
Stochastic Primal-Dual Method for Empirical Risk Minimization with $\mathcal{O}(1)$ Per-Iteration Complexity

Conghui Tan*

The Chinese University of Hong Kong
chtan@se.cuhk.edu.hk

Tong Zhang

Tencent AI Lab
tongzhang@tongzhang-ml.org

Shiqian Ma

University of California, Davis
sqma@math.ucdavis.edu

Ji Liu

Tencent AI Lab, University of Rochester
ji.liu.uwisc@gmail.com

Abstract

Regularized empirical risk minimization problem with linear predictor appears frequently in machine learning. In this paper, we propose a new stochastic primal-dual method to solve this class of problems. Different from existing methods, our proposed methods only require $\mathcal{O}(1)$ operations in each iteration. We also develop a variance-reduction variant of the algorithm that converges linearly. Numerical experiments suggest that our methods are faster than existing ones such as proximal SGD, SVRG and SAGA on high-dimensional problems.

1 Introduction

In this paper, we consider the convex regularized empirical risk minimization with linear predictors:

$$\min_{x \in X} \left\{ P(x) \triangleq \frac{1}{n} \sum_{i=1}^n \phi_i(a_i^\top x) + g(x) \right\}, \quad (1)$$

where $X \subset \mathbb{R}^d$ is a convex closed feasible set, $a_i \in \mathbb{R}^d$ is the i -th data sample, ϕ_i is its corresponding convex closed loss function, and $g(x) : X \rightarrow \mathbb{R}$ is a convex closed regularizer for model parameter x . Here we assume the feasible set X and the regularizer function $g(x)$ are both separable, i.e.,

$$X = X_1 \times \cdots \times X_d \quad \text{and} \quad g(x) = \sum_{j=1}^d g_j(x_j). \quad (2)$$

Problem (1) with structure (2) generalizes many well-known classification and regression problems. For example, support vector machine is with this form by choosing $\phi_i(u) = \max\{0, 1 - b_i u\}$ and $g(x) = \frac{\lambda}{2} \|x\|_2^2$. Other examples include ℓ_1 logistic regression, ℓ_2 logistic regression, and LASSO. One popular choice for solving (1) is the proximal stochastic gradient descent method (PSGD). In each iteration of PSGD, an index i is randomly sampled from $\{1, 2, \dots, n\}$, and then the iterates are updated using only the information of a_i and ϕ_i . As a result, the per-iteration cost of PSGD is $\mathcal{O}(d)$ and independent of n .

It is well known that PSGD converges at a sub-linear rate [9] even for strongly convex problems, due to non-diminishing variance of the stochastic gradients. One line of research tried is dedicated to improve the convergence rate of PSGD by utilizing the finite sum structure in (1). Some representative

*This work was done while Conghui Tan was a research intern at Tencent AI lab.

works include SVRG [5, 16], SDCA [13, 14], SAGA [4] and SPDC [18]. All these accelerated variants enjoy linear convergence when $g(x)$ is strongly convex and all ϕ_i 's are smooth.

Since all of these algorithms need to sample at least one data a_i in each iteration (so their per-iteration cost is at least $\mathcal{O}(d)$), their potential drawbacks include: 1) they are not suitable for the distributed learning with features distributed; 2) it may incur heavy computation per iteration in the high-dimensional case, i.e., when d is very large.

