_i \right)^{-1} \sum_{i=0}^{m-1} \gamma_i x_i^{(L)}$, it suffices to note that minibatch gradient observation $H(x, \omega^{(L)})$ is Lipschitz-continuous with Lipschitz constant ν , and that $H(x_*, \omega^{(L)})$ is sub-Gaussian with parameter σ_*^2 replaced with $\bar{\sigma}_{*,L}^2 \lesssim \frac{\Theta \sigma_*^2}{L}$, see Lemma A.3. \square

A.2 Deviation inequalities

Let us assume that $(\xi_i, \mathcal{F}_i)_{i=1,2,\dots}$ is a sequence of sub-Gaussian random variables satisfying⁴

$$\mathbf{E}_{i-1} \{ e^{t\xi_i} \} \leq e^{t\mu_i + \frac{t^2 s_i^2}{2}}, \quad a.s. \quad (30)$$

for some *nonrandom* $\mu_i, s_i, s_i \leq \bar{s}$. We denote by $S_n = \sum_{i=1}^n \xi_i - \mu_i$, $r_n = \sum_{i=1}^n s_i^2$, $v_n = \sum_{i=1}^n s_i^4$, $M_n = \sum_{i=1}^n \xi_i^2 - (s_i^2 + \mu_i^2)$, and $h_n = \sum_{i=1}^n 2\mu_i^2 s_i^2$. The following well known result is provided for reader's convenience.

⁴Here, same as above, we denote \mathbf{E}_{i-1} the expectation conditional to \mathcal{F}_{i-1} .

Lemma A.1 For all $x > 0$ one has

$$\text{Prob} \{S_n \geq \sqrt{2xr_n}\} \leq e^{-x}, \quad (31a)$$

$$\text{Prob} \left\{ M_n \geq 2\sqrt{x(v_n + h_n)} + 2x\bar{s}^2 \right\} \leq e^{-x}. \quad (31b)$$

Proof. The inequality (31a) is straightforward. To prove (31b), note that for $t < \frac{1}{2}\bar{s}^{-2}$ and $\eta \sim \mathcal{N}(0, 1)$ independent of ξ_0, \dots, ξ_n , we have:

$$\begin{aligned} \mathbf{E}_{i-1} \left\{ e^{t\xi_i^2} \right\} &= \mathbf{E}_{i-1} \left\{ \mathbf{E}_\eta \left\{ e^{\sqrt{2i}\xi_i\eta} \right\} \right\} = \mathbf{E}_\eta \left\{ \mathbf{E}_{i-1} \left\{ e^{\sqrt{2i}\xi_i\eta} \right\} \right\} \\ &\leq \mathbf{E}_\eta \left\{ \exp \left\{ \sqrt{2i}\eta\mu_i + t\eta^2 s_i^2 \right\} \right\} = (1 - 2ts_i^2)^{-1/2} \exp \left\{ \frac{t\mu_i^2}{1 - 2ts_i^2} \right\} \quad \text{a.s.}, \end{aligned}$$

and because, cf [31, Lemma 1],

$$-\frac{1}{2}\ln(1 - 2ts_i^2) + \frac{t\mu_i^2}{1 - 2ts_i^2} - t(s_i^2 + \mu_i^2) \leq \frac{t^2 s_i^2 (s_i^2 + 2\mu_i^2)}{1 - 2ts_i^2} \leq \frac{t^2 s_i^2 (s_i^2 + 2\mu_i^2)}{1 - 2t\bar{s}^2},$$

one has for $t < \frac{1}{2}\bar{s}^{-2}$

$$\mathbf{E} \left\{ e^{tM_n} \right\} \leq \exp \left\{ \frac{t^2(v_n + h_n)}{1 - 2t\bar{s}^2} \right\}.$$

By Lemma 8 of [6], this implies that

$$\text{Prob} \left\{ M_n \geq 2\sqrt{x(v_n + h_n)} + 2x\bar{s}^2 \right\} \leq e^{-x}$$

for all $x > 0$. □

Now, suppose that ζ_i , $i = 1, 2, \dots$ is a sequence of random variables satisfying

$$\mathbf{E}_{i-1} \{\zeta_i\} = \mu_i, \quad \mathbf{E}_{i-1} \{\zeta_i^2\} \leq \sigma_i^2, \quad |\zeta_i| \leq 1 \quad \text{a.s.} \quad (32)$$

Denote $M_n = \sum_{i=1}^n [\zeta_i - \mu_i]$ and $q_n = \sum_{i=1}^n \sigma_i^2$. Note that $q_n \leq n$.

Lemma A.2 Let $x \geq 1$; one has

$$\text{Prob} \left\{ M_n \geq \sqrt{2xq_n} + x \right\} \leq \left[e \left(2x \ln \left[\frac{9n}{2x} \right] + 1 \right) + 1 \right] e^{-x}.$$

In particular, for $x \geq 4\sqrt{2 + \ln n}$ one has

$$\text{Prob} \left\{ M_n \geq \sqrt{2xq_n} + x \right\} \leq 2e^{-x/2}.$$

Proof. In the premise of the lemma, applying Bernstein's inequality for martingales [4, 15] we obtain for all $x > 0$ and $u > 0$,

$$\text{Prob} \left\{ M_n \geq \sqrt{2xu} + \frac{x}{3}, \quad q_n \leq u \right\} \leq e^{-x}.$$

We conclude that

$$\text{Prob} \left\{ M_n \geq x, \quad q_n \leq \frac{2x}{9} \right\} \leq e^{-x},$$

and for any $u > 0$

$$\text{Prob} \left\{ M_n \geq \sqrt{2(x+1)q_n} + \frac{x}{3}, \quad u \leq q_n \leq u(1+1/x) \right\} \leq e^{-x},$$

so that

$$\delta_n(x; u) := \text{Prob} \left\{ M_n \geq \sqrt{2xq_n} + \frac{x}{3}, \quad u \leq q_n \leq u(1+1/x) \right\} \leq e^{-x+1}.$$

Let now $u_0 = 2x/9$, $u_j = \min\{n, (1 + 1/x)^j u_0\}$, $j = 0, \dots, J$, with

$$J = \lfloor \ln [n/u_0] \ln^{-1}[1 + 1/x] \rfloor.$$

Note that $\ln[1 + 1/x] \geq 1/(2x)$ for $x \geq 1$, so that

$$J \leq \ln [n/u_0] \ln^{-1}[1 + 1/x] + 1 \leq 2x \ln [n/u_0] + 1.$$

On the other hand,

$$\begin{aligned} \text{Prob} \left\{ M_n \geq \sqrt{2xq_n} + x \right\} &\leq e^{-x} + \sum_{j=1}^J \delta_n(x; u_j) \leq e^{-x} + J e^{-x+1} \\ &\leq \left[e \left(2x \ln \left[\frac{9n}{2x} \right] + 1 \right) + 1 \right] e^{-x} \end{aligned}$$

