



HAL
open science

Stochastic Variance Reduction Methods for Saddle-Point Problems

P Balamurugan, Francis Bach

► **To cite this version:**

P Balamurugan, Francis Bach. Stochastic Variance Reduction Methods for Saddle-Point Problems. 2016. hal-01319293v1

HAL Id: hal-01319293

<https://hal.science/hal-01319293v1>

Preprint submitted on 20 May 2016 (v1), last revised 2 Nov 2016 (v2)

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 Variance Reduction Methods for Saddle-Point Problems

P. Balamurugan and Francis Bach
INRIA - Sierra project-team
Département d'Informatique de l'École Normale Supérieure
Paris, France

May 20, 2016

Abstract

We consider convex-concave saddle-point problems where the objective functions may be split in many components, and extend recent stochastic variance reduction methods (such as SVRG or SAGA) to provide the first large-scale linearly convergent algorithms for this class of problems which is common in machine learning. While the algorithmic extension is straightforward, it comes with challenges and opportunities: (a) the convex minimization analysis does not apply and we use the notion of monotone operators to prove convergence, showing in particular that the same algorithm applies to a larger class of problems, such as variational inequalities, (b) there are two notions of splits, in terms of functions, or in terms of partial derivatives, (c) the split does need to be done with convex-concave terms, (d) non-uniform sampling is key to an efficient algorithm, both in theory and practice, and (e) these incremental algorithms can be easily accelerated using a simple extension of the “catalyst” framework, leading to an algorithm which is always superior to accelerated batch algorithms.

1 Introduction

When using optimization in machine learning, leveraging the natural separability of the objective functions has led to many algorithmic advances; the most common example is the separability as a sum of individual loss terms corresponding to individual observations, which leads to stochastic gradient descent techniques. Several lines of work have shown that the plain Robbins-Monro algorithm could be accelerated for strongly-convex finite sums, e.g., SAG [1], SDCA [2], SVRG [3], SAGA [4]. However, these only apply to separable objective functions.

In order to tackle non-separable losses or regularizers, we consider the saddle-point problem:

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} K(x, y) + M(x, y), \quad (1)$$

where the functions K and M are “convex-concave”, that is, convex with respect to the first variable, and concave with respect to the second variable, with M potentially non-smooth but “simple” (e.g., for which the proximal operator is easy to compute), and K smooth. These problems occur naturally within convex optimization through Lagrange or Fenchel duality [5]; for example the bilinear saddle-point problem $\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} f(x) + y^\top Kx - g(y)$ corresponds to a supervised learning problem with design matrix K , a loss function g^* (the Fenchel conjugate of g) and a regularizer f .

We assume that the function K may be split into a potentially large number of components. Many problems in machine learning exhibit that structure in the saddle-point formulation, but not in the associated convex minimization and concave maximization problems (see examples in Section 2.1).

Like for convex minimization, gradient-based techniques that are blind to this separable structure need to access all the components at every iteration. We show that algorithms such as SVRG [3] and SAGA [4] may be readily extended to the saddle-point problem. While the algorithmic extension is straightforward, it comes with challenges and opportunities. We make the following contributions:

- We provide the first convergence analysis for these algorithms for saddle-point problems, which differs significantly from the associated convex minimization set-up. In particular, we use in Section 6 the interpretation of saddle-point problems as finding the zeros of a monotone operator, and only use the monotonicity properties to show linear convergence of our algorithms, thus showing that they extend beyond saddle-point problems, e.g., to variational inequalities [6, 7].
- We show that the saddle-point formulation (a) allows two different notions of splits, in terms of functions, or in terms of partial derivatives, (b) does need splits into convex-concave terms (as opposed to convex minimization), and (c) that non-uniform sampling is key to an efficient algorithm, both in theory and practice (see experiments in Section 7).
- We show in Section 5 that these incremental algorithms can be easily accelerated using a simple extension of the “catalyst” framework of [8], thus leading to an algorithm which is always superior to accelerated batch algorithms.

2 Composite Decomposable Saddle-Point Problems

We consider the saddle-point problem defined in Eq. (1), with the following assumptions:

- (A) M is strongly (λ, γ) -convex-concave, that is, the function $(x, y) \mapsto M(x, y) - \frac{\lambda}{2}\|x\|^2 + \frac{\gamma}{2}\|y\|^2$ is convex-concave. Moreover, we assume that we may compute the proximal operator of M , i.e., for any $(x', y') \in \mathbb{R}^{n+d}$:

$$\text{prox}_M^\sigma(x', y') = \arg \min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} \sigma M(x, y) + \frac{\lambda}{2}\|x - x'\|^2 - \frac{\gamma}{2}\|y - y'\|^2. \quad (2)$$

The values of λ and γ lead to the definition of a weighted Euclidean norm on \mathbb{R}^{n+d} defined as $\Omega(x, y)^2 = \lambda\|x\|^2 + \gamma\|y\|^2$, with dual norm defined through $\Omega^*(x, y)^2 = \lambda^{-1}\|x\|^2 + \gamma^{-1}\|y\|^2$. Dealing with the two different scaling factors λ and γ is crucial for good performance, as these may be very different, depending on the many arbitrary ways to set-up a saddle-point problem.

- (B) K is convex-concave and has Lipschitz-continuous gradients; it is natural to quantify this by considering the *gradient operator* $B : \mathbb{R}^{n+d} \rightarrow \mathbb{R}^{n+d}$ defined as $B(x, y) = (\partial_x K(x, y), -\partial_y K(x, y)) \in \mathbb{R}^{n+d}$ and to consider $L = \sup_{\Omega(x-x', y-y')=1} \Omega^*(B(x, y) - B(x', y'))$. The quantity L represents the condition number of the problem.
- (C) The vector-valued function $B(x, y) = (\partial_x K(x, y), -\partial_y K(x, y)) \in \mathbb{R}^{n+d}$ may be split into a family of vector-valued functions as $B = \sum_{i \in \mathcal{J}} B_i$, where the only constraint is that each B_i is Lipschitz-continuous (with constant L_i). There is no need to assume the existence of $K_i : \mathbb{R}^{n+d} \rightarrow \mathbb{R}$ such that $B_i = (\partial_x K_i, -\partial_y K_i)$.

We will also consider splits which are adapted to the saddle-point nature of the problem, that is, of the form $B(x, y) = (\sum_{k \in \mathcal{K}} B_k^x(x, y), \sum_{j \in \mathcal{J}} B_j^y(x, y))$, which is a subcase of the above with $\mathcal{J} = \mathcal{J} \times \mathcal{K}$, $B_{jk}(x, y) = (p_j B_k^x(x, y), q_k B_j^y(x, y))$, for p and q sequences that sum to one. This substructure, which we refer to as “factored”, will only make a difference when storing the values of these operators in Section 4 for our SAGA algorithm.

Given assumptions (A)-(B), the saddle-point problem in Eq. (1) has a unique solution (x_*, y_*) such that $K(x_*, y) + M(x_*, y) \leq K(x_*, y_*) + M(x_*, y_*) \leq K(x, y_*) + M(x, y_*)$, for all (x, y) ; moreover $\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} K(x, y) + M(x, y) = \max_{y \in \mathbb{R}^n} \min_{x \in \mathbb{R}^d} K(x, y) + M(x, y)$ (see, e.g., [9, 5]).

The main generic examples for the functions $K(x, y)$ and $M(x, y)$ are:

- **Bilinear saddle-point problems:** $K(x, y) = y^\top Kx$ for a matrix¹ $K \in \mathbb{R}^{n \times d}$, for which the vector-valued function $B(x, y)$ is linear, i.e., $B(x, y) = (K^\top y, -Kx)$. Then $L = \|K\|_{\text{op}}/\sqrt{\gamma\lambda}$, where $\|K\|_{\text{op}}$ is the largest singular value of K .

There are two natural potential splits with $\mathcal{J} = \{1, \dots, n\} \times \{1, \dots, k\}$, with $B = \sum_{j=1}^n \sum_{k=1}^d B_{jk}$: (a) the split into individual elements $B_{jk}(x, y) = K_{jk}(y_j, -x_k)$, where every element is the gradient operator of a bi-linear function, and (b) the “factored” split into rows/columns $B_{jk}(x, y) = (q_k y_j K_{j.}^\top, -p_j x_k K_{.k})$, where $K_{j.}$ and $K_{.k}$ are the j -th row and k -th column of K , p and q are any set of vectors summing to one, and every element is not the gradient operator of any function. These splits correspond to several “sketches” of the matrix K [10], adapted to subsampling of K , but other sketches could be considered.

- **Separable functions:** $M(x, y) = f(x) - g(y)$ where f is any λ -strongly-convex and g is γ -strongly convex, for which the proximal operators $\text{prox}_f^\sigma(x') = \arg \min_{x \in \mathbb{R}^d} \sigma f(x) + \frac{\lambda}{2} \|x - x'\|^2$ and $\text{prox}_g^\sigma(y') = \arg \max_{y \in \mathbb{R}^d} -\sigma g(y) - \frac{\gamma}{2} \|y - y'\|^2$ are easy to compute. In this situation, $\text{prox}_M^\sigma(x', y') = (\text{prox}_f^\sigma(x'), \text{prox}_g^\sigma(y'))$. Following the usual set-up of composite optimization [11], no smoothness assumption is made on M and hence on f or g .

2.1 Examples in machine learning

Many learning problems are formulated as convex optimization problems, and hence by duality as saddle-point problems. We now give examples where our new algorithms are particularly adapted.

Supervised learning with non-separable losses or regularizers. For regularized linear supervised learning, with n d -dimensional observations put in a design matrix $K \in \mathbb{R}^{n \times d}$, the predictions are parameterized by a vector $x \in \mathbb{R}^d$ and lead to a vector of predictions $Kx \in \mathbb{R}^n$. Given a loss function defined through its Fenchel conjugate g^* from \mathbb{R}^n to \mathbb{R} , and a regularizer $f(x)$, we obtain exactly a bi-linear saddle-point problem. When the loss g^* or the regularizer f is separable, i.e., a sum of functions of individual variables, we may apply existing fast gradient-techniques [1, 3, 4] to the primal problem $\min_{x \in \mathbb{R}^d} g^*(Kx) + f(x)$ or the dual problem $\max_{y \in \mathbb{R}^n} -g(y) - f^*(K^\top y)$, as well as methods dedicated to separable saddle-point problems [12]. When the loss g^* and the regularizer f are not separable (but have a simple proximal operator), our new fast algorithms are the only ones that can be applied from the class of large-scale linearly convergent algorithms.

Non-separable *losses* may occur when (a) predicting by affine functions of the inputs and not penalizing the constant terms (in this case defining the loss functions as the minimum over the constant term, which becomes non-separable) or (b) using structured output prediction methods that lead to convex surrogates to the area under the ROC curve (AUC) or other precision/recall quantities [13, 14, 15]. These come often with efficient proximal operators (see Section 7 for an example).

Non-separable *regularizers* with available efficient proximal operators are numerous, such as grouped-norms with potentially overlapping groups, norms based on submodular functions, or total variation (see [16] and references therein, and an example in Section 7).

Robust optimization. The framework of robust optimization [17] aims at optimizing an objective function with uncertain data. Given that the aim is then to minimize the maximal value of the objective function given the uncertainty, this leads naturally to saddle-point problems.

¹We identify here a matrix with the associated bilinear function.

Convex relaxation of unsupervised learning. Unsupervised learning leads to convex relaxations which often exhibit structures naturally amenable to saddle-point problems, e.g, for discriminative clustering [18] or matrix factorization [19].

2.2 Existing batch algorithms

In this section, we review existing algorithms aimed at solving the composite saddle-point problem in Eq. (1), without using the sum-structure. Note that it is often possible to apply batch algorithms for the associated primal or dual problems (which are not separable in general).

Forward-backward (FB) algorithm. The main iteration is

$$\begin{aligned} (x_t, y_t) &= \text{prox}_M^\sigma \left[(x_{t-1}, y_{t-1}) - \sigma \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} B(x_{t-1}, y_{t-1}) \right] \\ &= \text{prox}_M^\sigma (x_{t-1} - \sigma \lambda^{-1} \partial_x K(x_{t-1}, y_{t-1}), y_{t-1} + \sigma \gamma^{-1} \partial_y K(x_{t-1}, y_{t-1})). \end{aligned}$$

Intuitively, the algorithm aims at simultaneously minimizing with respect to x while maximizing with respect to y (when $M(x, y)$ is the sum of isotropic quadratic terms and indicator functions, we get simultaneous projected gradient descent). This algorithm is known not to converge in general [9], but is linearly convergent for *strongly*-convex-concave problems, when $\sigma = 1/L^2$, with the rate $(1 - 1/(1 + L^2))^t$ [20] (see simple proof in Appendix B.1). This is the one we are going to adapt to stochastic variance reduction.

When $M(x, y) = f(x) - g(y)$, we obtain the two parallel updates $x_t = \text{prox}_f^\sigma (x_{t-1} - \lambda^{-1} \sigma \partial_x K(x_{t-1}, y_{t-1}))$ and $y_t = \text{prox}_g^\sigma (y_{t-1} + \gamma^{-1} \sigma \partial_y K(x_{t-1}, y_{t-1}))$, which can be done serially by replacing the second one by $y_t = \text{prox}_g^\sigma (y_{t-1} + \gamma^{-1} \sigma \partial_y K(x_t, y_{t-1}))$. This is often referred to as the Arrow-Hurwicz method (see [21] and references therein).

Extragradient algorithm. When M is not assumed to be strongly convex-concave, the forward-backward algorithm may not be convergent. The extragradient algorithm performs two updates at each iteration (see, e.g., [22, 23]), and is then convergent. It is convergent for our problem with a similar worst-case convergence rate [23], but is more robust to step-size selection.