Our contributions. In this paper, we explore the possibility of accelerating PSGD by making each iteration more light-weighted: only one coordinate of one data is sampled, i.e., one entry a_{ij} , in each iteration of the algorithm. This leads to a new algorithm, named SPD1 (stochastic primal-dual method with $\mathcal{O}(1)$ per-iteration complexity), whose per-iteration cost is only $\mathcal{O}(1)$. We prove that the convergence rate of the new method is $\mathcal{O}(1/\sqrt{t})$ for convex problems and $\mathcal{O}(\ln t/t)$ for strongly convex and smooth problems, where t is the iteration counter. Moreover, the overall computational cost is the same as PSGD in high-dimensional settings. Therefore, we managed to reduce the per-iteration complexity from $\mathcal{O}(d)$ to $\mathcal{O}(1)$ while keep the total computational cost at the same order. Furthermore, by incorporating the variance reduction technique, we develop a variant of SPD1, named SPD1-VR, that converges linearly for strongly convex problems. Comparing with existing methods, our SPD1 and SPD1-VR algorithms are more suitable for distributed systems by allowing the flexibility of either feature distributed or data distributed. An additional advantage of our $\mathcal{O}(1)$ per-iteration complexity algorithms is that they are more favorable by asynchronous parallelization and bring more speedup since they admit much better tolerance to the staleness caused by asynchrony. Our numerical tests indicate that our methods are faster than both PSGD and SVRG on high-dimensional problems, even in single-machine setting.

We notice that [6] and [17] used similar ideas in the sense that in each iteration only one coordinate of the iterate is updated using one sampled data. However, we need to point out that their algorithms still need to sample the full vector a_i to compute the directional gradient, and thus the per-iteration cost is still $\mathcal{O}(d)$.

1.1 Notation

We use $A \in \mathbb{R}^{n \times d}$ to denote the data matrix, whose rows are denoted by $a_i, i = 1, 2, \dots, n$. We use a_{ij} to denote the j -th entry of a_i . $[n]$ denotes the set $\{1, 2, \dots, n\}$. x^t denotes the iterate in the t -th iteration and x_j^t is its j -th entry. We always use $\|w\|$ to denote the ℓ_2 norm of w unless otherwise specified.

For function $f(z)$ whose domain is Z , its proximal mapping is defined as

$$\text{prox}_f(z) = \arg \min_{z' \in Z} \left\{ f(z') + \frac{1}{2} \|z' - z\|^2 \right\}. \quad (3)$$

We use $\partial f(z)$ to denote the subdifferential of f at point z . f is said to be μ -strongly convex ($\mu > 0$) if

$$f(z') \geq f(z) + s^\top (z' - z) + \frac{\mu}{2} \|z' - z\|^2 \quad \forall s \in \partial f(z), \forall z, z' \in Z.$$

The conjugate function of $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\phi_i^*(y) = \sup_{x \in \mathbb{R}} \{y \cdot x - \phi_i(x)\}. \quad (4)$$

Function ϕ_i is L -Lipschitz continuous if

$$|\phi_i(x) - \phi_i(y)| \leq L|x - y|, \quad \forall x, y \in \mathbb{R},$$

which is equivalent to

$$|s| \leq L, \quad \forall s \in \partial \phi_i(x), \forall x.$$

ϕ_i is said to be $(1/\gamma)$ -smooth if it is differentiable and its derivative is $(1/\gamma)$ -Lipschitz continuous.

We use $R \triangleq \max_{i \in [n]} \|a_i\|$ to denote the maximum row norm of A , and $R' \triangleq \max_{j \in [d]} \|a'_j\|$ to denote the maximum column norm of A , where a'_1, \dots, a'_d are the columns of matrix A .

2 Stochastic Primal-Dual Method with $\mathcal{O}(1)$ Per-Iteration Cost

Our algorithm solves the following equivalent primal-dual reformulation of (1):

$$\min_{x \in X} \max_{y \in Y} \left\{ F(x, y) \triangleq \frac{1}{n} y^\top A x - \frac{1}{n} \sum_{i=1}^n \phi_i^*(y_i) + g(x) \right\}, \quad (5)$$

where $Y = Y_1 \times \cdots \times Y_n$ and $Y_i = \{y_i \in \mathbb{R}^n \mid \phi_i^*(y_i) < \infty\}$ is the dual feasible set resulted by the conjugate function ϕ_i^* . For example, when ϕ_i is L -Lipschitz continuous, we have $Y_i \subset [-L, L]$ (see Lemma 4 in the appendix.)