Finally, we verify explicitly that for $x \geq 4\sqrt{2 + \ln n}$ one has

$$\left[e \left(2x \ln \left[\frac{9n}{2x} \right] + 1 \right) + 1 \right] e^{-x/2} \leq 2,$$

implying that for such x

$$\text{Prob} \left\{ M_n \geq \sqrt{2xq_n} + x \right\} \leq 2e^{-x/2}. \quad \square$$

Let $(\xi_i)_{i=1, \dots}$ be a sequence of independent random vectors in \mathbf{R}^n such that

$$\mathbf{E}_{i-1} \left\{ \exp \left(\frac{\|\xi_i\|_*^2}{s^2} \right) \right\} \leq \exp(1),$$

and let $\eta = \sum_{i=1}^m \xi_i$, $m \in \mathbf{Z}_+$. We are interested in “sub-Gaussian characteristics” of r.v. $\zeta = \langle u, \eta \rangle$ for some $u \in \mathbf{R}^n$, $\|u\| \leq R$, and of $\tau = \|\eta\|_*$.

Because $\mathbf{E}\{\langle u, \xi_i \rangle\} = 0$ and $|\langle u, \xi_i \rangle| \leq \|u\| \|\xi_i\|_*$, for all t one has (cf., e.g., Proposition 4.2 of [28])

$$\mathbf{E} \left\{ e^{t\langle u, \eta \rangle} \right\} = \prod_{i=1}^m \mathbf{E} \left\{ e^{t\langle u, \xi_i \rangle} \right\} \leq \prod_{i=1}^m \exp \left(\frac{3}{4} t^2 s^2 \right) = \exp \left(\frac{3}{4} m t^2 s^2 \right).$$

Let ξ_ℓ , $\ell = 1, 2, \dots$ be a sequence of independent random vectors $\xi_\ell \in E$, such that $\mathbf{E}\{\xi_\ell\} = 0$ and $\mathbf{E}\left\{ e^{\|\xi_\ell\|_*^2/s^2} \right\} \leq \exp(1)$. Denote $\eta_j = \sum_{\ell=1}^j \xi_\ell$. We have the following result.

Lemma A.3

$$\forall L \in \mathbf{Z}_+ \quad \mathbf{E} \left\{ \exp \left(\frac{\|\eta_L\|_*^2}{10\Theta s^2 L} \right) \right\} \leq \exp(1) \quad (33)$$

where $\Theta = \max_{\|z\| \leq 1} \theta(z)$ for the d.-g.f. θ of the unit ball of norm $\|\cdot\|$ in E , as defined in Section 2.2.

Proof. Let for $\eta \in E$, $\pi(\eta) = \sup_{\|z\| \leq 1} [\langle \eta, z \rangle - \theta(z)]$. Observe that for all $\beta > 0$,

$$\|\eta_L\|_* = \sup_{\|z\| \leq 1} \langle \eta_L, z \rangle \leq \max_{\|z\| \leq 1} \beta \theta(z) + \beta \pi(\eta_L/\beta) \leq \beta \Theta + \beta \pi \left(\frac{\eta_L}{\beta} \right). \quad (34)$$

On the other hand, we know (cf. [38, Lemma 1]) that π is smooth with $\|\nabla \pi\| \leq 1$, and $\nabla \pi$ is Lipschitz-continuous w.r.t. to $\|\cdot\|_*$, i.e.,

$$\|\nabla \pi(z) - \nabla \pi(z')\| \leq \|z - z'\|_* \quad \forall z, z' \in E.$$

As a consequence of Lipschitz continuity of π , when denoting $\pi_\beta(\eta) = \beta \pi(\frac{\eta}{\beta})$, we have

$$\pi_\beta(\eta_{j-1} + \xi_j) - \pi_\beta(\eta_{j-1}) \leq \|\xi_j\|_*,$$

so that $\mathbf{E} \left\{ \exp \left([\pi_\beta(\eta_j) - \pi_\beta(\eta_{j-1})]^2 / s^2 \right) \right\} \leq \exp(1)$. Furthermore,

$$\pi_\beta(\eta_{j-1} + \xi_j) \leq \pi_\beta(\eta_{j-1}) + \langle \nabla \pi_\beta(\eta_{j-1}), \xi_j / \beta \rangle + \|\xi_j\|_*^2 / \beta,$$

and, because η_{j-1} does not depend on ξ_j and $\mathbf{E}\{\|\xi_j\|_*^2\} \leq s^2$, we get

$$\mathbf{E}_{j-1} \{ \pi_\beta(\eta_j) - \pi_\beta(\eta_{j-1}) \} \leq s^2 / \beta.$$

By [28, Proposition 4.2] we conclude that random variables $\delta_j = \pi_\beta(\eta_j) - \pi_\beta(\eta_{j-1})$ satisfy for all $t \geq 0$,

$$\mathbf{E}_{j-1} \{ e^{t\delta_j} \} \leq \exp \left(ts^2\beta^{-1} + \frac{3}{4}t^2s^2 \right).$$

Consequently,

$$\mathbf{E} \left\{ e^{t\pi_\beta(\eta_L)} \right\} \leq \mathbf{E} \left\{ e^{t\pi_\beta(\eta_{L-1})} \right\} \exp \left(ts^2\beta^{-1} + \frac{3}{4}t^2s^2 \right) \leq \exp \left(ts^2L\beta^{-1} + \frac{3}{4}t^2s^2L \right).$$

When substituting the latter bound into (34), we obtain for $\beta^2 = s^2L/\Theta$

$$\mathbf{E} \left\{ e^{t\|\eta_L\|_*} \right\} \leq \exp \left(ts\sqrt{\Theta L} + \frac{3}{4}t^2s^2L \right) \quad \forall t \geq 0. \quad (35)$$