Accelerated forward-backward algorithm. The forward-backward algorithm may be accelerated by a simple extrapolation step, similar to Nesterov’s acceleration for convex minimization [24]. The algorithm from [21], which *only applies to bilinear functions* K , and which we extend from separable M to our more general set-up in Appendix B.2, has the following iteration:

$$(x_t, y_t) = \text{prox}_M^\sigma \left[(x_{t-1}, y_{t-1}) - \sigma \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} B(x_{t-1} + \theta(x_{t-1} - x_{t-2}), y_{t-1} + \theta(y_{t-1} - y_{t-2})) \right].$$

With $\sigma = 1/(2L)$ and $\theta = L/(L + 1)$, we get an improved convergence rate, where $(1 - 1/(1 + L^2))^t$ is replaced by $(1 - 1/(1 + 2L))^t$. This is always a strong improvement when L is large (ill-conditioned problems), as illustrated in Section 7. Note that our acceleration technique in Section 5 may be extended to get a similar rate for the batch set-up for non-linear K .

2.3 Existing stochastic algorithms

Forward-backward algorithms have been studied with added noise [25], leading to a convergence rate in $O(1/t)$ after t iterations for strongly-convex-concave problems. In our setting, we replace $B(x, y)$ in our algorithm with $\frac{1}{\pi_i} B_i(x, y)$, where $i \in \mathcal{J}$ is sampled from the probability vector $(\pi_i)_i$. We have $\mathbb{E} B_i(x, y) = B(x, y)$; the main iteration is then

$$(x_t, y_t) = \text{prox}_M^{\sigma_t} \left[(x_{t-1}, y_{t-1}) - \sigma_t \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} \frac{1}{\pi_{i_t}} B_{i_t}(x_{t-1}, y_{t-1}) \right],$$

with i_t selected independently at random in \mathcal{J} with probability vector π . In Appendix C, we show that using $\sigma_t = 2/(t+1+8\bar{L}(\pi)^2)$ leads to a convergence rate in $O(1/t)$, where $\bar{L}(\pi)$ is a smoothness constant explicated below. For saddle-point problems, it leads to the complexities shown in Table 1. Like for convex minimization, it is fast early on but the performance levels off. Such schemes are typically used in sublinear algorithms [26].

2.4 Sampling probabilities, convergence rates and running-time complexities

In order to characterize running-times, we denote by $T(A)$ the complexity of computing $A(x, y)$ for any operator A and $(x, y) \in \mathbb{R}^{n+d}$, while we denote by $T_{\text{prox}}(M)$ the complexity of computing $\text{prox}_M^\sigma(x, y)$. In our motivating example of bilinear functions $K(x, y)$, we assume that $T_{\text{prox}}(M)$ takes times proportional to $n+d$ and getting a single element of K is $O(1)$.

In order to characterize the convergence rate, we need the Lipschitz-constant L defined earlier as well as a smoothness constant adapted to our sampling schemes:

$$\bar{L}(\pi)^2 = \sup_{(x,y,x',y')} \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \Omega^*(B_i(x, y) - B_i(x', y'))^2 \text{ such that } \Omega(x - x', y - y')^2 \leq 1.$$

We always have the bounds $L^2 \leq \bar{L}(\pi)^2 \leq \max_{i \in \mathcal{J}} L_i^2 \times \sum_{i \in \mathcal{J}} \frac{1}{\pi_i}$. However, in structured situations (like in bilinear saddle-point problems), we get much improved bounds, as described below.

Bi-linear saddle-point. The constant L is equal to $\|K\|_{\text{op}}/\sqrt{\lambda\gamma}$, and we will consider as well the Frobenius norm $\|K\|_F$ defined through $\|K\|_F^2 = \sum_{j,k} K_{jk}^2$, and the norm $\|K\|_{\text{max}}$ defined as $\|K\|_{\text{max}}^2 = \max\{\sup_j (KK^\top)_{jj}^{1/2}, \sup_k (K^\top K)_{kk}^{1/2}\}$. Among the norms above, we always have:

$$\|K\|_{\text{max}} \leq \|K\|_{\text{op}} \leq \|K\|_F \leq \sqrt{\max\{n, d\}} \|K\|_{\text{max}} \leq \sqrt{\max\{n, d\}} \|K\|_{\text{op}}, \quad (3)$$

which allows to show below that some algorithms have better bounds than others.

There are several schemes to choose the probabilities π_{jk} (individual splits) and $\pi_{jk} = p_j q_k$ (factored splits). For the factored formulation where we select random rows and columns, we consider the non-uniform schemes $p_j = (KK^\top)_{jj}/\|K\|_F^2$ and $q_k = (K^\top K)_{kk}/\|K\|_F^2$, leading to $\bar{L}(\pi) \leq \|K\|_F/\sqrt{\lambda\gamma}$, or uniform, leading to $\bar{L}(\pi) \leq \sqrt{\max\{n, d\}} \|K\|_{\text{max}}/\sqrt{\lambda\gamma}$. For the individual formulation where we select random elements, we consider $\pi_{jk} = K_{jk}^2/\|K\|_F^2$, leading to $\bar{L}(\pi) \leq \sqrt{\max\{n, d\}} \|K\|_F/\sqrt{\lambda\gamma}$, or uniform, leading to $\bar{L}(\pi) \leq \sqrt{nd} \|K\|_{\text{max}}/\sqrt{\lambda\gamma}$ (in these situations, it is important to select several elements simultaneously, which our analysis supports).

We characterize convergence with the quantity $\varepsilon = \Omega(x - x_*, y - y_*)^2/\Omega(x_0 - x_*, y_0 - y_*)^2$, where (x_0, y_0) is the initialization of our algorithms (typically $(0, 0)$ for bilinear saddle-points). In Table 1 we give a summary of the complexity of all algorithms discussed in this paper: we recover the same type of speed-ups as for convex minimization. A few points are worth mentioning:

- Given the bounds between the various norms on K in Eq. (3), SAGA/SVRG with non-uniform sampling always has convergence bounds superior to SAGA/SVRG with uniform sampling, which is always superior to batch forward-backward. Note however, that in practice, SAGA/SVRG with uniform sampling may be inferior to accelerated batch method (see Section 7).
- Accelerated SVRG with non-uniform sampling is the most efficient method, which is confirmed in our experiments.

Algorithms	Complexity
Batch FB	$\log(1/\varepsilon) \times (nd + nd\ K\ _{\text{op}}^2/(\lambda\gamma))$
Batch FB-accelerated	$\log(1/\varepsilon) \times (nd + nd\ K\ _{\text{op}}/\sqrt{\lambda\gamma})$
Stochastic FB-non-uniform	$(1/\varepsilon) \times (\max\{n, d\}\ K\ _F^2/(\lambda\gamma))$
Stochastic FB-uniform	$(1/\varepsilon) \times (nd\ K\ _{\text{max}}^2/(\lambda\gamma))$
SAGA/SVRG-uniform	$\log(1/\varepsilon) \times (nd + nd\ K\ _{\text{max}}^2/(\lambda\gamma))$
SAGA/SVRG-non-uniform	$\log(1/\varepsilon) \times (nd + \max\{n, d\}\ K\ _F^2/(\lambda\gamma))$
SVRG-non-uniform-accelerated	$\log(1/\varepsilon) \times (nd + \sqrt{nd} \max\{n, d\}\ K\ _F/\sqrt{\lambda\gamma})$

Table 1: Summary of convergence results for the strongly (λ, γ) -convex-concave bilinear saddle-point problem with matrix K and individual splits (and $n + d$ updates per iteration). For factored splits (little difference), see Appendix D.4. For accelerated SVRG, we omitted the logarithmic term (see Section 5).

3 SVRG: Stochastic Variance Reduction for Saddle Points

Following [3, 27], we consider a stochastic-variance reduced estimation of the finite sum $B(x, y) = \sum_{i \in \mathcal{J}} B_i(x, y)$. This is achieved by assuming that we have an iterate (\tilde{x}, \tilde{y}) with a known value of $B(\tilde{x}, \tilde{y})$, and we consider the estimate of $B(x, y)$:

$$B(\tilde{x}, \tilde{y}) + \frac{1}{\pi_i} B_i(x, y) - \frac{1}{\pi_i} B_i(\tilde{x}, \tilde{y}),$$

which has the correct expectation when i is sampled from \mathcal{J} with probability π , but with a reduced variance. Since we need to refresh (\tilde{x}, \tilde{y}) regularly, the algorithm works in epochs (we allow to sample m elements per updates, i.e., a mini-batch of size m), with an algorithm that shares the same structure than SVRG for convex minimization; see Algorithm 1. Note that we provide an explicit number of iterations per epoch, proportional to $(L^2 + 3\bar{L}^2/m)$. We have the following theorem, shown in Appendix D.1 (see also a dicussion of the proof in Section 6).

Theorem 1 *Assume (A)-(B)-(C). After v epochs of Algorithm 2, we have:*

$$\mathbb{E}[\Omega(x_v - x_*, y_v - y_*)^2] \leq (3/4)^v \Omega(x_0 - x_*, y_0 - y_*)^2.$$

The computational complexity to reach precision ε is proportional to $[T(B) + (mL^2 + \bar{L}^2) \max_{i \in \mathcal{J}} T(B_i) + (1 + L^2 + \bar{L}^2/m) T_{\text{prox}}(M)] \log \frac{1}{\varepsilon}$. Note that by taking m large, we can alleviate the complexity of the proximal operator prox_M if too large. Moreover, if L^2 is too expensive to compute, we may replace it by \bar{L}^2 but with a worse complexity bound.

Bilinear saddle-point problems. When using a mini-batch size $m = 1$ with the factored updates, or $m = n + d$ for the individual updates, we get the same complexities proportional to $[nd + \max\{n, d\}\|K\|_F^2/(\lambda\gamma)] \log(1/\varepsilon)$ for non-uniform sampling, which improve significantly over (non-accelerated) batch methods (see Table 1).

4 SAGA: Online Stochastic Variance Reduction for Saddle Points

Following [4], we consider a stochastic-variance reduced estimation of $B(x, y) = \sum_{i \in \mathcal{J}} B_i(x, y)$. This is achieved by assuming that we store values $g^i = B_i(x^{\text{old}(i)}, y^{\text{old}(i)})$ for an old iterate $(x^{\text{old}(i)}, y^{\text{old}(i)})$,

Algorithm 1 SVRG: Stochastic Variance Reduction for Saddle Points

Input: Functions $(K_i)_i$, M , probabilities $(\pi_i)_i$, smoothness $\bar{L}(\pi)$ and L , iterate (x, y)
number of epochs v , number of updates per iteration (mini-batch size) m

Set $\sigma = [L^2 + 3\bar{L}^2/m]^{-1}$

for $u = 1$ to v **do**

Initialize $(\tilde{x}, \tilde{y}) = (x, y)$ and compute $B(\tilde{x}, \tilde{y})$

for $k = 1$ to $\log 4 \times (L^2 + 3\bar{L}^2/m)$ **do**

Sample $i_1, \dots, i_m \in \mathcal{J}$ from the probability vector $(\pi_i)_i$ with replacement

$$(x, y) \leftarrow \text{prox}_M^\sigma \left[(x, y) - \sigma \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} (B(\tilde{x}, \tilde{y}) + \frac{1}{m} \sum_{k=1}^m \{ \frac{1}{\pi_{i_k}} B_{i_k}(x, y) - \frac{1}{\pi_{i_k}} B_{i_k}(\tilde{x}, \tilde{y}) \}) \right]$$

end for

end for

Output: Approximate solution (x, y)

and we consider the estimate of $B(x, y)$:

$$\sum_{j \in \mathcal{J}} g^j + \frac{1}{\pi_i} B_i(x, y) - \frac{1}{\pi_i} g^i,$$

which has the correct expectation when i is sampled from \mathcal{J} with probability π . At every iteration, we also refresh the operator values $g^i \in \mathbb{R}^{n+d}$, for the same index i or with a new index i sampled uniformly at random. This leads to Algorithm 2, and we have the following theorem showing linear convergence, proved in Appendix D.2. Note that for bi-linear saddle-points, the initialization at $(0, 0)$ has zero cost (which is not possible for convex minimization).

Theorem 2 *Assume (A)-(B)-(C). After t iterations of Algorithm 2 (with the option of resampling when using non-uniform sampling), we have:*

$$\mathbb{E}[\Omega(x_t - x_*, y_t - y_*)^2] \leq 2(1 - (\max\{\frac{3|\mathcal{J}|}{2m}, 1 + \frac{L^2}{\mu^2} + 3\frac{\bar{L}^2}{m\mu^2}\})^{-1})^t \Omega(x_0 - x_*, y_0 - y_*)^2.$$

Resampling or re-using the same gradients. For the bound above to be valid for non-uniform sampling, like for convex minimization [28], we need to resample m operators after we make the iterate update. In our experiments, following [28], we considered a mixture of uniform and non-uniform sampling, without a resampling step.

SAGA vs. SVRG. The difference between the two algorithms is the same as for convex minimization (see, e.g., [29] and references therein), that is SVRG has no storage, but works in epochs and requires slightly more accesses to the oracles, while SAGA is a pure online method with fewer parameters but requires some storage (for bi-linear saddle-point problems, we only need to store $O(n+d)$ elements for the factored splits, while we need $O(dn)$ for the individual splits). Overall they have the same running-time complexity for individual splits; for factored splits, see Appendix D.4.

Factored splits. When using factored splits, we need to store the two parts of the operator values separately and update them independently, leading in Theorem 2 to replacing $|\mathcal{J}|$ by $\max\{|\mathcal{J}|, |\mathcal{K}|\}$.