Our SPD1 algorithm for solving (5) is presented in Algorithm 1. In each iteration of the algorithm, only one coordinate of x and one coordinate of y are updated by using a randomly sampled data entry $a_{i_t j_t}$. Therefore, SPD1 only requires $\mathcal{O}(1)$ time per iteration. Because of this, SPD1 is a kind of randomized coordinate descent method.

Algorithm 1 Stochastic Primal-Dual Method with $\mathcal{O}(1)$ Per-Iteration Cost (SPD1)

Parameters: primal step sizes $\{\eta_t\}$, dual step sizes $\{\tau_t\}$

Initialize $x^0 = \arg \min_{x \in X} g(x)$ and $y_i^0 = \arg \min_{y_i \in Y_i} \phi_i^*(y_i)$ for all $i \in [n]$

for $t = 0, 1, \dots, T - 1$ **do**

 Randomly sample $i_t \in [n]$ and $j_t \in [d]$ independently and uniformly

$$x_j^{t+1} = \begin{cases} \text{prox}_{\eta_t g_j} (x_j^t - \eta_t \cdot a_{i_t j_t} y_{i_t}^t) & \text{if } j = j_t \\ x_j^t & \text{if } j \neq j_t \end{cases}$$

$$y_i^{t+1} = \begin{cases} \text{prox}_{(\tau_t/d)\phi_i^*} (y_i^t + \tau_t \cdot a_{i_t j_t} x_{j_t}^t) & \text{if } i = i_t \\ y_i^t & \text{if } i \neq i_t \end{cases}$$

end for

Output: $\hat{x}^T = \frac{1}{T} \sum_{t=0}^{T-1} x^t$ and $\hat{y}^T = \frac{1}{T} \sum_{t=0}^{T-1} y^t$

The intuition of this algorithm is as follows. One can rewrite (5) into:

$$\min_{x \in X} \max_{y \in Y} \left\{ F(x, y) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^d \left[a_{ij} y_i x_j - \frac{1}{d} \phi_i^*(y_i) + g_j(x_j) \right] \right\},$$

which has a two-layer finite-sum structure. Then Algorithm 1 can be viewed as a primal-dual version of stochastic gradient descent (SGD) on this finite-sum problem, which samples a pair of induces (i_t, j_t) and then only utilizes the corresponding summand to do updates in each iteration. Hence, one can view SPD1 as a combination of randomized coordinate descent method and stochastic gradient descent method applied to the primal-dual reformulation (5).

Note that in the initialization stage, we need to minimize $g(x)$ and $\phi_i^*(y_i)$. Since we assume the proximal mappings of $g(x)$ and ϕ_i^* are easy to solve, these two direct minimization problems should be even easier, and thus would not bring any trouble in implementation of this algorithm. For example, when $g(x) = \lambda \|x\|_1$, it is well known its proximal mapping is the soft thresholding operator. While its direct minimizer, namely $x^0 = \min_x g(x)$, is simply $x^0 = 0$.

As a final remark, we point out that the primal-dual reformulation (5) is a convex-concave bilinear saddle point problem (SSP). This problem has drawn a lot of research attentions recently. For example, Chambolle and Pock developed an efficient primal-dual gradient method for solving bilinear SSP in [2], which has an accelerated rate in certain circumstances. Besides, in [3], Dang and Lan proposed a randomized algorithm for solving (5). However, their algorithm needs to use the full-dimensional gradient in each iteration, so the per-iteration cost is much higher than SPD1.