To complete the proof of the lemma, it remains to show that (35) implies (33). This is straightforward. Indeed, for $\chi \sim \mathcal{N}(0, 1)$, $\alpha > 0$ and $\zeta = \|\eta_L\|_*$ one has

$$\begin{aligned} \mathbf{E} \left\{ e^{\alpha\zeta^2} \right\} &= \mathbf{E} \left\{ \mathbf{E}_\eta \left(e^{\sqrt{2\alpha}\zeta\chi} \right) \right\} = \mathbf{E}_\chi \left\{ \mathbf{E} \left\{ e^{\sqrt{2\alpha}\zeta\chi} \right\} \right\} \\ &\leq \mathbf{E}_\chi \left\{ \exp \left(\sqrt{2\alpha\Theta L} s\chi + \frac{3}{2}\alpha N s^2 \chi^2 \right) \right\} = (1 - 3\alpha N s^2)^{-1/2} \exp \left\{ \frac{\alpha\Theta L s^2}{1 - 3\alpha N s^2} \right\} \end{aligned}$$

When setting $\alpha = (10\Theta s^2 L)^{-1}$, we conclude that

$$\mathbf{E} \left\{ e^{\alpha\zeta^2} \right\} \leq \exp(1)$$

due to $\Theta \geq 1/2$. □

A.3 Proof of Theorem 2.1

We start with analysing the behaviour of the approximate solution $\widehat{x}_{m_0}^k$ at the stages of the preliminary phase of the procedure.

Lemma A.4 *Let $m_0 = \lceil 64\delta^2\rho\nu s(4\Theta + 60t) \rceil$ (here $\lceil a \rceil$ stands for the smallest integer greater or equal to a), $\gamma = (4\nu)^{-1}$, and let t satisfy $t \geq 4\sqrt{2 + \log(m_0)}$.*

Suppose that $R \geq 2\delta\sigma_\sqrt{6\rho s/\nu}$, that initial condition x_0 of Algorithms 1 and 2 satisfies $\|x_0 - x_*\| \leq R$, and that at the stage k of the preliminary phase we choose*

$$\kappa_k = R_{k-1} \sqrt{\frac{\nu(4\Theta + 60t)}{\rho s m_0}} \quad (36)$$

where $(R_k)_{k \geq 0}$ is defined recursively:

$$R_{k+1} = \frac{1}{2}R_k + \frac{16\sigma_*^2\delta^2\rho s}{\nu R_k}, \quad R_0 = R.$$

Then the approximate solution $\widehat{x}_{m_0}^k$ at the end of the k th stage of the CSMD-SR algorithm satisfies, with probability $\geq 1 - 4ke^{-t}$

$$\|\widehat{x}_{m_0}^k - x_*\| \leq R_k \leq 2^{-k}R + 4\sigma_*\delta\sqrt{2\rho s/\nu}. \quad (37)$$

In particular, the estimate $\widehat{x}_{m_0}^{\overline{K}_1}$ after $\overline{K}_1 = \left\lceil \frac{1}{2} \log_2 \left(\frac{R^2\nu}{32\sigma_^2\delta^2\rho s} \right) \right\rceil$ stages satisfies with probability at least $1 - 4\overline{K}_1 e^{-t}$*

$$\|\widehat{x}_{m_0}^{\overline{K}_1} - x_*\| \leq 8\sigma_*\delta\sqrt{2\rho s/\nu}. \quad (38)$$

Proof of the lemma.

1^o. Note that initial point x_0 satisfies $x_0 \in X_R(x_*)$. Suppose that the initial point $x_0^k = \hat{x}_{m_0}^{k-1}$ of the k th stage of the method satisfy $x_0^k \in X_{R_{k-1}}(x_*)$ with probability $1 - 4(k-1)e^{-t}$. In other words, there is a set $\mathcal{B}_{k-1} \subset \Omega$, $\text{Prob}(\mathcal{B}_{k-1}) \geq 1 - 4(k-1)e^{-t}$, such that for all $\bar{\omega}^{k-1} = [\omega_1; \dots; \omega_{m_0(k-1)}] \in \mathcal{B}_{k-1}$ one has $x_0^k \in X_{R_{k-1}}(x_*)$. Let us show that upon termination of the k th stage $\hat{x}_{m_0}^k$ satisfy $\|\hat{x}_{m_0}^k - x_*\| \leq R_k$ with probability $1 - 4ke^{-t}$. By Proposition A.1 (with $h(x) = \kappa_k \|x\|$) we conclude that for some $\bar{\Omega}_k \subset \Omega$, $\text{Prob}(\bar{\Omega}_k) \geq 1 - 4e^{-t}$, solution $\hat{x}_{m_0}^k$ after m_0 iterations of the stage satisfies, for all for all $\omega^k = [\omega_{(k-1)m_0+1}, \dots, \omega_{km_0}] \in \bar{\Omega}_k$,

$$F(\hat{x}_{m_0}^k) - F(x_*) \leq \frac{1}{m_0} (\nu R_{k-1}^2 (4\Theta + 60t) + \kappa_k R_{k-1}) + \frac{\sigma_*^2}{\nu} \left(\frac{7}{4} + \frac{6t}{m_0} \right).$$

When using the relationship (14) of Assumption [RSC] we now get

$$\|\hat{x}_{m_0}^k - x_*\| \leq \delta \left[\rho s \kappa_k + \frac{R_{k-1}}{m_0} + \frac{\nu R_{k-1}^2}{\kappa_k m_0} (4\Theta + 60t) + \frac{\sigma_*^2}{\nu \kappa_k} \left(\frac{7}{4} + \frac{6t}{m_0} \right) \right]. \quad (39)$$

Note that κ_k as defined in (36) satisfies $\kappa_k \leq R_{k-1} (8\delta \rho s)^{-1}$, while $\kappa_k m_0 \geq 8\delta (4\Theta + 60t) R_{k-1} \nu$. Because $m_0 \geq 3840t$ due to $\rho \nu \geq 1$ and $\delta \geq 1$, one also has $\left(\frac{7}{4} + \frac{6t}{m_0} \right) \kappa_k^{-1} < 16\delta \rho s / R_{k-1}$. When substituting the above bounds into (39) we obtain

$$\|\hat{x}_{m_0}^k - x_*\| \leq \delta R_{k-1} \left(\frac{1}{4\delta} + \frac{1}{m_0} \right) + \frac{16\delta^2 \rho s \sigma_*^2}{R_{k-1} \nu} \leq \frac{1}{2} R_{k-1} + \frac{16\delta^2 \rho s \sigma_*^2}{R_{k-1} \nu} = R_k. \quad (40)$$

We conclude that $\hat{x}_{m_0}^k \in X_{R_k}(x_*)$ for all $\bar{\omega}^k \in \mathcal{B}_k = \mathcal{B}_{k-1} \cap \bar{\Omega}_k$, and

$$\text{Prob}(\mathcal{B}_k) \geq \text{Prob}(\mathcal{B}_{k-1}) - \text{Prob}(\bar{\Omega}_k^c) \geq 1 - 4ke^{-t}.$$