5 Acceleration

Following the “catalyst” framework of [8], we consider a sequence of saddle-point problems with added regularization; namely, given (\bar{x}, \bar{y}) , we use SVRG to solve approximately

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} K(x, y) + M(x, y) + \frac{\lambda\tau}{2} \|x - \bar{x}\|^2 - \frac{\gamma\tau}{2} \|y - \bar{y}\|^2, \quad (4)$$

Algorithm 2 SAGA: Online Stochastic Variance Reduction for Saddle Points

Input: Functions $(K_i)_i$, M , probabilities $(\pi_i)_i$, smoothness $\bar{L}(\pi)$ and L , iterate (x, y)
number of iterations t , number of updates per iteration (mini-batch size) m

Set $\sigma = [\max\{\frac{3|\mathcal{J}|}{2m} - 1, L^2 + 3\frac{\bar{L}^2}{m}\}]^{-1}$

Initialize $g^i = B_i(x, y)$ for all $i \in \mathcal{J}$ and $G = \sum_{i \in \mathcal{J}} g^i$

for $u = 1$ to t **do**

Sample $i_1, \dots, i_m \in \mathcal{J}$ from the probability vector $(\pi_i)_i$ with replacement

Compute $h_k = B_{i_k}(x, y)$ for $k \in \{1, \dots, m\}$

$(x, y) \leftarrow \text{prox}_M^\sigma [(x, y) - \sigma \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} (G + \frac{1}{m} \sum_{k=1}^m \{\frac{1}{\pi_{i_k}} h_k - \frac{1}{\pi_{i_k}} g^{i_k}\})]$

(optional) Sample $i_1, \dots, i_m \in \mathcal{J}$ uniformly with replacement

(optional) Compute $h_k = B_{i_k}(x, y)$ for $k \in \{1, \dots, m\}$

Replace $G \leftarrow G - \sum_{k=1}^m \{g^{i_k} - h_k\}$ and $g^{i_k} \leftarrow h_k$ for $k \in \{1, \dots, m\}$

end for

Output: Approximate solution (x, y)

for well-chosen τ and (\bar{x}, \bar{y}) . The main iteration of the algorithm differs from the original SVRG by the presence of the iterate (\bar{x}, \bar{y}) and different step-sizes (see details in Appendix D.3). The complexity to get an approximate solution of Eq. (4) (forgetting the complexity of the proximal operator and for a single update), up to logarithmic terms, is proportional, to $T(B) + \bar{L}^2(1 + \tau)^{-2} \max_{i \in \mathcal{J}} T(B_i)$.

The key difference with the convex optimization set-up is that the analysis is simpler, without the need for Nesterov acceleration machinery [24] to define a good value of (\bar{x}, \bar{y}) ; indeed, the solution of Eq. (4) is one iteration of the proximal-point algorithm, which is known to converge linearly [30] with rate $(1 + \tau^{-1})^{-1} = (1 - \frac{1}{1 + \tau})$. Thus the overall complexity is up to logarithmic terms equal to $T(B)(1 + \tau) + \bar{L}^2(1 + \tau)^{-1} \max_{i \in \mathcal{J}} T(B_i)$. The trade-off in τ is optimal for $1 + \tau = \bar{L} \sqrt{\max_{i \in \mathcal{J}} T(B_i) / T(B)}$, showing that there is a potential acceleration when $\bar{L} \sqrt{\max_{i \in \mathcal{J}} T(B_i) / T(B)} \geq 1$, leading to a complexity $\bar{L} \sqrt{T(B) \max_{i \in \mathcal{J}} T(B_i)}$.

Since the SVRG algorithm already works in epochs, this leads to a simple modification where every $\log(1 + \tau)$ epochs, we change the values of (\bar{x}, \bar{y}) . See Algorithm 3 in Appendix D.3. Moreover, we can adaptively update (\bar{x}, \bar{y}) more aggressively to speed-up the algorithm.

The following theorem gives the convergence rate of the method (see proof in Appendix D.3). With the value of τ defined above (corresponding to $\tau = \max\{0, \frac{\|K\|_F}{\sqrt{\lambda\gamma}} \sqrt{\max\{n^{-1}, d^{-1}\}} - 1\}$ for bilinear problems), we get the complexity $\bar{L} \sqrt{T(B) \max_{i \in \mathcal{J}} T(B_i)}$, up to the logarithmic term $\log(1 + \tau)$. For bilinear problems, this provides a significant acceleration, as shown in Table 1.

Theorem 3 Assume (A)-(B)-(C). After v epochs of Algorithm 3, we have, for any positive τ :

$$\mathbb{E}[\Omega(x_v - x_*, y_v - y_*)^2] \leq (1 - \frac{1}{\tau + 1})^v \Omega(x_0 - x_*, y_0 - y_*)^2.$$

While we provide a proof only for SVRG, the same scheme should work for SAGA. Moreover, the same idea also applies to the batch setting (by simply considering $|\mathcal{J}| = 1$, i.e., a single function), leading to an acceleration, but now valid for all functions K (not only bilinear).

6 Extension to Monotone Operators

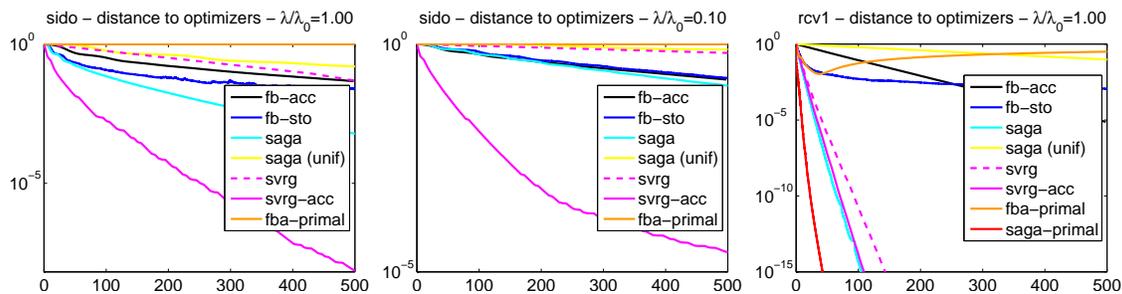
In this paper, we have chosen to focus on saddle-point problems because of their ubiquity in machine learning. However, it turns out that our algorithm, and, more importantly, our analysis extend to all set-valued *monotone operators* [9, 31]. We thus consider a maximal strongly-monotone operator A on a Euclidean space \mathcal{E} , as well as a finite family of Lipschitz-continuous (not necessarily monotone) operators B_i , $i \in \mathcal{J}$, with $B = \sum_{i \in \mathcal{J}} B_i$ monotone. Our algorithm then finds the zeros of $A + \sum_{i \in \mathcal{J}} B_i = A + B$, from the knowledge of the resolvent (“backward”) operator $(I + \sigma A)^{-1}$ (for a well chosen $\sigma > 0$) and the forward operators B_i , $i \in \mathcal{J}$. There several interesting examples (on which our algorithms apply):

- **Saddle-point problems:** We assume for simplicity that $\lambda = \gamma = \mu$ (this can be achieved by a simple change of variable). If we denote $B(x, y) = (\partial_x K(x, y), -\partial_y K(x, y))$ and the multi-valued operator $A(x, y) = (\partial_x M(x, y), -\partial_y M(x, y))$, then the proximal operator prox_M^σ may be written as $(\mu I + \sigma A)^{-1}(\mu x, \mu y)$, and we recover exactly our framework from Section 2.
- **Convex minimization:** $A = \partial g$ and $B_i = \partial f_i$ for a strongly-convex function g and smooth functions f_i : we recover proximal-SVRG [27] and SAGA [4], to minimize $\min_{z \in \mathcal{E}} g(z) + \sum_{i \in \mathcal{J}} f_i(z)$. However, this is a situation where the operators B_i have an extra property called co-coercivity [7], which we are not using because it is not satisfied for saddle-point problems. The extension of SAGA and SVRG to monotone operators was proposed earlier by [32], but only co-coercive operators are considered, and thus only convex minimization is considered (with important extensions beyond plain SAGA and SVRG), while our analysis covers a much broader set of problems. In particular, the step-sizes obtained with co-coercivity lead to divergence in the general setting.
Because we do not use co-coercivity, applying our results directly to convex minimization, we would get slower rates, while, as shown in Section 2.1, they can be easily cast as a saddle-point problem if the proximal operators of the functions f_i are known, and we then get the same rates than existing fast techniques which are dedicated to this problem [1, 2, 3, 4].
- **Variational inequality problems**, which are notably common in game theory (see, e.g., [6]).

7 Experiments

We consider binary classification problems with design matrix K and label vector in $\{-1, 1\}^n$, a non-separable strongly-convex regularizer with an efficient proximal operator (the sum of the squared norm $\lambda \|x\|^2/2$ and the clustering-inducing term $\sum_{i \neq j} |x_i - x_j|$, for which the proximal operator may be computed in $O(n \log n)$ by isotonic regression [33]) and a non-separable smooth loss (a surrogate to the area under the ROC curve, defined as proportional to $\sum_{i_+ \in I_+} \sum_{i_- \in I_-} (1 - y_i + y_j)^2$, where I_+/I_- are sets with positive/negative labels, for a vector of prediction y , for which an efficient proximal operator may be computed as well, see Appendix E).

Our upper-bounds depend on the ratio $\|K\|_F^2/(\lambda\gamma)$ where λ is the regularization strength and $\gamma \approx n$ in our setting where we minimize an average risk. Setting $\lambda = \lambda_0 = \|K\|_F^2/n^2$ corresponds to a regularization proportional to the average squared radius of the data divided by $1/n$ which is standard in this setting [1]. We also experiment with smaller regularization (i.e., $\lambda/\lambda_0 = 10^{-1}$), to make the problem more ill-conditioned (it turns out that the corresponding testing losses are sometimes slightly better). We consider two datasets, `sido` ($n = 10142$, $d = 4932$, non-separable losses and regularizers presented above) and `rcv1` ($n = 20242$, $d = 47236$, separable losses and regularizer described in Appendix F, so that we can compare with SAGA run in the primal). We report below the squared distance to optimizers which appears in our bounds, as a function of the number of passes on the data (for more details and experiments with primal-dual gaps and testing losses, see Appendix F). Unless otherwise specified, we always use non-uniform sampling.



We see that uniform sampling for SAGA does not improve on batch methods, SAGA and accelerated SVRG improve significantly over the existing methods, with a stronger gain for the accelerated version for ill-conditioned problems (middle vs. left plot). On the right plot, we compare to primal methods on a separable loss, showing that primal methods (here “fba-primal”, which is Nesterov acceleration) that do not use separability (and can thus be applied in all cases) are inferior, while SAGA run on the primal remains faster (but cannot be applied for non-separable losses).

8 Conclusion

We proposed the first linearly convergent incremental gradient algorithms for saddle-point problems, which improve both in theory and practice over existing batch or stochastic algorithms. While we currently need to know the strong convexity-concavity constants, we plan to explore in future work adaptivity to these constants like already obtained for convex minimization [4], paving the way to an analysis without strong convexity-concavity.

Acknowledgements

We would like to thank Simon Lacoste-Julien and Jalal Fadili for fruitful discussions related to saddle-point problems and monotone operators.

References

- [1] N. Le Roux, M. Schmidt, and F. Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Adv. NIPS*, 2012.
- [2] S. Shalev-Shwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.
- [3] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Adv. NIPS*, 2013.
- [4] A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Adv. NIPS*, 2014.
- [5] R. T. Rockafellar. Monotone operators associated with saddle-functions and minimax problems. *Nonlinear Functional Analysis*, 18(part 1):397–407, 1970.
- [6] P. T. Harker and J.-S. Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Math. Prog.*, 48(1-3):161–220, 1990.

- [7] D. L. Zhu and P. Marcotte. Co-coercivity and its role in the convergence of iterative schemes for solving variational inequalities. *SIAM Journal on Optimization*, 6(3):714–726, 1996.
- [8] H. Lin, J. Mairal, and Z. Harchaoui. A universal catalyst for first-order optimization. In *Adv. NIPS*, 2015.
- [9] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer Science & Business Media, 2011.
- [10] D. Woodruff. Sketching as a tool for numerical linear algebra. Technical Report 1411.4357, arXiv, 2014.
- [11] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [12] X. Zhu and A. J. Storkey. Adaptive stochastic primal-dual coordinate descent for separable saddle point problems. In *Machine Learning and Knowledge Discovery in Databases*, pages 645–658. Springer, 2015.
- [13] T. Joachims. A support vector method for multivariate performance measures. In *Proc. ICML*, 2005.
- [14] R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. In *Adv. NIPS*, 1999.
- [15] B. Taskar, S. Lacoste-Julien, and M. I. Jordan. Structured prediction, dual extragradient and Bregman projections. *The Journal of Machine Learning Research*, 7:1627–1653, 2006.
- [16] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- [17] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.
- [18] L. Xu, J. Neufeld, B. Larson, and D. Schuurmans. Maximum margin clustering. In *Adv. NIPS*, 2004.
- [19] F. Bach, J. Mairal, and J. Ponce. Convex sparse matrix factorizations. Technical Report 0812.1869, arXiv, 2008.
- [20] G. H. G. Chen and R. T. Rockafellar. Convergence rates in forward-backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997.
- [21] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [22] Korpelevich G. M. The extra-gradient method for finding saddle points and other problems. *Ekonomika i Matem. Metody*, 12(4):747–756, 1976.
- [23] P. Marcotte. Application of Khobotov’s algorithm to variational inequalities and network equilibrium problems. *Information Systems and Operational Research*, 29(4):258–270, 1991.
- [24] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer, 2004.
- [25] L. Rosasco, S. Villa, and B. C. Vũ. A stochastic forward-backward splitting method for solving monotone inclusions in hilbert spaces. Technical Report 1403.7999, arXiv, 2014.
- [26] K. L. Clarkson, E. Hazan, and D. P. Woodruff. Sublinear optimization for machine learning. *Journal of the ACM (JACM)*, 59(5):23, 2012.
- [27] L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.