3 SPD1 with Variance Reduction

As we have discussed above, SPD1 has some close connection to SGD. Hence, we can incorporate the variance reduction technique [5] to reduce the variance of the stochastic gradients so that to improve the convergence rate of SPD1. This new algorithm, named SPD1-VR, is presented in Algorithm 2. Similar to SVRG [5], SPD1-VR has a two-loop structure. In the outer loop, the snapshots of the full

gradients are computed for both \tilde{x}^k and \tilde{y}^k . In the inner loop, the updates are similar to Algorithm 1, but the stochastic gradient is replaced by its variance reduced version. That is, $a_{i_t j_t} y_{i_t}^t$ is replaced by $a_{i_t j_t} (y_{i_t}^t - \tilde{y}_{i_t}^k) + G_{x, j_t}^k$, where G_{x, j_t}^k is the j_t -th coordinate of the latest snapshot of full gradient. This variance reduced stochastic gradient is still an unbiased estimator of the full gradient along direction x_{j_t} , i.e.,

$$\mathbb{E}_{i_t} [a_{i_t j_t} (y_{i_t}^t - \tilde{y}_{i_t}^k) + G_{x, j_t}^k] = \mathbb{E}_{i_t} [a_{i_t j_t} y_{i_t}^t] = \frac{1}{n} \sum_{i=1}^n a_{i j_t} y_i^t.$$

Because of the variance reduction technique, fixed step sizes η and τ can be used instead of diminishing ones.

Algorithm 2 SPD1 with Variance Reduction (SPD1-VR)

Parameters: primal step size η , dual step size τ

Initialize $\tilde{x}^0 \in X$ and $\tilde{y}^0 \in Y$

for $k = 0, 1, \dots, K - 1$ **do**

 Compute full gradients $G_x^k = (1/n)A^\top \tilde{y}^k$ and $G_y^k = (1/d)A\tilde{x}^k$

 Let $(x^0, y^0) = (\tilde{x}^k, \tilde{y}^k)$

for $t = 0, 1, \dots, T - 1$ **do**

 Randomly sample $i_t, i'_t \in [n]$ and $j_t, j'_t \in [d]$ independently

$$\tilde{x}_j^t = \begin{cases} \text{prox}_{\eta_t g_j} [x_j^t - \eta \cdot (a_{i'_t j} (y_{i'_t}^t - \tilde{y}_{i'_t}^k) + G_{x, j}^k)] & \text{if } j = j_t \\ x_j^t & \text{if } j \neq j_t \end{cases}$$

$$\tilde{y}_i^t = \begin{cases} \text{prox}_{(\tau_t/d)\phi_i^*} [y_i^t + \tau \cdot (a_{i j'_t} (x_{j'_t}^t - \tilde{x}_{j'_t}^k)) + G_{y, i}^k] & \text{if } i = i_t \\ y_i^t & \text{if } i \neq i_t \end{cases}$$

$$x_j^{t+1} = \begin{cases} \text{prox}_{\eta_t g_j} [x_j^t - \eta \cdot (a_{i_t j} (\tilde{y}_{i_t}^t - \tilde{y}_{i_t}^k) + G_{x, j}^k)] & \text{if } j = j_t \\ x_j^t & \text{if } j \neq j_t \end{cases}$$

$$y_i^{t+1} = \begin{cases} \text{prox}_{(\tau_t/d)\phi_i^*} [y_i^t + \tau \cdot (a_{i j_t} (\tilde{x}_{j_t}^t - \tilde{x}_{j_t}^k) + G_{y, i}^k)] & \text{if } i = i_t \\ y_i^t & \text{if } i \neq i_t \end{cases}$$

end for

 Set $(\tilde{x}^{k+1}, \tilde{y}^{k+1}) = (x^T, y^T)$

end for

Output: \tilde{x}^K and \tilde{y}^K

Besides the variance reduction technique, another crucial difference between SPD1 and SPD1-VR is that the latter is in fact an extragradient method [7]. Note that each iteration of the inner loop of SPD1-VR consists two gradient steps: the first step is a normal gradient descent/ascent, while in the second step, it starts from x^t and y^t but uses the gradient estimations at $(\tilde{x}^t, \tilde{y}^t)$. For saddle point problems, extragradient method has stronger convergence guarantees than simple gradient methods [8]. Moreover, in each iteration of SPD1-VR, two independent pairs of random indices (i_t, j_t) and (i'_t, j'_t) are drawn. This is because two stochastic gradients are needed for the extragradient framework. Similar to the classical analysis of stochastic algorithms, we need the stochastic gradients to be independent. However, when updating \tilde{x}^t and x^{t+1} , we choose the same coordinate j_t , so the independence property is only required for two directional stochastic gradients along coordinate j_t .