2^o. Let now $a = 16\delta^2 \rho s \sigma_*^2 / \nu$, and let us study the behaviour of the sequence

$$R_k = \frac{R_{k-1}}{2} + \frac{a}{R_{k-1}} =: f(R_{k-1}), \quad R_0 = R \geq \sqrt{2a}.$$

Function f admits a fixed point at $R = \sqrt{2a}$ which is also the minimum of f , so one has $R_k \geq \sqrt{2a} \forall k$. Thus,

$$d_k := R_k - \sqrt{2a} = \frac{R_{k-1} - \sqrt{2a}}{2} + \frac{2a - \sqrt{2a} R_{k-1}}{2R_{k-1}} \leq \frac{1}{2} d_{k-1} \leq 2^{-k} d_0 \leq 2^{-k} (R - \sqrt{2a}).$$

We deduce that $R_k \leq 2^{-k} R_0 + \sqrt{2a}$ which is (37). Finally, after running \bar{K}_1 stages of the preliminary phase, the estimate $\hat{x}_{m_0}^{\bar{K}_1}$ satisfies

$$\|\hat{x}_{m_0}^{\bar{K}_1} - x_*\| \leq 8\delta \sigma_* \sqrt{2\rho s / \nu}. \quad \square$$

We turn next to the analysis of the asymptotic phase of Algorithm 2. We assume that the preliminary phase of the algorithm has been completed.

Lemma A.5 *Let t be such that $t \geq 4\sqrt{2 + \log(m_1)}$, with $m_1 = \lceil 81\delta^2 \rho s \nu (4\Theta + 60t) \rceil$, $\gamma = (4\nu)^{-1}$, and let $\ell_k = \lceil 10 \times 4^{k-1} \Theta \rceil$. We set*

$$\kappa_k = r_{k-1} \sqrt{\frac{\nu(4\Theta + 60t)}{\rho s m_1}}, \quad r_k = 2^{-k} r_0, \quad r_0 = 8\delta \sigma_* \sqrt{2\rho s / \nu}.$$

Then the approximate solution by Algorithm 2 $\hat{x}_{m_1}^k$ at the end of the k th stage of the asymptotic phase satisfies, with probability $\geq 1 - 4(\bar{K}_1 + k)e^{-t}$, $\|\hat{x}_{m_1}^k - x_\| \leq r_k$, implying that*

$$\|\hat{x}_{m_1}^k - x_*\| \lesssim \delta^2 \sigma_* \rho s \sqrt{\frac{\Theta(\Theta + t)}{N_k}}, \quad (41)$$

where $N_k = m_1 \sum_{i=1}^k \ell_i$ is the total count of oracle calls for k asymptotic stages.

Proof of the lemma. Upon terminating the preliminary phase, the initial condition $x_0 = \widehat{x}_{m_0}^{\overline{K}_1}$ of the asymptotic phase satisfies (38) with probability greater or equal to $1 - 4\overline{K}_1 e^{-t}$. We are to show that $\forall k \geq 1$, with probability at least $1 - 4(\overline{K}_1 + k)e^{-t}$,

$$\|\widehat{x}_{m_1}^k - x_*\| \leq r_k = 2^{-k} r_0, \quad r_0 = 8\delta\sigma_*\sqrt{2\rho s/\nu}.$$

The claim is obviously true for $k = 0$. Let us suppose that it holds at stage $k - 1 \geq 0$, and let us prove that it also holds at stage k . To this end, we reproduce the argument used in the proof of Lemma A.4, while taking into account that now ℓ_k observations are averaged at each iteration of the CSMD algorithm. Recall (cf. Lemma A.3) that this amounts to replacing sub-Gaussian parameter σ_*^2 with $\overline{\sigma}_*^2 = 10\Theta\sigma_*^2/\ell_k$. When applying the result of Proposition A.1 and the bound of (14) we conclude (cf. (39)) that, with probability $1 - (\overline{K}_1 + k)e^{-t}$,

$$\|\widehat{x}_{m_1}^k - x_*\| \leq \delta \left[\rho s \kappa_k + \frac{r_{k-1}}{m_1} + \frac{\nu r_{k-1}^2}{\kappa_k m_1} (4\Theta + 60t) + \frac{10\Theta\sigma_*^2}{\nu\kappa_k\ell_k} \left(\frac{7}{4} + \frac{6t}{m_1} \right) \right]$$

By simple algebra, we obtain the following analogue of (40):

$$\|\widehat{x}_{m_1}^k - x_*\| < \delta r_{k-1} \left(\frac{2}{9\delta} + \frac{1}{m_1} \right) + 10 \frac{4^{-k+1} \delta^2 \rho s \sigma_*^2}{r_{k-1} \nu} < \frac{r_{k-1}}{4} + \frac{r_{k-1}}{4} = r_k.$$

Observe that upon the end of the k th stage we used $N_k = m_1 \sum_{i=1}^k \ell_k < 3m_1\Theta \sum_{j=1}^k 4^{j-1} \leq 4^k \Theta m_1$ observations of the asymptotic stage. As a consequence, $4^{-k} < \Theta m_1 / N_k$ and

$$r_k = 2^{-k} r_0 \lesssim \delta^2 \sigma_* \sqrt{\frac{\Theta(\Theta + t) s \nu \rho}{N_k}}. \quad \square$$

Assuming that the preliminary phase of Algorithm 1 was completed, we now consider the asymptotic phase of the algorithm.

Lemma A.6 *Let $t \geq 4\sqrt{2 + \log m_k}$, $m_k = \lceil 4^{k+4}(4\Theta + 60t)\delta^2 \rho s \nu \rceil$,*

$$\gamma^k = \frac{r_{k-1}}{2\sigma_*} \sqrt{\frac{(4\Theta + 60t)}{2m_k}}, \quad \kappa_k^2 = \frac{5\sigma_* r_{k-1}}{\rho s} \sqrt{\frac{(4\Theta + 60t)}{m_k}} \quad (42)$$

where

$$r_k := 2^{-k} r_0, \quad r_0 = 8\delta\sigma_*\sqrt{2\rho s/\nu}.$$

Then the approximate solution $\widehat{x}_{m_k}^k$ upon termination of the k th asymptotic stage satisfies with probability $\geq 1 - 4(\overline{K}_1 + k)e^{-t}$

$$\|\widehat{x}_{m_k}^k - x_*\| \leq 2^{-k} r_0 \lesssim 2^{-k} \sigma_* \delta \sqrt{\rho s \nu^{-1}} \lesssim \delta^2 \sigma_* \rho s \sqrt{\frac{\Theta + t}{N_k}} \quad (43)$$

where $N_k = \sum_{j=1}^k m_j$ is the total iteration count of k stages of the asymptotic phase.