- [28] M. Schmidt, R. Babanezhad, M.O. Ahmed, A. Defazio, A. Clifton, and A. Sarkar. Non-uniform stochastic average gradient method for training conditional random fields. In *Proc. AISTATS*, 2015.
- [29] R. Harikandeh, M. O. Ahmed, A. Virani, M. Schmidt, J. Konečný, and S. Sallinen. Stop wasting my gradients: Practical SVRG. In *Adv. NIPS*, 2015.
- [30] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.
- [31] E. Ryu and S. Boyd. A primer on monotone operator methods. *Appl. Comput. Math.*, 15(1):3–43, 2016.
- [32] D. Davis. Smart: The stochastic monotone aggregated root-finding algorithm. Technical Report 1601.00698, arXiv, 2016.
- [33] X. Zeng and M. Figueiredo. Solving OSCAR regularization problems by fast approximate proximal splitting algorithms. *Digital Signal Processing*, 31:124–135, 2014.

A Formalization through Monotone Operators

Throughout the proofs, we will consider only maximal monotone operators on a Euclidean space \mathcal{E} , that is A is assumed to be a μ -strongly monotone (corresponding to M for saddle-points) and potentially set-valued, while B is monotone and L -Lipschitz-continuous with respect to the Euclidean norm (and hence single-valued). For an introduction to monotone operators, see [9, 31].

For simplicity, in this appendix, we will only consider a single-valued operator A (noting that the proof extends to any set-valued operator A), and we will mostly focus here on the monotonicity properties (noting that the “maximal” can be treated rigorously [9], in particular to ensure that the resolvent operator is defined everywhere). An operator is monotone if and only if for all (z, z') , $(A(z) - A(z'))^\top (z - z') \geq 0$. The most basic example is the subdifferential of a convex function. In this paper, we focus on saddle-point problems.

Application to saddle-point problems. For the saddle-point problems defined in Section 2 of the main paper, where we have $z = (x, y)$, we need to make a change of variable because of the two potentially different scaling factors λ and γ . We consider the operators

$$\begin{aligned} B(x, y) &= (\lambda^{-1/2} \partial_x K(\lambda^{-1/2} x, \gamma^{-1/2} y), -\gamma^{-1/2} \partial_y K(\lambda^{-1/2} x, \gamma^{-1/2} y)) \\ A(x, y) &= (\lambda^{-1/2} \partial_x M(\lambda^{-1/2} x, \gamma^{-1/2} y), -\gamma^{-1/2} \partial_y M(\lambda^{-1/2} x, \gamma^{-1/2} y)). \end{aligned}$$

The solutions of $A(x, y) + B(x, y) = 0$ are exactly the solutions of the problem in Eq. (1), rescaled by $\lambda^{1/2}$ and $\gamma^{1/2}$. Moreover, the operator A is 1-monotone, i.e., for any z, z' , we have $(A(z) - A(z'))^\top (z - z') \geq \|z - z'\|^2$. Finally, our definition of the smoothness constants for B and B_i in the main paper, exactly leads to a Lipschitz-constant of L with respect to the natural Euclidean norm (a similar result holds for the constant $\bar{L}(\pi)$ defined later). Moreover, convergence results in the Euclidean norm here transfer to convergence results in the norm Ω defined in the main paper. Note that because of our proofs through operators, it is not easily possible to get bounds on the primal and dual gaps.

Properties of monotone operators and resolvents. Given a maximal monotone operator A , we may define its *resolvent* operator as $z' = (I + \sigma A)^{-1}(z)$, which is defined as finding z' such that $z' + \sigma A(z') = z$. When A is the operator associated to the saddle-point function M as described above, then the resolvent operator is exactly the proximal operator of M defined in Eq. (2) of the main paper. Note that care has to be taken with the scaling factors λ and γ .

We will use the following properties (on top of Lipschitz-continuity) [9, 31]:

- Monotonicity property: for any (z, z') , $(B(z) - B(z'), z - z') \geq 0$.
- Contractivity of the resolvent operator for A μ -strongly-monotone: for any (z, z') , $\|(I + \sigma A)^{-1}(z) - (I + \sigma A)^{-1}(z')\| \leq (1 + \sigma\mu)^{-1} \|z - z'\|$.
- Firm non-expansiveness of the resolvent: for any (z, z') , $\|(I + \sigma A)^{-1}(z) - (I + \sigma A)^{-1}(z')\|^2 \leq (1 + \sigma\mu)^{-1}(z - z')^\top ((I + \sigma A)^{-1}(z) - (I + \sigma A)^{-1}(z'))$.

Moreover, given our strong-monotonicity assumption, $A + B$ has a unique zero $z_* \in \mathcal{E}$.

Finally in order to characterize the running-times, we will consider the complexity $T_{\text{fw}}(B)$ of computing the operator B and the complexity $T_{\text{bw}}(A)$ to compute the resolvent of A . For saddle-point problems, these correspond to $T(B)$ and $T_{\text{prox}}(M)$ from the main paper.

B Proof for Deterministic Algorithms

All proofs in this section will follow the same principle, by showing that at every step of our algorithms, a certain function (a ‘‘Lyapunov’’ function) is contracted by a factor strictly less than one. For the forward-backward algorithm, this will be the distance to optimum $\|z_t - z_*\|^2$; while for the accelerated version, it will be different.

B.1 Forward-backward algorithm

We consider the iteration $z_t = (I + \sigma A)^{-1}(z_{t-1} - \sigma B(z_{t-1}))$, with B being monotone L -Lipschitz-continuous and A being μ -strongly monotone. The optimum z_* (i.e., the zero of $A + B$) is invariant by this iteration. Note that this is the analysis of [20] and that we could improve by putting some of the strong-monotonicity in the operator B rather than in A .

We have:

$$\begin{aligned}
& \|z_t - z_*\|^2 \\
& \leq \frac{1}{(1 + \sigma\mu)^2} \|z_{t-1} - z_* - \sigma(B(z_{t-1}) - B(z_*))\|^2 \text{ by contractivity of the resolvent,} \\
& = \frac{1}{(1 + \sigma\mu)^2} \left[\|z_{t-1} - z_*\|^2 - 2\sigma(z_{t-1} - z_*)^\top (B(z_{t-1}) - B(z_*)) + \sigma^2 \|B(z_{t-1}) - B(z_*)\|^2 \right] \\
& \leq \frac{1}{(1 + \sigma\mu)^2} (1 + \sigma^2 L^2) \|z_{t-1} - z_*\|^2 \text{ by monotonicity of and Lipschitz-continuity of } B, \\
& \leq \left(\frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} \right)^t \|z_0 - z_*\|^2, \text{ by applying the recursion } t \text{ times.}
\end{aligned}$$

Thus we get linear (i.e., geometric) convergence as soon as $1 + \sigma^2 L^2 < (1 + \sigma\mu)^2$. If we consider $\eta = \frac{\sigma\mu}{1 + \sigma\mu} \in [0, 1)$, and the rate above becomes equal to:

$$\frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} = (1 - \eta)^2 + \eta^2 \frac{L^2}{\mu^2} = 1 - 2\eta + \eta^2 \left(1 + \frac{L^2}{\mu^2}\right),$$

thus the algorithm converges if $\eta < \frac{2}{1 + \frac{L^2}{\mu^2}}$, and with $\eta = \frac{1}{1 + \frac{L^2}{\mu^2}}$ which corresponds to $\sigma = \frac{1}{\mu} \frac{\eta}{1 - \eta} = \frac{\mu}{L^2}$, we get a linear convergence rate with constant $1 - \eta = \frac{L^2}{\mu^2 + L^2}$.

Thus the complexity to reach the precision $\varepsilon \times \|z_0 - z_*\|^2$ in squared distance to optimum $\|z_t - z_*\|^2$ is equal to $(1 + \frac{L^2}{\mu^2})[T_{\text{fw}}(B) + T_{\text{bw}}(A)] \log \frac{1}{\varepsilon}$.

Note that we obtain a slow convergence when applied to convex minimization, because we are not using any co-coercivity of B , which would lead to a rate $(1 - \mu/L)$ [7]. Indeed, co-coercivity means that $\|B(z) - B(z')\|^2 \leq L(B(z) - B(z'))^\top (z - z')$, and this allows to replace above the term $1 + \sigma^2 L^2$ by 1 if $\sigma \leq 2/L$, leading to linear convergence rate with constant $(1 + \mu/L)^{-2} \approx 1 - 2\mu/L$.

B.2 Accelerated forward-backward algorithm

We consider the iteration $z_t = (I + \sigma A)^{-1}(z_{t-1} - \sigma B[z_{t-1} + \theta(z_{t-1} - z_{t-2})])$, with B being monotone L -Lipschitz-continuous and *linear*, and A being μ -strongly monotone. Note that this is an extension of the analysis of [21] to take into account the general monotone operator situation. Again z_* is a fixed-point of the iteration.

Using the firm non-expansiveness of the resolvent operator, we get, with $\eta = \frac{\sigma\mu}{1+\sigma\mu}$, and then using the linearity of B :

$$\begin{aligned} \|z_t - z_*\|^2 &\leq \frac{1}{1 + \sigma\mu} (z_t - z_*)^\top \left[z_{t-1} - z_* - \sigma B[z_{t-1} - z_* + \theta(z_{t-1} - z_{t-2})] \right] \\ &= (z_t - z_*)^\top \left[(1 - \eta)(z_{t-1} - z_*) - \frac{\eta}{\mu} B[z_{t-1} - z_* + \theta(z_{t-1} - z_{t-2})] \right] \\ &= -\frac{1 - \eta}{2} \|z_t - z_{t-1}\|^2 + \frac{1 - \eta}{2} \|z_t - z_*\|^2 + \frac{1 - \eta}{2} \|z_{t-1} - z_*\|^2 \\ &\quad - \frac{\eta}{\mu} (z_t - z_*)^\top B[z_{t-1} - z_* + \theta(z_{t-1} - z_{t-2})] \\ &= -\frac{1 - \eta}{2} \|z_t - z_{t-1}\|^2 + \frac{1 - \eta}{2} \|z_t - z_*\|^2 + \frac{1 - \eta}{2} \|z_{t-1} - z_*\|^2 \\ &\quad - \frac{\eta}{\mu} (z_t - z_*)^\top B(z_{t-1} - z_*) - \theta \frac{\eta}{\mu} (z_t - z_*)^\top B(z_{t-1} - z_{t-2}), \end{aligned}$$

by regrouping terms. By using the Lipschitz-continuity of B , we get:

$$\begin{aligned} &\|z_t - z_*\|^2 \\ &\leq -\frac{1 - \eta}{2} \|z_t - z_{t-1}\|^2 + \frac{1 - \eta}{2} \|z_t - z_*\|^2 + \frac{1 - \eta}{2} \|z_{t-1} - z_*\|^2 - \frac{\eta}{\mu} (z_t - z_*)^\top B(z_{t-1} - z_t) \\ &\quad - \theta \frac{\eta}{\mu} (z_{t-1} - z_*)^\top B(z_{t-2} - z_{t-1}) + \theta \frac{\eta}{\mu} L \|z_t - z_{t-1}\| \|z_{t-1} - z_{t-2}\| \\ &\leq -\frac{1 - \eta}{2} \|z_t - z_{t-1}\|^2 + \frac{1 - \eta}{2} \|z_t - z_*\|^2 + \frac{1 - \eta}{2} \|z_{t-1} - z_*\|^2 - \frac{\eta}{\mu} (z_t - z_*)^\top B(z_{t-1} - z_t) \\ &\quad - \theta \frac{\eta}{\mu} (z_{t-1} - z_*)^\top B(z_{t-2} - z_{t-1}) + \frac{\theta L}{2} \frac{\eta}{\mu} [\alpha^{-1} \|z_t - z_{t-1}\|^2 + \alpha \|z_{t-1} - z_{t-2}\|^2], \end{aligned}$$

with a constant $\alpha > 0$ to be determined later. This leads to, with $\theta = \frac{1-\eta}{1+\eta}$, and by regrouping terms:

$$\begin{aligned} &\frac{1 + \eta}{2} \|z_t - z_*\|^2 + \left(\frac{1 - \eta}{2} - \frac{\theta\eta L}{2\mu} \alpha^{-1} \right) \|z_t - z_{t-1}\|^2 - \eta (z_t - z_*)^\top B(z_{t-1} - z_t) \\ &\leq \frac{1 - \eta}{2} \|z_{t-1} - z_*\|^2 + \left(\frac{\alpha\eta\theta L}{2\mu} \right) \|z_{t-1} - z_{t-2}\|^2 - \theta \frac{\eta}{\mu} (z_{t-1} - z_*)^\top B(z_{t-2} - z_{t-1}) \\ &\leq \theta \left[\frac{1 + \eta}{2} \|z_{t-1} - z_*\|^2 + \left(\frac{\eta\alpha L}{2\mu} \right) \|z_{t-1} - z_{t-2}\|^2 - \frac{\eta}{\mu} (z_{t-1} - z_*)^\top B(z_{t-2} - z_{t-1}) \right]. \end{aligned}$$

We get a Lyapunov function $\mathcal{L} : (z, z') \mapsto \frac{1+\eta}{2}\|z - z_*\|^2 + \left(\frac{1-\eta}{2} - \frac{\theta\eta L}{2\mu}\alpha^{-1}\right)\|z - z'\|^2 - \eta(z - z_*)^\top B(z' - z)$, such that $\mathcal{L}(z_t, z_{t-1})$ converges to zero geometrically, if $\frac{\alpha\eta L}{\mu} \leq 1 - \eta - \frac{\eta\theta L}{\mu}\alpha^{-1}$ and $\begin{pmatrix} 1 + \eta & -\eta L/\mu \\ \eta L/\mu & 1 - \eta - \eta\theta L\mu^{-1}\alpha^{-1} \end{pmatrix} \succ 0$. By setting $\eta = \frac{1}{1+2L/\mu}$, and thus $\theta = \frac{1-\eta}{1+\eta} = \frac{1}{1+\mu/L}$, $\sigma = \frac{1}{\mu} \frac{\eta}{1-\eta} = \frac{1}{2L}$, and $\alpha = 1$, we get the desired first property and the fact that the matrix above is greater than $\begin{pmatrix} 1/2 & 0 \\ 0 & 0 \end{pmatrix}$, which allows us to get a linear rate of convergence for $\|z_t - z_*\|^2 \leq 2\mathcal{L}(z_t, z_{t-1})$.