We note that every iteration of the inner loop only involves $\mathcal{O}(1)$ operations in SPD1-VR. Full gradients are computed in each outer loop, whose computational cost is $\mathcal{O}(nd)$.

Finally, we have to mention that [12] also developed a variance-reduction method for solving convex-concave saddle point problems, which is related to Algorithm 2. However, except for the common variance reduction ideas used in both methods, the method in [12] and SPD1-VR are quite different. First, there is no coordinate descent counterpart in their method, so the per-iteration cost is much higher than our method. Second, their method is a gradient method instead of extragradient method like SPD1-VR. Third, their method has quadratic dependence on the problem condition number unless extra acceleration technique is combined, while our method depends only linearly on condition number as shown in Section 4.

4 Iteration Complexity Analysis

4.1 Iteration Complexity of SPD1

In this subsection, we analyze the convergence rate of SPD1 (Algorithm 1). We measure the optimality of the solution by primal-dual gap, which is defined as

$$\mathcal{G}(\hat{x}^T, \hat{y}^T) \triangleq \sup_{y \in Y} F(\hat{x}^T, y) - \inf_{x \in X} F(x, \hat{y}^T).$$

Note that primal-dual gap equals 0 if and only if (\hat{x}^T, \hat{y}^T) is a pair of primal-dual optimal solutions to problem (5). Besides, primal-dual gap is always an upper bound of primal sub-optimality:

$$\begin{aligned} \mathcal{G}(\hat{x}^T, \hat{y}^T) &\geq \sup_{y \in Y} F(\hat{x}^T, y) - \sup_{y \in Y} \inf_{x \in X} F(x, y) \\ &\geq \sup_{y \in Y} F(\hat{x}^T, y) - \inf_{x \in X} \sup_{y \in Y} F(x, y) \\ &= P(\hat{x}^T) - \inf_{x \in X} P(x). \end{aligned}$$

Our main result for the iteration complexity of SPD1 is summarized in Theorem 1.

Theorem 1. *Assume each ϕ_i is L -Lipschitz continuous, and the primal feasible set X is bounded, i.e.,*

$$D \triangleq \sup_{x \in X} \|x\| < \infty.$$

If we choose the step sizes in SPD1 as

$$\eta_t = \frac{\sqrt{2d}D}{LR\sqrt{t+1}} \quad \text{and} \quad \tau_t = \frac{\sqrt{2d}Ln}{DR'\sqrt{t+1}},$$

then we have the following convergence rate for SPD1:

$$\mathbb{E} [\mathcal{G}(\hat{x}^T, \hat{y}^T)] \leq \frac{\sqrt{2d}LD \cdot (R + R')}{\sqrt{T}}. \quad (6)$$

Note that when the problem is high-dimensional, i.e., $d \geq n$, it usually holds that $R \geq R'$. In this case, Theorem 1 implies that SPD1 requires

$$\mathcal{O}\left(\frac{dL^2D^2R^2}{\epsilon^2}\right) \quad (7)$$

iterations to ensure that the primal-dual gap is smaller than ϵ .

Under the same assumptions, if we directly apply the classical result by Nemirovski et al. [9] for PSGD on the primal problem (1), the number of iterations needed by PSGD is

$$\mathcal{O}\left(\frac{L^2D^2R^2}{\epsilon^2}\right),$$

in order to reduce the primal sub-optimality to be smaller than ϵ . Considering that each iteration of PSGD costs $\mathcal{O}(d)$ computation, its overall complexity is actually the same as (7).

If we further impose strong convexity and smoothness assumptions, we get an improved iteration complexity shown in Theorem 2.