Proof of the lemma.

We are to show that $\forall k \geq 0$, $\|\widehat{x}_{m_k}^k - x_*\| \leq r_k$ with probability $\geq 1 - 4(\overline{K}_1 + k)e^{-t}$ is true. By Lemma A.4, the claim is true for $k = 0$ (at the start of the asymptotic phase, the initial condition $x_0 = \widehat{x}_{m_0}^{\overline{K}_1}$ satisfies the bound (38)). We now assume it to hold for $k - 1 \geq 0$, our objective is to implement the recursive step $k - 1 \rightarrow k$ of the proof. First, observe that the choice of γ^k in (42) satisfies $\gamma^k \leq (4\nu)^{-1}$, $k = 1, \dots$, so that Proposition A.1 can be applied. From the result of the proposition and bound (14) we conclude (cf. (39)) that it holds, with probability $1 - (\overline{K}_1 + k)e^{-t}$,

$$\|\widehat{x}_{m_k}^k - x_*\| \leq \delta \left[\rho s \kappa_k + \frac{r_{k-1}}{m_k} + \frac{r_{k-1}^2 (4\Theta + 60t)}{\gamma^k \kappa_k m_k} + 8 \frac{\gamma^k \sigma_*^2}{\kappa_k} \right]$$

When substituting the value of γ^k from (42) we obtain

$$\|\hat{x}_{m_k}^k - x_*\| \leq \delta \left[\rho s \kappa_k + \frac{r_{k-1}}{m_k} + \frac{4\sigma_* r_{k-1}}{\kappa_k} \sqrt{\frac{2(4\Theta + 60t)}{m_k}} \right],$$

which, by the choice of κ_k in (42), results in results in

$$\|\hat{x}_{m_k}^k - x_*\|^2 \leq 2\delta^2 \left[10\rho s \sigma_* r_{k-1} \sqrt{\frac{4\Theta + 60t}{m_k}} + \frac{r_{k-1}^2}{m_k^2} \right] \leq \frac{r_{k-1}^2}{4} = r_k^2.$$

It remains to note that the total number $N_k = \sum_{j=1}^k m_j$ of iterations during k stages of the asymptotic phase satisfies $N_k \lesssim 4^k(\Theta + t)\delta^2\rho s\nu$, and $2^{-k} \lesssim \delta\sqrt{\frac{(\Theta+t)\rho s\nu}{N_k}}$, which along with definition of r_0 implies (43). \square

Proof of Theorem 2.1. We can now terminate the proof of the theorem. Let us prove the accuracy bound of the theorem for the minibatch variant of the procedure.

Assume that the “total observation budget” N is such that only the preliminary phase of the procedure is implemented. This is the case when either $m_0\bar{K}_1 \geq N$, or $m_0\bar{K}_1 < N$ and $m_0\bar{K}_1 + m_1\ell_1 > N$. The output \hat{x}_N of the algorithm is then the last update of the preliminary phase, and by Lemma A.4 it satisfies $\|\hat{x}_N - x_*\| \leq R2^{-k}$ where k is the count of completed stages. In the case of $m_0\bar{K}_1 \geq N$ this clearly implies that (recall that $N \geq m_0$) that $k \geq cN/m_0$ and, with probability $\geq 1 - 4ke^{-t}$

$$\|\hat{x}_N - x_*\| \lesssim R \exp \left\{ -\frac{c'N}{\delta^2\rho s\nu(\Theta + t)} \right\}. \quad (44)$$

On the other hand, when $m_0\bar{K}_1 < N < m_0\bar{K}_1 + m_1\ell_1$, by definition of m_1 and ℓ_1 , one has $N \leq Cm_0\bar{K}_1$, so that bound (44) still holds in this case.

Now, consider the case where at least one asymptotic stage has been completed. When $m_0\bar{K}_1 > \frac{N}{2}$ we still have $N \leq Cm_0\bar{K}_1$, so that the bound (44) holds for the approximate solution $\hat{x}_N^{(b)}$ at the end of the asymptotic stage. Otherwise, the number of oracle calls N_k of asymptotic stages satisfies $N_k \geq N/2$, and by (41) this implies that with probability $\geq 1 - 4(\bar{K}_1 + \bar{K}_2)e^{-t}$,

$$\|\hat{x}_N^{(b)} - x_*\| \lesssim \delta^2\sigma_*\rho s \sqrt{\frac{\Theta(\Theta + t)}{N}}.$$

To summarize, in both cases, the bound of Theorem 2.1 holds with probability at least $1 - 4(\bar{K}_1 + \bar{K}_2)e^{-t}$.

The proof of the accuracy bound for the “standard” solution \hat{x}_N is completely analogous, making use of the bound (43) of Lemma A.6 instead of (41). \square

Remark A.1 *Theorem 2.1 as stated in Section 2.3 does not say anything about convergence of $g(\hat{x}_N)$ to $g(x_*)$. Such information can be easily extracted from the proof of the theorem. Indeed, observe that at the end of a stage of the method, one has, with probability $1 - Cke^{-t}$,*

$$F_{\kappa_k}(\hat{x}^k) - F_{\kappa_k} \leq v_k,$$

or

$$g(\hat{x}^k) - g(x_*) \leq v_k + \kappa_k(\|\hat{x}^k\| - \|x_*\|) \leq v_k + \kappa_k\|\hat{x}^k - x_*\|$$

where \hat{x}^k is the approximate solution at the end of the stage k . On the other hand, at the end of the k th stage of the preliminary phase one has $\|\hat{x}^k - x_*\| \leq R_k \leq 2^{-k}R$, with $\kappa_k \lesssim R_k(\delta\rho s)^{-1} \leq 2^{-k}R(\delta\rho s)^{-1}$ and $v_k \lesssim \frac{4^{-k}R^2}{\delta^2\rho s}$ implying that

$$g(\hat{x}^k) - g(x_*) \lesssim v_k + \frac{R_k^2}{\delta^2\rho s} \lesssim (\delta^{-2} + \delta^{-1}) \frac{R^2}{\rho s} \exp \left\{ -\frac{c}{\delta\rho\nu} \frac{N}{s(\Theta + t)} \right\}$$