C Proof for Existing Stochastic Algorithms

We follow [25], but with a specific step-size that leads to a simple result, which also applies to non-uniform sampling from a finite pool. We consider the iteration $z_t = (I + \sigma_t A)^{-1}(z_{t-1} - \sigma_t(Bz_{t-1} + C_t z_{t-1}))$, with B being monotone L -Lipschitz-continuous and A being μ -strongly monotone, and C_t a random operator (*not necessarily monotone*) such that $\mathbb{E}C_t(z) = 0$ for all z . We assume that all random operators C_t are independent, and we denote by \mathcal{F}_t the σ -field generated by C_1, \dots, C_t , i.e., the information up to time t .

We have with $\text{Lip}(C_t)$ the Lipschitz-constant of C_t :

$$\begin{aligned} \|z_t - z_*\|^2 &\leq \frac{1}{(1 + \sigma_t \mu)^2} \|z_{t-1} - z_* - \sigma_t(B(z_{t-1}) - B(z_*)) - \sigma_t C_t(z_{t-1})\|^2 \\ &\quad \text{by contractivity of the resolvent,} \\ &= \frac{1}{(1 + \sigma_t \mu)^2} \left[\|z_{t-1} - z_*\|^2 - 2\sigma_t(z_{t-1} - z_*)^\top (B(z_{t-1}) - B(z_*)) \right. \\ &\quad \left. + \sigma_t^2 \|B(z_{t-1}) - B(z_*) + C_t(z_{t-1})\|^2 + 2\sigma_t(C_t(z_{t-1}))^\top (z_{t-1} - z_*) \right]. \end{aligned}$$

By taking conditional expectations, we get:

$$\begin{aligned} \mathbb{E}(\|z_t - z_*\|^2 | \mathcal{F}_{t-1}) &\leq \frac{1}{(1 + \sigma_t \mu)^2} \left[(1 + \sigma_t^2 L^2) \|z_{t-1} - z_*\|^2 + \sigma_t^2 \mathbb{E}(\|C_t(z_{t-1})\|^2 | \mathcal{F}_{t-1}) \right] \\ &\quad \text{by monotonicity and Lipschitz-continuity of } B, \\ &\leq \frac{1}{(1 + \sigma_t \mu)^2} \left[(1 + \sigma_t^2 L^2) \|z_{t-1} - z_*\|^2 + 2\sigma_t^2 \mathbb{E}(\|C_t(z_*)\|^2 | \mathcal{F}_{t-1}) \right. \\ &\quad \left. + 2\sigma_t^2 \|z_{t-1} - z_*\|^2 \mathbb{E}\left(\sup_{\|z-z'\|=1} \|C_t(z) - C_t(z')\|^2 | \mathcal{F}_{t-1}\right) \right] \\ &= \frac{1}{(1 + \sigma_t \mu)^2} \left[(1 + \sigma_t^2 L^2) \|z_{t-1} - z_*\|^2 + 2\sigma_t^2 \mathbb{E}(\|C_t(z_*)\|^2 | \mathcal{F}_{t-1}) \right. \\ &\quad \left. + 2\sigma_t^2 \|z_{t-1} - z_*\|^2 \mathbb{E}(\text{Lip}(C_t)^2 | \mathcal{F}_{t-1}) \right] \\ &= \frac{1}{(1 + \sigma_t \mu)^2} \left[(1 + \sigma_t^2 L^2 + 2\sigma_t^2 \mathbb{E}(\text{Lip}(C_t)^2 | \mathcal{F}_{t-1})) \|z_{t-1} - z_*\|^2 + 2\sigma_t^2 \mathbb{E}(\|C_t(z_*)\|^2 | \mathcal{F}_{t-1}) \right]. \end{aligned}$$

By denoting $\eta_t = \frac{\sigma_t \mu}{1 + \sigma_t \mu} \in [0, 1)$, we get

$$\mathbb{E}\|z_t - z_*\|^2 \leq \left(1 - 2\eta_t + \eta_t^2 + 2\eta_t^2 \frac{L^2}{\mu^2} + 2\eta_t^2 \frac{1}{\mu^2} \mathbb{E}(\text{Lip}(C_t)^2 | \mathcal{F}_{t-1})\right) \|z_{t-1} - z_*\|^2 + 2\frac{\eta_t^2}{\mu^2} \mathbb{E}(\|C_t z_*\|^2 | \mathcal{F}_{t-1}).$$

By selecting $\eta_t = \frac{2}{(t+1)+4\frac{L^2}{\mu^2} + \frac{4}{\mu^2}\mathbb{E}(\text{Lip}(C_t)^2|\mathcal{F}_{t-1})} = \frac{2}{t+1+A}$, with $A = 4\frac{L^2}{\mu^2} + \frac{4}{\mu^2}\mathbb{E}(\text{Lip}(C_t)^2|\mathcal{F}_{t-1})$, we get:

$$\begin{aligned}
\mathbb{E}\|z_t - z_*\|^2 &\leq (1 - \eta_t)\mathbb{E}\|z_{t-1} - z_*\|^2 + 2\frac{\eta_t^2}{\mu^2}\mathbb{E}(\|C_t z_*\|^2) \\
&= \frac{t-1+A}{t+1+A}\mathbb{E}\|z_{t-1} - z_*\|^2 + \frac{8}{(t+1+A)^2}\frac{1}{\mu^2}\mathbb{E}(\|C_t z_*\|^2) \\
&\leq \frac{A(1+A)}{(t+1+A)(t+A)}\|z_0 - z_*\|^2 + \frac{8}{\mu^2}\sum_{u=1}^t \frac{(u+A)(u+1+A)}{(t+1+A)(t+A)}\frac{1}{(u+1+A)^2}\mathbb{E}(\|C_u z_*\|^2) \\
&\quad \text{by expanding the recursion } t \text{ times,} \\
&\leq \frac{A(1+A)}{(t+1+A)(t+A)}\|z_0 - z_*\|^2 + \frac{8}{\mu^2}\sum_{u=1}^t \frac{1}{(t+1+A)(t+A)}\mathbb{E}(\|C_u z_*\|^2) \\
&\leq \frac{(1+A)^2}{(t+A)^2}\|z_0 - z_*\|^2 + \frac{8}{\mu^2(t+A)}\sup_{u \in \{1, \dots, t\}} \mathbb{E}(\|C_u z_*\|^2).
\end{aligned}$$

The overall convergence rate is in $O(1/t)$ and the constant depends on the noise in the operator values at the optimum. Note that initial conditions are forgotten at a rate $O(1/t^2)$.

Application to sampling from a finite family. When sampling from $|\mathcal{J}|$ operators B_i , $i \in \mathcal{J}$, and selecting i_t with probability vector π , then we have $\mathbb{E}(\text{Lip}(C_t)^2|\mathcal{F}_{t-1}) \leq \bar{L}(\pi)^2 = \bar{L}^2$ defined as $\sup_{\|z-z'\| \leq 1} \sqrt{\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z) - B_i(z')\|^2}$. Thus, we can take the step-size $\frac{2}{t+1+4\frac{\bar{L}^2+L^2}{\mu^2}}$, and thus $\sigma_t = \frac{2/\mu}{t+1+4\frac{\bar{L}^2+L^2}{\mu^2}}$. Moreover, if L is unknown (or hard to compute), we can take \bar{L} instead.

We may further bound: $\mathbb{E}(\|C_u z_*\|^2) \leq 2\mathbb{E}(\|C_u z_0\|^2) + 2\mathbb{E}(\text{Lip}(C_t)^2)\|z_0 - z_*\|^2$, and thus, if we start from an initial point z_0 such that $C_u z_0 = 0$, which is always possible for bi-linear problems, we get an overall bound of (taking $L = \bar{L}$ for simplicity)

$$\left(\frac{(1 + 8\bar{L}^2/\mu^2)^2}{(t + 8\bar{L}^2/\mu^2)^2} + \frac{16\bar{L}^2/\mu^2}{t + 8\bar{L}^2/\mu^2} \right) \|z_0 - z_*\|^2 \leq \frac{1 + 24\bar{L}^2/\mu^2}{t + 8\bar{L}^2/\mu^2} \|z_0 - z_*\|^2.$$

We thus get an overall $O(1/t)$ convergence rate.

D Proof for New Stochastic Algorithms

We also consider the monotone operator set-up, since this is the only assumption that we use. We follow the proof of the corresponding convex minimization algorithms, with key differences which we highlight below. In particular, (a) we do not use function values, and (b) we use shorter step-sizes to tackle the lack of co-coercivity.

D.1 SVRG: Stochastic-Variance reduced saddle-point problems (Theorem 1)

We only analyze a single epoch starting from the reference estimate \tilde{z} , and show that the expected squared distance to optimum is shrunk by a factor of 3/4 if the number of iterations per epoch is well-chosen. The epoch is started with $z_0 = \tilde{z}$.

We denote by \mathcal{F}_{t-1} the information up to time $t-1$. We consider sampling $i_{t1}, \dots, i_{tm} \in \mathcal{J}$ with replacement at time t . By using the contractivity of the resolvent operator of A , and the fact that $z_* = (I + \sigma A)^{-1}(z_* - \sigma B(z_*))$, we get:

$$\begin{aligned} \|z_t - z_*\|^2 &\leq \frac{1}{(1 + \sigma\mu)^2} \left\| z_{t-1} - z_* - \sigma[B(\tilde{z}) - B(z_*) + \frac{1}{m} \sum_{k=1}^m \frac{1}{\pi_{i_{tk}}} (B_{i_{tk}}(z_{t-1}) - B_{i_{tk}}(\tilde{z}))] \right\|^2 \\ &= \frac{1}{(1 + \sigma\mu)^2} \left\| z_{t-1} - z_* \right. \\ &\quad \left. - \sigma[B(z_{t-1}) - B(z_*) + \frac{1}{m} \sum_{k=1}^m \frac{1}{\pi_{i_{tk}}} (B_{i_{tk}}(z_{t-1}) - B_{i_{tk}}(\tilde{z})) - (B(z_{t-1}) - B(\tilde{z}))] \right\|^2. \end{aligned}$$

Expanding the squared norm, taking conditional expectations with $\mathbb{E}(\frac{1}{\pi_{i_{tk}}} B_{i_{tk}} | \mathcal{F}_{t-1}) = B$, and using the independence of i_{t1}, \dots, i_{tm} , we get:

$$\begin{aligned} &\mathbb{E}[\|z_t - z_*\|^2 | \mathcal{F}_{t-1}] \\ &\leq \frac{1}{(1 + \sigma\mu)^2} (\|z_{t-1} - z_*\|^2 - 2\sigma(z_{t-1} - z_*)^\top (B(z_{t-1}) - B(z_*)) + \sigma^2 \|B(z_{t-1}) - B(z_*)\|^2) \\ &\quad + \frac{1}{m} \mathbb{E} \left[\frac{1}{(1 + \sigma\mu)^2} \left\| \frac{1}{\pi_{i_t}} (B_{i_t}(z_{t-1}) - B_{i_t}(\tilde{z})) - (B(z_{t-1}) - B(\tilde{z})) \right\|^2 \middle| \mathcal{F}_{t-1} \right]. \end{aligned}$$

Using the monotonicity of B and the Lipschitz-continuity of B (like in Appendix B.1), we get the bound

$$\frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} \|z_{t-1} - z_*\|^2 + \frac{1}{m} \mathbb{E} \left[\frac{1}{(1 + \sigma\mu)^2} \left\| \frac{1}{\pi_{i_t}} (B_{i_t}(z_{t-1}) - B_{i_t}(\tilde{z})) - (B(z_{t-1}) - B(\tilde{z})) \right\|^2 \middle| \mathcal{F}_{t-1} \right].$$

We denote by \bar{L}^2 the quantity $\bar{L}^2 = \sup_{z, z' \in \mathcal{E}} \frac{1}{\|z - z'\|^2} \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z) - B_i(z')\|^2$. We then have (using the fact that a variance is less than the second-order moment):

$$\mathbb{E} \left[\left\| \frac{1}{\pi_{i_t}} (B_{i_t}(z_{t-1}) - B_{i_t}(\tilde{z})) - (B(z_{t-1}) - B(\tilde{z})) \right\|^2 \middle| \mathcal{F}_{t-1} \right] \leq \mathbb{E} \left[\left\| \frac{1}{\pi_{i_t}} (B_{i_t}(z_{t-1}) - B_{i_t}(\tilde{z})) \right\|^2 \middle| \mathcal{F}_{t-1} \right],$$

which is less than $\bar{L}^2 \|z_{t-1} - \tilde{z}\|^2$ because we sample i from π . This leads to

$$\begin{aligned} \mathbb{E}[\|z_t - z_*\|^2 | \mathcal{F}_{t-1}] &\leq \frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} \|z_{t-1} - z_*\|^2 + \frac{1}{(1 + \sigma\mu)^2} \frac{\bar{L}^2}{m} \|z_{t-1} - \tilde{z}\|^2 \\ &\leq \left(1 - 2\eta + \eta^2 + \eta^2 \frac{L^2}{\mu^2} + \frac{(1+a)\eta^2 \bar{L}^2}{\mu^2 m} \right) \|z_{t-1} - z_*\|^2 \\ &\quad + \frac{(1+a^{-1})\eta^2 \bar{L}^2}{\mu^2 m} \|\tilde{z} - z_*\|^2, \end{aligned}$$

with $\eta = \frac{\sigma\mu}{1 + \sigma\mu} \in [0, 1)$ and $a > 0$ to be determined later. Assuming that $\eta \left(1 + \frac{L^2}{\mu^2} + \frac{(1+a)\bar{L}^2}{\mu^2 m} \right) \leq 1$, and taking full expectations, this leads to:

$$\mathbb{E}\|z_t - z_*\|^2 \leq (1 - \eta) \mathbb{E}\|z_{t-1} - z_*\|^2 + \frac{(1+a^{-1})\eta^2 \bar{L}^2}{\mu^2 m} \|\tilde{z} - z_*\|^2,$$

that is we get a shrinking of the expected distance to optimum with additional noise that depends on the distance to optimum of the reference point \tilde{z} . The difference with the convex minimization set-up of [27] is that the proof is more direct, and we get a shrinkage directly on the iterates (we

have no choice for monotone operators), without the need to do averaging of the iterates. Moreover, we never use any monotonicity of the operators B_i , thus allowing any type of splits (as long as the sum B is monotone).