Theorem 2. *We assume the same assumptions as in Theorem 1. Moreover, we assume that $g(x)$ is μ -strongly convex ($\mu > 0$), and all ϕ_i are $(1/\gamma)$ -smooth ($\gamma > 0$). If the step sizes in SPD1 are chosen as*

$$\eta_t = \frac{2}{\mu(t+4)} \quad \text{and} \quad \tau_t = \frac{2nd}{\gamma(t+4)},$$

we have the following convergence rate for SPD1:

$$\mathbb{E} [\mathcal{G}(\hat{x}^T, \hat{y}^T)] \leq \frac{4dD^2\mu + 4L^2\gamma + 2L^2R^2/\mu \cdot \ln(T+4) + 2dD^2R'^2/\gamma \cdot \ln(T+4)}{T}. \quad (8)$$

Comparing with the classical convergence rate result of PSGD, we note that the convergence rate of SPD1 not only depends on μ , but also depends on the dual strong convexity parameter γ . This is reasonable because SPD1 has stochastic updates for both primal and dual variables, and this is why $\gamma > 0$ is necessary for ensuring the $\mathcal{O}(\ln T/T)$ convergence rate. Furthermore, we believe that the factor $\ln T$ is removable in (8) so $\mathcal{O}(1/T)$ convergence rate can be obtained, by applying more sophisticated analysis technique such as optimal averaging [15]. We do not dig into this to keep the paper more succinct.

4.2 Iteration Complexity of SPD1-VR

In this subsection, we analyze the iteration complexity of SPD1-VR (Algorithm 2). Before stating the main result, we first introduce the notion of condition number. When each ϕ_i is $(1/\gamma)$ -smooth, and $g(x)$ is μ -strongly convex, the condition number of the primal problem (1) defined in the literature is (see, e.g., [18]):

$$\kappa = \frac{R^2}{\mu\gamma}.$$

Here we also define another condition number:

$$\kappa' = \frac{dR'^2}{n\mu\gamma}.$$

Since R is the maximum row norm and R' is the maximum column norm of data matrix $A \in \mathbb{R}^{n \times d}$, usually R^2 and $(d/n)R'^2$ should be in the same order, which means $\kappa' \approx \kappa$. Without loss of generality, we assume that $\kappa \geq 1$ and $\kappa' \geq 1$.

Theorem 3. *Assume each ϕ_i is $(1/\gamma)$ -smooth ($\gamma > 0$), and $g(x)$ is μ -strongly convex ($\mu > 0$). If we choose the step sizes in SPD1-VR as*

$$\eta = \frac{\gamma}{128R^2} \cdot \min \left\{ \frac{d\kappa}{n\kappa'}, 1 \right\} \quad \text{and} \quad \tau = \frac{n\mu}{128R'^2} \cdot \min \left\{ \frac{n\kappa'}{d\kappa}, 1 \right\}, \quad (9)$$

and let $T \geq c \cdot \max\{d\kappa, n\kappa'\}$ for some uniform constant $c > 0$ independent of problem setting, SPD1-VR converges linearly in expectation:

$$\mathbb{E} [\Delta_K] \leq \left(\frac{3}{5} \right)^K \cdot \Delta_0,$$

where

$$\Delta_k = \frac{\|\tilde{x}^k - x^*\|^2}{\eta} + \frac{\|\tilde{y}^k - y^*\|^2}{\tau}$$

and (x^*, y^*) is a pair of primal-dual optimal solution to (5).

This theorem implies that we need $\mathcal{O}(\log(1/\epsilon))$ outer loops, or $\mathcal{O}(\max\{d\kappa, n\kappa'\} \log(1/\epsilon))$ inner loops to ensure $\mathbb{E} [\Delta_k] \leq \epsilon$. Considering that there are $\Theta(nd)$ extra computation cost for computing the full gradient in every outer loop, the total complexity of this algorithm is

$$\mathcal{O} \left((nd + \max\{d\kappa, n\kappa'\}) \log \frac{1}{\epsilon} \right). \quad (10)$$

As a comparison, the complexity of SVRG under same setting is [16]:

$$\mathcal{O} \left(d(n + \kappa) \log \frac{1}{\epsilon} \right), \quad (11)$$

Since usually it holds $\kappa' \approx \kappa$, $d\kappa$ will dominate $n\kappa'$ when $d > n$. In this case, the two complexity bounds (10) and (11) are the same.