where N is the current iteration count. Furthermore, at the end of the k th asymptotic stage, one has, with probability $1 - (\bar{K}_1 + k)e^{-t}$, $\|\hat{x}^k - x_*\| \leq R_k \lesssim \delta^2 \sigma_* \rho s \sqrt{\frac{\Theta+t}{m_k}}$, while $\kappa_k \asymp 2^{-k} \delta \sigma_* (\rho \nu s)^{-1/2} \lesssim \delta \sigma_* \sqrt{\frac{\Theta+t}{m_k}}$, and $v_k \lesssim \delta^2 \sigma_*^2 \rho s (\Theta + t) / m_k$. As a result, the corresponding \hat{x}^k satisfies

$$g(\hat{x}^k) - g(x_*) \leq v_k + \kappa_k \|\hat{x}^k - x_*\| \lesssim (\delta^2 + \delta^3) \rho \sigma_*^2 s \frac{\Theta + t}{m_k}.$$

When putting the above bounds together, assuming that at least 1 stage of the algorithm was completed, we arrive at the bound after N steps:

$$g(\hat{x}_N) - g(x_*) \lesssim (\delta^{-2} + \delta^{-1}) \frac{R^2}{\rho s} \exp \left\{ -\frac{c}{\delta^2 \rho \nu} \frac{N}{s(\Theta + t)} \right\} + (\delta^2 + \delta^3) \rho \sigma_*^2 s \frac{\Theta + t}{N} \quad (45)$$

with probability $1 - (\bar{K}_1 + \bar{K}_2)e^{-t}$.

A.4 Proof of Proposition 3.1

1°. Recall that \mathfrak{r} is \bar{r} -Lipschitz continuous, i.e., for all $t, t' \in \mathbf{R}^m$

$$|\mathfrak{r}(t) - \mathfrak{r}(t')| \leq \bar{r}|t - t'|.$$

As a result, for all $x, x' \in X$,

$$\|\phi[\mathfrak{r}(\phi_i^T x) - \mathfrak{r}(\phi_i^T x')]\|_\infty \leq \bar{r} \|\phi_i\|_\infty |\phi_i^T(x - x')| \leq \bar{r} \|\phi_i\|_\infty^2 \|x - x'\|_1 \leq \bar{r} \bar{\nu}^2 \|x - x'\|_1,$$

so that $\nabla G(x, \omega) = \phi[\mathfrak{r}(\phi^T x) - \eta]$ is Lipschitz continuous w.r.t. ℓ_1 -norm with Lipschitz constant $\mathcal{L}(\omega) \leq \bar{r} \bar{\nu}^2$.

2°. Due to strong monotonicity of \mathfrak{r} ,

$$\begin{aligned} g(x) - g(x_*) &= \int_0^1 \nabla g(x_* + t(x - x_*))^T (x - x_*) dt \\ &= \int_0^1 \mathbf{E} \left\{ \phi[\mathfrak{r}(\phi^T(x_* + t(x - x_*))) - \mathfrak{r}(\phi^T x_*)] \right\}^T (x - x_*) dt \\ &\geq \int_0^1 \underline{r} \mathbf{E} \{ (\phi^T(x - x_*))^2 \} t dt = \frac{1}{2} \underline{r} \|x - x_*\|_\Sigma^2, \end{aligned}$$

what is (16).

3°. The sub-Gaussianity in the “batchless” case is readily given by $\nabla G(x_*, \omega_i) = \sigma \phi_i \xi_i$ with $\|\phi_i \xi_i\|_\infty \leq \|\phi_i\|_\infty \|\xi_i\| \leq \bar{\nu} \|\xi_i\|_2$ and

$$\mathbf{E} \left\{ \exp \left(\frac{\|\nabla G(x_*, \omega_i)\|_\infty^2}{\sigma^2 \bar{\nu}^2} \right) \right\} \leq e$$

due to $\mathbf{E}\{e^{\xi_i^2}\} \leq \exp(1)$. Because Θ variation of the d.-g.f. θ , as defined in (20), is bounded with $C \ln n$, by Lemma A.3 we conclude that batch observation

$$H(x_*, \omega_i^{(L)}) = \frac{1}{L} \sum_{\ell=1}^L \nabla G(x_*, \omega_i^\ell) = \frac{1}{L} \sum_{\ell=1}^L \sigma \phi_i^\ell \xi_i^\ell$$

is sub-Gaussian with parameter $\lesssim \sigma^2 \bar{\nu}^2 \ln n$.

4°. In the situation of Section 3.1, Σ is positive definite, $\Sigma \succeq \kappa_\Sigma I$, $\kappa_\Sigma > 0$, and condition $\mathbf{Q}(\lambda, \psi)$ is satisfied with $\lambda = \kappa_\Sigma$ and $\psi = 1$. Because quadratic minoration condition (17) for g is verified with $\mu \geq \underline{r}$ due to (16), when applying the result of Lemma 3.1, we conclude that Assumption [RSC] holds with $\delta = 1$ and $\rho = (\kappa_\Sigma \underline{r})^{-1}$.⁵ \square

B Properties of sparsity structures

B.1 Sparsity structures

The scope of results of Section 2 is much broader than “vanilla” sparsity optimization. We discuss here general notion of *sparsity structure* which provides a proper application framework for these results.

In what follows we assume to be given a *sparsity structure* [21] on E —a family \mathcal{P} of projector mappings $P = P^2$ on E such that

A1.1 every $P \in \mathcal{P}$ is assigned a linear map \bar{P} on E such that $P\bar{P} = 0$ and a nonnegative weight $\pi(P)$;

A1.2 whenever $P \in \mathcal{P}$ and $f, g \in E$ such that $\|f\|_* \leq 1$, $\|g\|_* \leq 1$,

$$\|P^*f + \bar{P}^*g\|_* \leq 1$$

where for a linear map $Q : E \rightarrow F$, $Q^* : F \rightarrow E$ is the conjugate mapping.

Following [21], we refer to a collection of the just introduced entities and *sparsity structure on E* . For a nonnegative real s we set

$$\mathcal{P}_s = \{P \in \mathcal{P} : \pi(P) \leq s\}.$$

Given $s \geq 0$ we call $x \in E$ *s-sparse* if there exists $P \in \mathcal{P}_s$ such that $Px = x$.