Then, using the fact that $z_0 = \tilde{z}$, and expanding the recursion:

$$\begin{aligned}\mathbb{E}\|z_t - z_*\|^2 &\leq (1 - \eta)^t \|z_0 - z_*\|^2 + \left(\sum_{u=0}^{t-1} (1 - \eta)^u \right) \frac{(1 + a^{-1})\eta^2 \bar{L}^2}{\mu^2 m} \|\tilde{z} - z_*\|^2 \\ &\leq \left((1 - \eta)^t + \frac{(1 + a^{-1})\eta \bar{L}^2}{\mu^2 m} \right) \|\tilde{z} - z_*\|^2.\end{aligned}$$

If we take $a = 2$, $\eta = \frac{1}{1 + L^2 + 3\bar{L}^2/(m\mu^2)}$, which corresponds to $\sigma = \frac{1}{\mu} \frac{\eta}{1 - \eta} = \frac{\mu}{L^2 + \frac{3}{m}\bar{L}^2}$ and $t = \log 4/\eta = \log 4 \times (1 + \frac{L^2}{\mu^2} + 3\frac{\bar{L}^2}{m\mu^2})$, we obtain a bound of $3/4$, that is, after t step in an epoch, we obtain $\mathbb{E}\|z_t - z_*\|^2 \leq \frac{3}{4} \|\tilde{z} - z_*\|^2$, which is the desired result.

In terms of running-time, we therefore need a time proportional to $T(B) + (1 + \frac{L^2}{\mu^2} + 3\frac{\bar{L}^2}{m\mu^2})(m \max_i T(B_i) + T_{\text{prox}}(A))$, times $\log \frac{1}{\varepsilon}$ to reach precision ε .

Note that if L^2 is too expensive to compute (because it is a global constant), we may replace it by \bar{L}^2 and get a worse bound (but still a valid algorithm).

D.2 SAGA: Online stochastic-variance reduced saddle-point problems (Theorem 2)

The proof follows closely the one of SVRG above. Following the same arguments, we get, by contractivity of the resolvent operator:

$$\begin{aligned}\|z_t - z_*\|^2 &\leq \frac{1}{(1 + \sigma\mu)^2} \left\| z_{t-1} - z_* - \sigma \left[\sum_{i \in \mathcal{J}} g_{t-1}^i - B(z_*) + \frac{1}{m} \sum_{k=1}^m \frac{1}{\pi_{i_{tk}}} (B_{i_{tk}}(z_{t-1}) - g_{t-1}^{i_{tk}}) \right] \right\|^2 \\ &= \frac{1}{(1 + \sigma\mu)^2} \left\| z_{t-1} - z_* - \sigma [B(z_{t-1}) - B(z_*) \right. \\ &\quad \left. + \frac{1}{m} \sum_{k=1}^m \frac{1}{\pi_{i_{tk}}} (B_{i_{tk}}(z_{t-1}) - g_{t-1}^{i_{tk}}) - (B(z_{t-1}) - \sum_{i \in \mathcal{J}} g_{t-1}^i)] \right\|^2.\end{aligned}$$

Then, using independence, monotonicity and Lipschitz-continuity of B , we get (note that we never use any monotonicity of B_i), like in the proof of Theorem 1:

$$\begin{aligned}\mathbb{E}[\|z_t - z_*\|^2 | \mathcal{F}_{t-1}] &\leq \frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} \|z_{t-1} - z_*\|^2 \\ &\quad + \frac{1}{m} \mathbb{E} \left[\frac{1}{(1 + \sigma\mu)^2} \left\| \frac{1}{\pi_{i_t}} (B_{i_t}(z_{t-1}) - g_{t-1}^{i_t}) - (B(z_{t-1}) - \sum_{i \in \mathcal{J}} g_{t-1}^i) \right\|^2 \middle| \mathcal{F}_{t-1} \right] \\ &\leq \frac{1 + \sigma^2 L^2}{(1 + \sigma\mu)^2} \|z_{t-1} - z_*\|^2 + \frac{1}{m} \frac{1}{(1 + \sigma\mu)^2} \left(\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_{t-1}) - g_{t-1}^i\|^2 \right) \\ &\leq \left(1 - 2\eta + \eta^2 + \eta^2 \frac{L^2}{\mu^2} + \frac{(1 + a)\eta^2 \bar{L}^2}{\mu^2 m} \right) \|z_{t-1} - z_*\|^2 \\ &\quad + \frac{(1 + a^{-1})\eta^2}{\mu^2 m} \left(\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 \right),\end{aligned}$$

with $\eta = \frac{\sigma\mu}{1+\sigma\mu}$. Assuming $\eta(1 + \frac{L^2}{\mu^2} + \frac{(1+a)\bar{L}^2}{\mu^2 m}) \leq 1$, we get

$$\mathbb{E}[\|z_t - z_*\|^2 | \mathcal{F}_{t-1}] \leq (1 - \eta)\|z_{t-1} - z_*\|^2 + \frac{(1+a^{-1})\eta^2}{\mu^2 m} \left(\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 \right).$$

Like in the SVRG proof above, we get a contraction of the distance to optimum, with now an added noise that depends on the difference between our stored operator values and the operator values at the global optimum. We thus need to control this distance by adding the proper factors to a Lyapunov function. Note that we never use any monotonicity of the operators B_i , thus allowing any type of splits (as long as the sum B is monotone).

We assume that we update (at most m because we are sampling with replacement and we may sample the same gradient twice) “gradients” g_t^i uniformly at random (when we consider uniform sampling, we can reuse the same gradients as dependence does not impact the bound), by replacing them by $g_t^i = B_i(z_{t-1})$. Thus:

$$\begin{aligned} & \mathbb{E} \left(\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_t^i\|^2 \middle| \mathcal{F}_{t-1} \right) \\ &= \mathbb{E} \left(\sum_{i \text{ selected}} \frac{1}{\pi_i} \|B_i(z_*) - B_i(z_{t-1})\|^2 + \sum_{i \text{ non selected}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 \middle| \mathcal{F}_{t-1} \right) \\ &= \mathbb{E} \left(\sum_{i \text{ selected}} \frac{1}{\pi_i} \left(\|B_i(z_*) - B_i(z_{t-1})\|^2 - \|B_i(z_*) - g_{t-1}^i\|^2 \right) + \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 \middle| \mathcal{F}_{t-1} \right). \end{aligned}$$

Since we sample uniformly *with* replacement, the marginal probabilities of selecting an element i is equal to $\rho = 1 - (1 - \frac{1}{|\mathcal{J}|})^m$. We thus get

$$\begin{aligned} \mathbb{E} \left(\sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_t^i\|^2 \middle| \mathcal{F}_{t-1} \right) &\leq (1 - \rho) \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 + \rho \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - B_i(z_{t-1})\|^2 \\ &\leq (1 - \rho) \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2 + \rho \bar{L}^2 \|z_{t-1} - z_*\|^2. \end{aligned}$$

Therefore, overall, we have, for a scalar $b > 0$ to be chosen later:

$$\begin{aligned} & \mathbb{E} \left(\|z_t - z_*\|^2 + b \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_t^i\|^2 \middle| \mathcal{F}_{t-1} \right) \\ &\leq \left(1 - 2\eta + \eta^2 + \eta^2 \frac{L^2}{\mu^2} + \frac{(1+a)\eta^2 \bar{L}^2}{\mu^2 m} + b\rho \bar{L}^2 \right) \|z_{t-1} - z_*\|^2 \\ &\quad + b \left(1 - \rho + b^{-1} \frac{(1+a^{-1})\eta^2}{m\mu^2} \right) \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2. \end{aligned}$$

If we take $a = 2$, $\eta = \frac{1}{\max\{\frac{3|\mathcal{J}|}{2m}, 1 + \frac{L^2}{\mu^2} + 3\frac{\bar{L}^2}{m\mu^2}\}}$, which corresponds to $\sigma = \frac{1}{\mu} \frac{\eta}{1-\eta} = \frac{\mu}{\max\{\frac{3|\mathcal{J}|}{2m} - 1, \frac{L^2}{\mu^2} + 3\frac{\bar{L}^2}{m\mu^2}\}}$, with $b\rho \bar{L}^2 = \frac{3\eta}{4}$, then we get the bound (using $\eta \leq 1/(\bar{L}^2/(3m))$):

$$(1 - \frac{\eta}{4}) \|z_{t-1} - z_*\|^2 + (1 - \frac{\rho}{3}) \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g_{t-1}^i\|^2,$$

which shows that the function $(z, g) \mapsto \|z - z_*\|^2 + b \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - g^i\|^2$ is a good Lyapunov function for the problem that shrinks geometrically in expectation (it resembles the one from convex minimization, but without the need for function values).

Finally, since we assume that $m \leq |\mathcal{J}|$, we have $\rho = 1 - (1 - 1/|\mathcal{J}|)^m \geq 1 - \exp(-m/|\mathcal{J}|) \geq m/(2|\mathcal{J}|)$. This leads to, after t iterations

$$\mathbb{E}\|z_t - z_*\|^2 \leq (1 - \min\{\frac{\eta}{4}, \frac{m}{6|\mathcal{J}}\})^t \left[\|z_0 - z_*\|^2 + \frac{3\eta}{4\rho L^2} \sum_{i \in \mathcal{J}} \frac{1}{\pi_i} \|B_i(z_*) - B_i(z_0)\|^2 \right].$$

We have $\eta \leq 2m/(3|\mathcal{J}|)$ and $3\eta/(4\rho) \leq \frac{3}{4} \frac{2m}{3|\mathcal{J}|} \frac{2|\mathcal{J}|}{m} \leq 1$, leading to

$$\mathbb{E}\|z_t - z_*\|^2 \leq 2(1 - \frac{\eta}{4})^t \|z_0 - z_*\|^2,$$

which is the desired result.

Note that we get the same overall running-time complexity than for SVRG.

Factored splits. Note that when applying to saddle-points with factored splits, we need to use a Lyapunov function that considers these splits. The only difference is to treat separately the two parts of the vectors, leading to replacing everywhere $|\mathcal{J}|$ by $\max\{|\mathcal{J}|, |\mathcal{K}|\}$.

D.3 Acceleration

We also consider in this section a proof based on monotone operators. We first give the algorithm for saddle-point problems.

Algorithms for saddle-point problems. At each iteration, we need solve the problem in Eq. (4) of the main paper, with the SVRG algorithm applied to $\tilde{K}(x, y) = K(x, y) - \lambda\tau x^\top \bar{x} + \gamma\tau y^\top \bar{y}$, and $\tilde{M}(x, y) = M(x, y) + \frac{\lambda\tau}{2}\|x\|^2 - \frac{\gamma\tau}{2}\|y\|^2$. These functions lead to constants $\lambda = \lambda(1 + \tau)$, $\tilde{\gamma} = \gamma(1 + \tau)$ and $\tilde{L} = L/(1 + \tau)$, $\tilde{\sigma} = \sigma(1 + \tau)^2$. We thus get the iteration, for a single selected operator,

$$(x, y) \leftarrow \text{prox}_{\tilde{M}}^{\tilde{\sigma}} \left[(x, y) - \tilde{\sigma} \begin{pmatrix} 1/\tilde{\lambda} & 0 \\ 0 & 1/\tilde{\gamma} \end{pmatrix} (\tilde{B}(\tilde{x}, \tilde{y}) + \left\{ \frac{1}{\pi_i} \tilde{B}_i(x, y) - \frac{1}{\pi_i} \tilde{B}_i(\tilde{x}, \tilde{y}) \right\}) \right].$$

A short calculation shows that $\text{prox}_{\tilde{M}}^{\tilde{\sigma}}(x, y) = \text{prox}_M^{\sigma(1+\tau)/(1+\sigma\tau(1+\tau))}((x, y)/(1 + \sigma\tau(1 + \tau)))$, leading to the update (with σ the step-size from the regular SVRG algorithm in Section 3):

$$(x, y) \leftarrow \text{prox}_{\tilde{M}}^{\tilde{\sigma}} \left[(x, y) + \sigma\tau(1 + \tau)(\bar{x}, \bar{y}) - \sigma(1 + \tau) \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} (B(\tilde{x}, \tilde{y}) + \left\{ \frac{1}{\pi_i} \tilde{B}_i(x, y) - \frac{1}{\pi_i} \tilde{B}_i(\tilde{x}, \tilde{y}) \right\}) \right].$$

This leads to Algorithm 3, where differences with the SVRG algorithm, e.g., Algorithm 1, are highlighted in red. Given the value of τ , the estimate (\bar{x}, \bar{y}) is updated every $\log(1 + \tau)$ epochs of SVRG. While this leads to a provably better convergence rate, in practice, this causes the algorithm to waste time solving with too high precision the modified problem. We have used the simple heuristic of changing (\bar{x}, \bar{y}) one epoch after the primal-dual gap has been reduced from the previous change of (\bar{x}, \bar{y}) .

Proof of Theorem 3 using monotone operators. We consider $\tau \geq 0$, and we consider the following algorithm, which is the transposition of the algorithm presented above. We consider a mini-batch $m = 1$ for simplicity. We consider a set of SVRG epochs, where \bar{z} remains fixed. These epochs are initialized by $\tilde{z} = \bar{z}$.

For each SVRG epoch, given z and \tilde{z} , and starting from $z = \tilde{z}$, we run t iterations of:

$$z \leftarrow (I + \sigma(\tau I + A))^{-1} \left(z - \sigma [B\tilde{z} + \frac{1}{\pi_i} (B_i z - B_i \tilde{z}) - \tau \bar{z}] \right),$$

Algorithm 3 Accelerated Stochastic Variance Reduction for Saddle Points

Input: Functions $(K_i)_i$, M , probabilities $(\pi_i)_i$, smoothness $\bar{L}(\pi)$ and L , iterate (x, y)

number of epochs v , number of updates per iteration m , **acceleration factor** τ

Set $\sigma = [L^2 + 3\bar{L}^2/m]^{-1}$ and $(\bar{x}, \bar{y}) = (x, y)$

for $u = 1$ to v **do**

If $u = 0 \pmod{[2 + 2\log(1 + \tau)/(\log 4/3)]}$, **set** $(\bar{x}, \bar{y}) = (\bar{x}, \bar{y})$

Initialize $(\tilde{x}, \tilde{y}) = (x, y)$ and compute $B(\tilde{x}, \tilde{y})$

for $k = 1$ to $\log 4 \times (L^2 + 3\bar{L}^2/m)(1 + \tau)^2$ **do**

Sample $i_1, \dots, i_m \in \mathcal{J}$ from probability vector $(\pi_i)_i$ with replacement

$z \leftarrow (x, y) + \sigma\tau(1 + \tau)(\bar{x}, \bar{y}) - \sigma(1 + \tau) \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1/\gamma \end{pmatrix} (B(\tilde{x}, \tilde{y}) + \{\frac{1}{\pi_i}\tilde{B}_i(x, y) - \frac{1}{\pi_i}\tilde{B}_i(\tilde{x}, \tilde{y})\})$

$(x, y) \leftarrow \text{prox}_M^{\sigma(1+\tau)/(1+\sigma\tau(1+\tau))}(z/(1 + \sigma\tau(1 + \tau)))$

end for

end for

Output: Approximate solution (x, y)

and then update \tilde{z} as z at the end of the SVRG epoch. It corresponds exactly to running the SVRG algorithm to find $(\tau I + A + B)^{-1}(\tau \bar{z})$ approximately, we know from the proof of Theorem 1 that after $\log 4(1 + \frac{L^2}{\mu^2(1+\tau)^2} + \frac{\bar{L}^2}{\mu^2(1+\tau)^2})$ iterations, we have an iterate z such that $\mathbb{E}\|z - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2 \leq \frac{3}{4}\mathbb{E}\|\tilde{z} - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2$. Thus, if we run s epochs where we update \tilde{z} (but not \bar{z}) at each start of epoch, we get an iterate z such that $\mathbb{E}\|z - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2 \leq (\frac{3}{4})^s \mathbb{E}\|\tilde{z} - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2$, and thus

$$\begin{aligned} \mathbb{E}\|z - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2 &\leq \left(\frac{3}{4}\right)^s \mathbb{E}\|\tilde{z} - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2 \\ &= \left(\frac{3}{4}\right)^s \mathbb{E}\|\tilde{z} - z_* - (\tau I + A + B)^{-1}(\tau \bar{z}) + (\tau I + A + B)^{-1}(\tau z_*)\|^2 \\ &\quad \text{using } z_* = (\tau I + A + B)^{-1}(\tau z_*), \\ &= \left(\frac{3}{4}\right)^s \mathbb{E}\|\tilde{z} - z_* - (I + \tau^{-1}(A + B))^{-1}(\bar{z}) + (I + \tau^{-1}(A + B))^{-1}(z_*)\|^2. \end{aligned}$$

We may now use the fact that for any multi-valued maximal monotone operator C , $I - (I + C)^{-1} = (I + C^{-1})^{-1}$, which shows that $I - (I + C)^{-1}$ is 1-Lipschitz-continuous. Thus

$$\mathbb{E}\|z - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2 \leq \left(\frac{3}{4}\right)^s \mathbb{E}\|\tilde{z} - z_*\|^2.$$

This implies, by Minkowski's inequality,

$$\begin{aligned} (\mathbb{E}\|z - z_*\|^2)^{1/2} &\leq (\mathbb{E}\|z - (\tau I + A + B)^{-1}(\tau \bar{z})\|^2)^{1/2} + (\mathbb{E}\|(\tau I + A + B)^{-1}(\tau \bar{z}) - z_*\|^2)^{1/2} \\ &\leq \left(\frac{3}{4}\right)^{s/2} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2} + (\mathbb{E}\|(\tau I + A + B)^{-1}(\tau \bar{z}) - (\tau I + A + B)^{-1}(\tau z_*)\|^2)^{1/2} \\ &= \left(\frac{3}{4}\right)^{s/2} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2} + (\mathbb{E}\|(I + \tau^{-1}(A + B))^{-1}(\bar{z}) - (I + \tau^{-1}(A + B))^{-1}(z_*)\|^2)^{1/2} \\ &\leq \left(\frac{3}{4}\right)^{s/2} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2} + \frac{1}{1 + \tau^{-1}\mu} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2} \\ &= \left(\frac{3}{4}\right)^{s/2} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2} + \frac{\tau}{\tau + \mu} (\mathbb{E}\|\tilde{z} - z_*\|^2)^{1/2}, \end{aligned}$$

using the fact that the contractivity of resolvents of strongly monotone operators. Thus after $s = 2 + 2\frac{\log(1 + \frac{\tau}{\mu})}{\log \frac{4}{3}}$, we get a decrease by $(1 - \frac{\mu}{\tau + \mu})$, and thus the desired result.

D.4 Factored splits and bi-linear models

In the table below, we report the running-time complexity for the factored splits which we used in simulations. Note that SAGA and SVRG then have different bounds. Moreover, all these schemes are adapted when n is close to d . For n much different from d , one could imagine to (a) either complete with zeros or (b) to regroup the data in the larger dimension so that we get as many blocks as for the lower dimension.

Algorithms	Complexity
Stochastic FB-non-uniform	$(1/\varepsilon) \times (\max\{n, d\} \ K\ _F^2 / (\lambda\gamma))$
Stochastic FB-uniform	$(1/\varepsilon) \times (\max\{n, d\}^2 \ K\ _{\max}^2 / (\lambda\gamma))$
SVRG-uniform	$\log(1/\varepsilon) \times (nd + \max\{n, d\}^2 \ K\ _{\max}^2 / (\lambda\gamma))$
SAGA-uniform	$\log(1/\varepsilon) \times (\max\{n, d\}^2 + \max\{n, d\}^2 \ K\ _{\max}^2 / (\lambda\gamma))$
SVRG-non-uniform	$\log(1/\varepsilon) \times (nd + \max\{n, d\} \ K\ _F^2 / (\lambda\gamma))$
SAGA-non-uniform	$\log(1/\varepsilon) \times (\max\{n, d\}^2 + \max\{n, d\} \ K\ _F^2 / (\lambda\gamma))$
SVRG-non-uniform-acc.	$\log(1/\varepsilon) \times (nd + \max\{n, d\}^{3/2} \ K\ _F / \sqrt{\lambda\gamma})$

Table 2: Summary of convergence results for the strongly (λ, γ) -convex-concave bilinear saddle-point problem with matrix K and factored splits, with access to a single row and a single column per iteration. The difference with the individual splits from Table 1 is highlighted in red.

E Surrogate to Area Under the ROC Curve

We consider the following loss function on \mathbb{R}^n , given a vector of positive and negative labels, which corresponds to a convex surrogate to the number of misclassified pairs [13, 14]:

$$\begin{aligned}
\ell(u) &= \frac{1}{2n_+n_-} \sum_{i_+ \in I_+} \sum_{i_- \in I_-} (1 - u_{i_-} + u_{i_+})^2 \\
&= \frac{1}{2n_+n_-} \sum_{i_+ \in I_+} \sum_{i_- \in I_-} \{1 + u_{i_-}^2 + u_{i_+}^2 - 2u_{i_-} + 2u_{i_+} - 2u_{i_-}u_{i_+}\} \\
&= \frac{1}{2} + \frac{1}{n_+} \sum_{i_+ \in I_+} u_{i_+} - \frac{1}{n_-} \sum_{i_- \in I_-} u_{i_-} + \frac{1}{2n_-} \sum_{i_- \in I_-} u_{i_-}^2 + \frac{1}{2n_+} \sum_{i_+ \in I_+} u_{i_+}^2 - \frac{1}{n_+n_-} \sum_{i_+ \in I_+} \sum_{i_- \in I_-} u_{i_-}u_{i_+} \\
&= \frac{1}{2} + \frac{1}{n_+} e_+^\top u - \frac{1}{n_-} e_-^\top u + \frac{1}{2} u^\top \text{Diag}\left(\frac{1}{n_+} e_+ + \frac{1}{n_-} e_-\right) u - \frac{1}{2n_+n_-} u^\top (e_+ e_-^\top + e_- e_+^\top) u \\
&= \frac{1}{2} - a^\top u + \frac{1}{2} u^\top A u,
\end{aligned}$$

with $e_+ \in \mathbb{R}^n$ the indicator vector of I_+ and $e_- \in \mathbb{R}^n$ the indicator vector of I_- . We have $A = \text{Diag}\left(\frac{1}{n_+} e_+ + \frac{1}{n_-} e_-\right) - \frac{1}{n_+n_-} [e_+ e_-^\top + e_- e_+^\top]$ and $a = e_+/n_+ - e_-/n_-$. A short calculation shows that the largest eigenvalue of A is $\frac{1}{M} = \frac{1}{n_+} + \frac{1}{n_-}$.

We consider the function $h(u) = \frac{1}{2} u^\top A u$. It is $(1/M)$ -smooth, its Fenchel conjugate is equal to

$$\frac{1}{2} v^\top A^{-1} v,$$

and our function g will be equal to $v \mapsto \frac{1}{2}v^\top A^{-1}v - \frac{M}{2}\|v\|^2$. Given that $\mathbf{1}$ is a singular vector of A , $g(v)$ is finite only when $v^\top \mathbf{1}_n = 0$.

We need to be able to compute $g(v)$, i.e., solve the system $A^{-1}v$, and to compute the proximal operator

$$\min_v \frac{1}{2}\|v - v_0\|^2 + \sigma g(v) = \min_v \frac{1}{2}\|v - v_0\|^2 + \frac{\sigma}{2}v^\top (A^{-1} - MI)v,$$

which leads to the system: $(A^{-1} - MI + \sigma^{-1}I)v = \sigma^{-1}v_0$, which is equivalent to: $(I - MA + \sigma^{-1}A)v = \sigma^{-1}Av_0$. We thus need to compute efficiently Aw , and $(I + \kappa A)^{-1}w$ with $\kappa > -M$. We have

$$\begin{aligned} I + \kappa A &= \text{Diag}((1 + \kappa/n_+)e_+ + (1 + \kappa/n_-)e_-) - \frac{\kappa}{n_+n_-} [e_+e_-^\top + e_-e_+^\top] \\ &= \text{Diag}((1 + \kappa/n_+)e_+ + (1 + \kappa/n_-)e_-)^{1/2} \\ &\quad \left[I - \frac{\kappa}{n_+n_-} \left(\left[\frac{1}{\sqrt{1 + \kappa/n_+}}e_+ \right] \left[\frac{1}{\sqrt{1 + \kappa/n_-}}e_- \right]^\top - \left[\frac{1}{\sqrt{1 + \kappa/n_-}}e_- \right] \left[\frac{1}{\sqrt{1 + \kappa/n_+}}e_+ \right]^\top \right) \right] \\ &\quad \text{Diag}((1 + \kappa/n_+)e_+ + (1 + \kappa/n_-)e_-)^{1/2} \\ &= D^{1/2}(I - \alpha u_+u_-^\top - \alpha u_-u_+^\top)D^{1/2}, \end{aligned}$$

with $u_+^\top u_- = 0$ and $u_+ = \frac{e_+}{\sqrt{n_+}}$, $u_- = \frac{e_-}{\sqrt{n_-}}$ of norm 1 and $D = \text{Diag}((1 + \kappa/n_+)e_+ + (1 + \kappa/n_-)e_-)$. We have:

$$\begin{aligned} I - \alpha u_+u_-^\top - \alpha u_-u_+^\top &= I - u_+u_+^\top - u_-u_-^\top + (u_+, u_-) \begin{pmatrix} 1 & -\alpha \\ -\alpha & 1 \end{pmatrix} (u_+, u_-)^\top \\ (I - \alpha u_+u_-^\top - \alpha u_-u_+^\top)^{-1} &= I - u_+u_+^\top - u_-u_-^\top + \frac{1}{1 - \alpha^2} (u_+, u_-) \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix} (u_+, u_-)^\top \\ &= I + (1/(1 - \alpha^2) - 1)u_+u_+^\top + (1/(1 - \alpha^2) - 1)u_-u_-^\top + \frac{\alpha}{1 - \alpha^2}(u_+u_-^\top + u_-u_+^\top) \\ &= I + (1/(1 - \alpha^2) - 1)\frac{1}{n_+}e_+e_+^\top + (1/(1 - \alpha^2) - 1)\frac{1}{n_-}e_-e_-^\top \\ &\quad + \frac{\alpha}{1 - \alpha^2} \frac{1}{\sqrt{n_+n_-}}(e_+e_-^\top + e_-e_+^\top). \end{aligned}$$

We have here $\alpha = \frac{\kappa}{n_+n_-} \sqrt{\frac{n_+}{1 + \kappa/n_+}} \sqrt{\frac{n_-}{1 + \kappa/n_-}}$. Thus

$$(I + \kappa A)^{-1} = D^{-1/2} \left[I - u_+u_+^\top - u_-u_-^\top + \frac{1}{1 - \alpha^2} (u_+, u_-) \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix} (u_+, u_-)^\top \right] D^{-1/2},$$

which can be done in $O(n)$.