Typically, one is interested in the following “standard examples”:

1. “Vanilla (usual)” sparsity: in this case $E = \mathbf{R}^n$ with the standard inner product, \mathcal{P} is comprised of projectors on all coordinate subspaces of \mathbf{R}^n , $\pi(P) = \text{rank}(P)$, and $\|\cdot\| = \|\cdot\|_1$.
2. Group sparsity: $E = \mathbf{R}^n$, and we partition the set $\{1, \dots, n\}$ of indices into K nonoverlapping subsets I_1, \dots, I_K , so that to every $x \in \mathbf{R}^n$ we associate blocks x^k with corresponding indices in I_k , $k = 1, \dots, K$. Now \mathcal{P} is comprised of projectors $P = P_I$ onto subspaces $E_I = \{[x^1, \dots, x^K] \in \mathbf{R}^n : x^k = 0 \forall k \notin I\}$ associated with subsets I of the index set $\{1, \dots, K\}$. We set $\pi(P_I) = \text{card}I$, and define $\|x\| = \sum_{k=1}^K \|x^k\|_2$ —*block ℓ_1/ℓ_2 -norm*.
3. Low rank structure: in this example $E = \mathbf{R}^{p \times q}$ with, for the sake of definiteness, $p \geq q$, and the Frobenius inner product. Here \mathcal{P} is the set of mappings $P(x) = P_\ell x P_r$ where P_ℓ and P_r are, respectively, $q \times q$ and $p \times p$ orthoprojectors, $\bar{P}(x) = (I - P_\ell)x(I - P_r)$, and $\|\cdot\|$ is the nuclear norm $\|x\| = \sum_{i=1}^q \sigma_i(x)$ where $\sigma_1(x) \geq \sigma_2(x) \geq \dots \geq \sigma_q(x)$ are singular values of x , $\|\cdot\|_*$ is the spectral norm, so that $\|x\|_* = \sigma_1(x)$, and $\pi(P) = \max[\text{rank}(P_\ell), \text{rank}(P_r)]$.

In this case, for $\|f\|_* \leq 1$ and $\|g\|_* \leq 1$ one has

$$\|P^*(f)\|_* = \|P_\ell f P_r\|_* \leq 1, \quad \|\bar{P}^*(g)\|_* = \|(I - P_\ell)g(I - P_r)\|_* \leq 1,$$

and because the images and orthogonal complements to the kernels of P and \bar{P} are orthogonal to each other, $\|P^*(f) + \bar{P}^*(g)\|_* \leq 1$.

⁵We refer to Section B.2 and Lemma B.1 for the proof of Lemma 3.1.

B.2 Condition $\mathbf{Q}(\lambda, \psi)$

We say that a positive semidefinite mapping $\Sigma : E \rightarrow E$ satisfies condition $\mathbf{Q}(\lambda, \psi)$ for given $s \in \mathbf{Z}_+$ if for some $\psi, \lambda > 0$ and all $P \in \mathcal{P}_s$ and $z \in E$

$$\|Pz\| \leq \sqrt{s/\lambda}\|z\|_\Sigma + \|\bar{P}z\| - \psi\|z\|. \quad (46)$$

Lemma B.1 *Suppose that x_* is an optimal solution to (5) such that for some $P \in \mathcal{P}_s$, $\|(I - P)x_*\| \leq \delta$, and that condition $\mathbf{Q}(\lambda, \psi)$ is satisfied. Furthermore, assume that objective g of (5) satisfies the following minoration condition*

$$g(x) - g(x_*) \geq \mu(\|x - x_*\|_\Sigma)$$

where $\mu(\cdot)$ is monotone increasing and convex. Then a feasible solution $\hat{x} \in \mathcal{X}$ to (7) such that

$$\text{Prob}\{F_\kappa(\hat{x}) - F_\kappa(x_*) \leq v\} \geq 1 - \epsilon.$$

satisfies, with probability at least $1 - \epsilon$,

$$\|\hat{x} - x_*\| \leq \frac{\mu^*\left(\kappa\sqrt{s/\lambda}\right) + v}{\kappa\psi} + \frac{2\delta}{\psi} \quad (47)$$

where $\mu^* : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ is conjugate to $\mu(\cdot)$, $\mu^*(t) = \sup_{u \geq 0}[tu - \mu(u)]$.

Proof. When setting $z = \hat{x} - x_*$ one has

$$\begin{aligned} \hat{x} &= \|x_* + z\| = \|Px_* + (I - P)x_* + z\| \geq \|Px_* + z\| - \|(I - P)x_*\| \\ &\geq \|Px_*\| + \|\bar{P}z\| - \|Pz\| - \delta \end{aligned}$$

where we used the relation

$$\|Px_* + z\| \geq \|Px_*\| - \|Pz\| + \|\bar{P}z\|$$

(cf. Lemma 3.1 of [21] applied to $w = Px_*$). When using condition $\mathbf{Q}(\lambda, \psi)$ we obtain

$$\|\hat{x}\| \geq \|Px_*\| - \sqrt{s/\lambda}\|z\|_\Sigma + \psi\|z\| - \delta,$$

so that $F_\kappa(\hat{x}) \leq F_\kappa(x_*) + v$ implies

$$\begin{aligned} \kappa(\|Px_*\| + \psi\|z\| - \delta) &\leq \frac{1}{2}[g(x_*) - g(\hat{x})] + \kappa\sqrt{s/\lambda}\|z\|_\Sigma + \kappa\|x_*\| + v \\ &\leq -\frac{1}{2}\mu(\|z\|_\Sigma) + \kappa\sqrt{s/\lambda}\|z\|_\Sigma + \kappa\|x_*\| + v \\ &\leq \frac{1}{2}\mu^*(2\kappa\sqrt{s/\lambda}) + \kappa\|x_*\| + v, \end{aligned}$$

and we conclude that

$$\kappa\psi\|z\| \leq \frac{1}{2}\mu^*(2\kappa\sqrt{s/\lambda}) + 2\kappa\delta + v$$

due to $\|x_*\| - \|Px_*\| \leq \|(I - P)x_*\| \leq \delta$. □

Note that when $\mu(u) = \frac{\mu}{2}u^2$, one has $\mu^*(t) = \frac{1}{2\mu}t^2$, and in the case of $\|\cdot\| = \|\cdot\|_1$, with probability $1 - \epsilon$,

$$\|\hat{x} - x_*\|_1 \leq \frac{s\kappa}{\mu\lambda\psi} + \frac{v}{\kappa\psi} + \frac{2\delta}{\psi}.$$

This, in particular, implies bound (18) of Lemma 3.1.

Remark B.1 We discuss implications of condition $\mathbf{Q}(\lambda, \psi)$ and result of Lemma B.1 for “usual” sparsity in Section 3 of the paper. Now, let us consider the case of the low rank sparsity. Let $z \in \mathbf{R}^{p \times q}$ with $p \geq q$ for the sake of definiteness. In this case, $\|\cdot\|$ is the nuclear norm, and we put $P(z) = P_\ell z P_r$ where P_ℓ and P_r are orthoprojectors of rank $s \leq q$ such that $\|(I - P)(x)\| = \|x_* - P_\ell x_* P_r\| \leq \delta$.⁶

Furthermore, for a $p \times q$ matrix z let us put

$$\sigma^{(k)}(z) = \sum_{i=1}^k \sigma_i(z), \quad 1 \leq k \leq q.$$

With the sparsity parameter s being a nonnegative integer,

$$\forall (z \in \mathbf{R}^{p \times q}, P \in \mathcal{P}_s) : \quad \|P(z)\| \leq \sigma^{(s)}(z), \quad \|\bar{P}(z)\| \geq \|z\| - \sigma^{(2s)}(z).⁷$$

and we conclude that in the present situation condition

$$\sigma^{(s)}(z) + \sigma^{(2s)}(z) \leq \sqrt{s/\lambda} \|z\|_\Sigma + (1 - \psi) \|z\| \quad (48)$$

is sufficient for the validity of $\mathbf{Q}(\lambda, \psi)$. As a result, condition (48) with $\psi > 0$ is sufficient for applicability of the bound of Lemma B.1. It may also be compared to the necessary and sufficient condition of “ s -goodness of Σ ” in [43]:

$$\exists \psi > 0 : 2\sigma^{(s)}(z) \leq (1 - \psi) \|z\| \quad \forall z \in \text{Ker}(\Sigma).$$


C Supplementary numerical experiments

This section complements the numerical results appearing on the main body of the paper. We consider the setting in Section 3.3 of sparse recovery problem from GLR model observations (15). In the experiments below, we consider the choice (19) of activation function $\tau_\alpha(t)$ with values $\alpha = 1$ and $\alpha = 1/10$; value $\alpha = 1$ corresponds to linear regression with $\tau(t) = t$, whereas when $\alpha = 0.1$ activation have a flatter curve with rapidly decreasing with r modulus of strong convexity for $|t| \leq r$. Same as before, in our experiments, the dimension of the parameter space is $n = 500\,000$, the sparsity level of the optimal point x_* is $s = 100$; we use the minibatch Algorithm 2 with the maximal number of oracle calls is $N = 250\,000$. In Figures 4 and 5 we report results for $\kappa_\Sigma \in \{0.1, 1\}$ and $\sigma \in \{0.001, 0.1\}$; the simulations are repeated 10 times, we trace the median of the estimation error $\|\hat{x}_i - x_*\|_1$ along with its first and the last deciles against the number of oracle calls.


In our experiments, multistage algorithms exhibit linear convergence on initial iterations. Surprisingly, “standard” (non-Euclidean) SMD also converges fast in the “preliminary” regime. This may be explained by the fact that iteration x_i of the SMD obtained by the “usual” proximal mapping $\text{Prox}(\gamma_{i-1} \nabla G(x_{i-1}, \omega_i), x_{i-1})$ is computed as a solution to the optimization problem with “penalty” $\theta(x) = c \|x\|_p^p$, $p = 1 + 1/\ln n$ which results in a “natural” sparsification of x_i . As iterations progress, such “sparsification” becomes insufficient, and the multistage routine eventually outperforms the SMD. Implementing the method for “flatter” nonlinear activation $\tau(t)$ or increased condition number of the regressor covariance matrix Σ requires increasing the length m_0 of the stage of the algorithm.

⁶E.g., choose P_ℓ and P_r as left and right projectors on the space generated by s principal left and right singular vectors of x_* , so that $\|x_* - P_\ell x_* P_r\| = \|(I - P_\ell)x_*(I - P_r)\| = \sum_{i=s+1}^q \sigma_i \leq \delta$.


⁷Indeed, let $P \in \mathcal{P}_s$, so that $\text{rank}(P_\ell) \leq s$ and $\text{rank}(P_r) \leq s$, and $\|P(z)\| = \|P_\ell z P_r\| \leq \sigma^{(s)}(z)$. Since the matrix $\bar{P}(z)$ differs from z by a matrix of rank at most $2s$, by the Singular Value Interlacing theorem we have $\sigma_i(\bar{P}(z)) \geq \sigma_{i+2s}(z)$, whence $\|\bar{P}(z)\| \geq \|z\| - \sigma^{(2s)}(z)$.



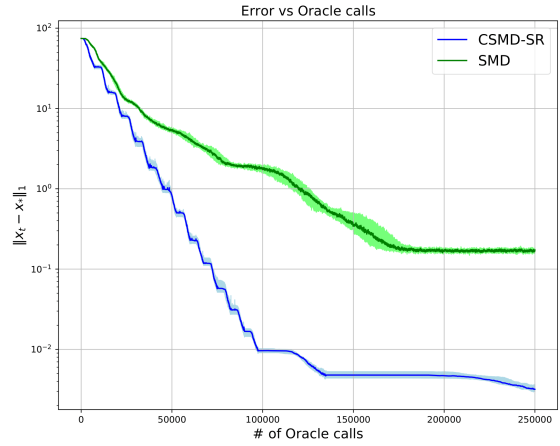
(a) $\kappa_\Sigma = 1, \sigma = 0.1, m_0 = 5000$



(b) $\kappa_\Sigma = 1, \sigma = 0.001, m_0 = 5000$




(c) $\kappa_\Sigma = 0.1, \sigma = 0.1, m_0 = 7500$




(d) $\kappa_\Sigma = 0.1, \sigma = 0.001, m_0 = 7500$


Figure 4: CSMD-SR and “vanilla” SMD in Linear Regression problem (activation function $\tau(t) = t$); ℓ_1 error as a function of the number of oracle calls




(a) $\kappa_{\Sigma} = 1, \sigma = 0.1, m_0 = 8000$



(b) $\kappa_{\Sigma} = 1, \sigma = 0.001, m_0 = 8000$



(c) $\kappa_{\Sigma} = 0.1, \sigma = 0.1, m_0 = 10000$



(d) $\kappa_{\Sigma} = 0.1, \sigma = 0.001, m_0 = 10000$

Figure 5: CSMD-SR and “vanilla” SMD in Generalized Linear Regression problem: activation function $\tau_{1/10}(t)$; ℓ_1 error as a function of the number of oracle calls