Moreover, we have

$$\begin{aligned} A &= \text{Diag}((1/n_+)e_+ + (1/n_-)e_-)^{1/2} \\ &\quad \left[I - \frac{1}{n_+n_-} \left(\left[\sqrt{n_+}e_+ \right] \left[\sqrt{n_-}e_- \right]^\top - \left[\sqrt{n_-}e_- \right] \left[\sqrt{n_+}e_+ \right]^\top \right) \right] \\ &\quad \text{Diag}((1/n_+)e_+ + (1/n_-)e_-)^{1/2} \\ &= D^{1/2}(I - u_+u_-^\top - u_-u_+^\top)D^{1/2} \end{aligned}$$

with $u_+^\top u_- = 0$ and u_+, u_- of norm 1. Thus we have

$$\begin{aligned} I - u_+u_-^\top - u_-u_+^\top &= I - u_+u_+^\top - u_-u_-^\top + (u_+, u_-) \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} (u_+, u_-)^\top \\ (I - u_+u_-^\top - u_-u_+^\top)^{-1} &= I - u_+u_+^\top - u_-u_-^\top + \frac{1}{0} (u_+, u_-) \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} (u_+, u_-)^\top. \end{aligned}$$

Thus, if $v^\top \mathbf{1}_n = 0$, we get:

$$v^\top A^{-1}v = v^\top \text{Diag}(n_+e_+ + n_-e_-)v - (v^\top e_+)^2 - (v^\top e_-)^2,$$

which has running-time complexity $O(n)$.

Optimization problem. With a regularizer $f(x) + \frac{\lambda}{2}\|x\|^2$, we obtain the problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^d} \frac{\lambda}{2}\|x\|^2 + f(x) + \frac{1}{2} - a^\top Kx + \frac{1}{2}(Kx)^\top A(Kx) \\ \min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^n} \frac{\lambda}{2}\|x\|^2 + f(x) + \frac{1}{2} - a^\top Kx + y^\top Kx - \frac{M}{2}\|y\|^2 - \frac{1}{2}y^\top (A^{-1} - MI)y, \end{aligned}$$

with $g(y) = \frac{1}{2}y^\top (A^{-1} - MI)y$.

F Additional Experimental Results

We complement the results of the main paper in several ways: (a) by providing all test losses, the distance to optimum $\Omega(x - x_*, y - y_*)$ in log-scale, as well as the primal-dual gaps in log-scale, as a function of the number of passes on the data. We consider the three machine learning settings:

- Figure 1: `sido` dataset, AUC loss and cluster norm (plus squared-norm) regularizer (both non separable).
- Figure 2: `sido` dataset, square loss and ℓ_1 -norm (plus squared-norm) regularizer (both separable).
- Figure 3: `rcv1` dataset, square loss and ℓ_1 -norm (plus squared-norm) regularizer (both separable).

We consider the following methods in all cases (all methods are run with the step-sizes proposed in their respective convergence analysis):

- fb-acc: accelerated forward-backward saddle-point method from Section 2.2,
- fb-sto: stochastic forward-backward saddle-point method from Section 2.3,
- saga: our new algorithm from Section 4, with non-uniform sampling, and sampling of a single row and column per iteration,
- saga (unif): our new algorithm from Section 4, with uniform sampling, and sampling of a single row and column per iteration,
- svrg: our new algorithm from Section 3, with non-uniform sampling, and sampling of a single row and column per iteration,
- svrg-acc: our new accelerated algorithm from Section 3, with non-uniform sampling, and sampling of a single row and column per iteration,
- fba-primal: accelerated proximal method [11], which can be applied to the primal version of our problem (which is the sum of a smooth term and a strongly convex term).

Moreover, for the separable cases, we add:

- saga-primal: SAGA with non-uniform sampling [28], which can only be run with separable losses.

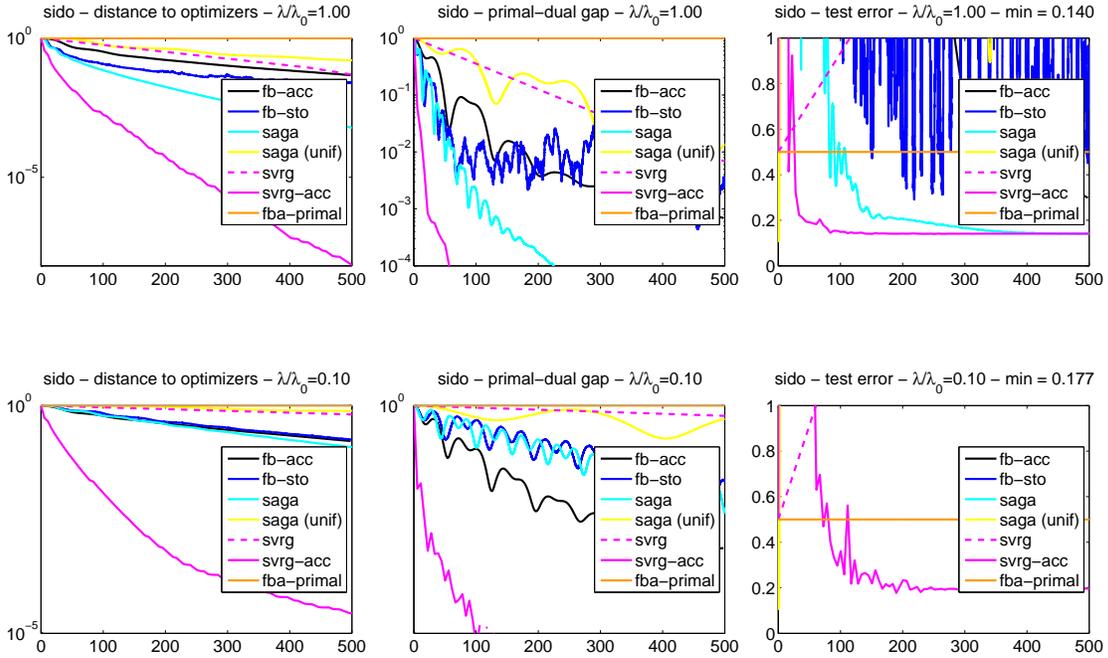


Figure 1: *sido* dataset. Top: $\lambda = \lambda_0 = \|K\|_F^2/n^2$, Bottom: $\lambda = \lambda_0/10 = \frac{1}{10}\|K\|_F^2/n^2$. AUC loss and cluster-norm regularizer. Distances to optimum, primal-dual gaps and test losses, as a function of the number of passes on the data. Note that the primal SAGA (with non-uniform sampling) cannot be used because the loss is not separable. Best seen in color.

We can make the following observations:

- Non-uniform sampling is key to good performance.
- The distance to optimum (left plots) exhibits a clear linear convergence behavior (which is predicted by our analysis), which is not the case for the primal-dual gap, which does converge, but more erratically. It would be interesting to provide bounds for these as well.
- When λ decreases (bottom plots, more ill-conditioned problems), the gains of accelerated methods with respect to non-accelerated ones are unsurprisingly larger. Note that for two out of three settings, the final test loss is smaller for the smaller regularization, and non-accelerated methods need more passes on the data to reach good testing losses.
- Primal methods which are not using separability (here “fba-primal”) can be run on all instances, but are not competitive. Note that in some situations, they achieve early on good performances (e.g., Figure 2), before getting caught up by stochastic-variance-reduced saddle-point techniques (note also that since these are not primal-dual methods, we compute dual candidates through the gradient of the smooth loss functions, which is potentially disadvantageous).
- Primal methods that use separability (here “saga-primal”) cannot be run on non-separable problems, but when they can run, they are still significantly faster than our saddle-point techniques. We believe that this is partly due to adaptivity to strong convexity (the convergence bounds for the two sets of techniques are the same for this problem).

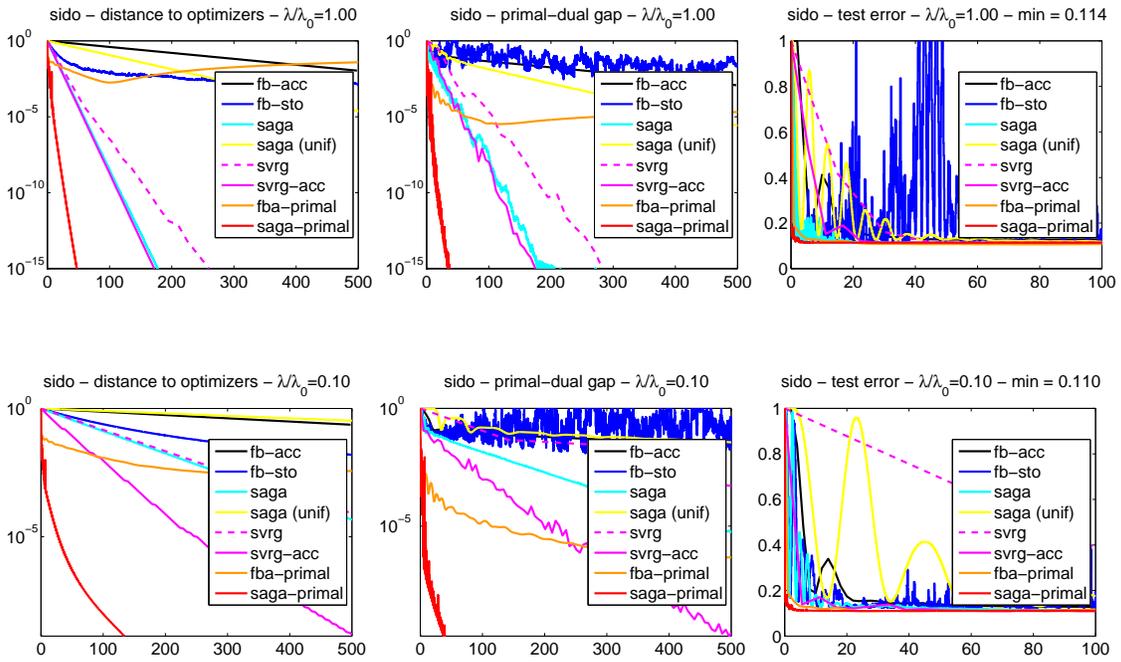


Figure 2: `sido` dataset. Top: $\lambda = \lambda_0 = \|K\|_F^2/n^2$, Bottom: $\lambda = \lambda_0/10 = \frac{1}{10}\|K\|_F^2/n^2$. Squared loss, with ℓ_1 -regularizer. Distances to optimum, primal-dual gaps and test losses, as a function of the number of passes on the data. Note that the primal SAGA (with non-uniform sampling) can only be used because the loss is separable. Best seen in color.

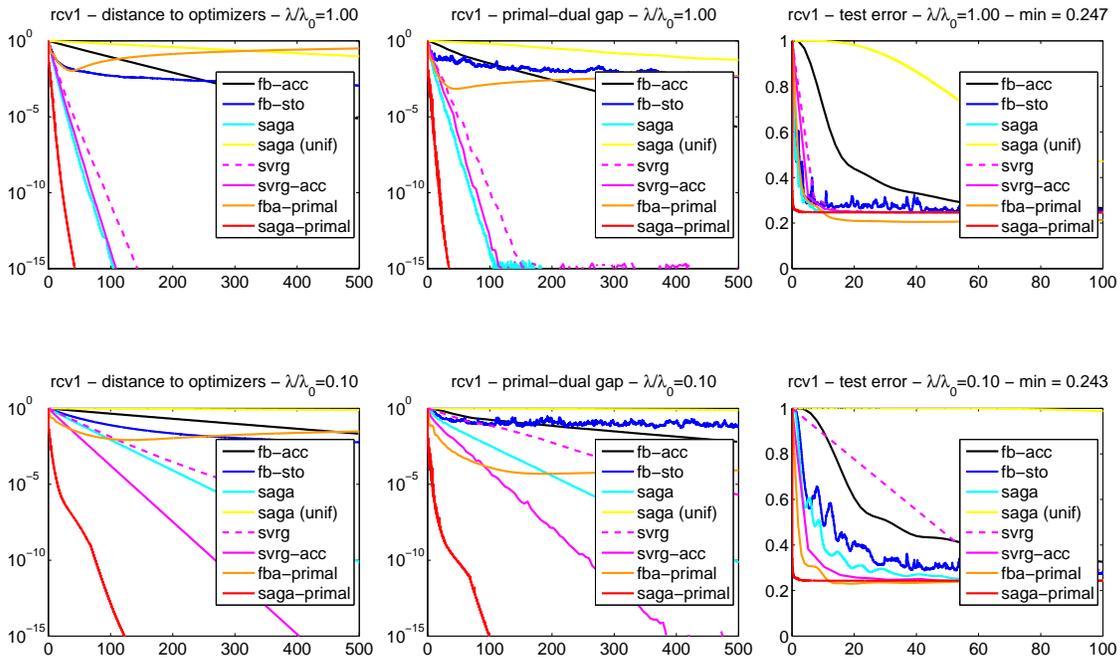


Figure 3: rcv1 dataset. Top: $\lambda = \lambda_0 = \|K\|_F^2/n^2$, Bottom: $\lambda = \lambda_0/10 = \frac{1}{10}\|K\|_F^2/n^2$. Squared loss, with ℓ_1 -regularizer. Distances to optimum, primal-dual gaps and test losses, as a function of the number of passes on the data. Note that the primal SAGA (with non-uniform sampling) can only be used because the loss is separable. Best seen in color.