Although the theoretical complexity of SPD1-VR is same as SVRG when $d \geq n$, we empirically found that SPD1-VR is significantly faster than SVRG in high-dimensional problems (see Section 5), by allowing much larger step sizes than the ones in (9) suggested by theory. We conjecture that this is due to the power of coordinate descent. Nesterov's seminal work [11] has rigorously proved that coordinate descent can reduce the Lipschitz constant of the problem, and thus allows larger step




Figure 1: Experimental results on synthetic data. n is set to 1000 in all figures, while d varies from 100 to 10000. λ is fixed as $\lambda = 10^{-3}$. The y -axis here is the primal sub-optimality, namely $P(x^t) - P(x^*)$.

sizes than gradient descent. However, due to the sophisticated coupling of primal and dual variable updates in our algorithm, our analysis is currently unable to reflect this property.

We point out that the existing accelerated algorithms such as SPDC [18] and Katyusha [1] have better complexity given by

$$\mathcal{O}\left(d(n + \sqrt{n\kappa}) \log \frac{1}{\epsilon}\right). \quad (12)$$

These accelerated algorithms employ Nesterov’s extrapolation techniques [10] to accelerate the algorithms. We believe that it is also possible to incorporate the same technique to further accelerate SPD1-VR, but we leave this as a future work at this moment.

5 Experiments

In this section, we conduct numerical experiments of our proposed algorithms. Due to space limitation, we only present part of the results here, more experiments can be found in appendix.

Here we consider solving a classification problem with logistic loss function

$$\phi_i(u) = \log(1 + \exp\{-b_i u\}),$$

where $b_i \in \{\pm 1\}$ is the class label. Note that this loss function is smooth. Although this ϕ_i^* does not admit closed-form solution for its proximal mapping, following [13], we apply Newton’s method to compute its proximal mapping, which can achieve very high accuracy in very few (say, 5) steps. Since the proximal sub-problem of ϕ_i^* is a 1-dimensional optimization problem, using Newton’s method here is actually quite cheap. Besides, We use $g(x) = \frac{\lambda}{2}\|x\|^2$ as the regularizer.

We compare our SPD1 and SPD1-VR with some standard stochastic algorithms for solving (1), including PSGD, SVRG and SAGA. We always set $T = nd$ for SPD1-VR and $T = n$ for SVRG, where T is the number of inner loops in each outer loop.

5.1 Results on Synthetic Data

Since our theory in Section 4 suggest that the performance of our proposed methods relies on the relationship between n and d . Here we will test our methods on synthetic dataset with different n and d to see their effects to the performance. To generate the data, we first randomly sample the data matrix A and a vector $\bar{x} \in \mathbb{R}^d$ with entries i.i.d. drawn from $\mathcal{N}(0, 1)$, and then generate the labels as

$$b_i = \text{sign}(a_i^\top \bar{x} + \varepsilon_i), \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2),$$

for some constant $\sigma^2 > 0$. Since the focus here is the relationship between n and d , in order to simplify the experiments, we fix n as $n = 1000$, but vary the value of d with values chosen from $d \in \{100, 1000, 10000\}$.

The results are presented in Figure 1. When $d = 100 < n$, it is clear that SPD1 is slower than PSGD, and SPD1-VR is also inferior than both SVRG and SAGA. While when $d = 1000 = n$, even though SPD1 falls behind PSGD at the beginning, their final performance is quite close at last, and SPD1-VR

Table 1: Summary of datasets

Dataset	n	d	λ
colon-cancer	62	2,000	1
gisette	6,000	5,000	10^{-2}
rcv1.binary	20,242	47,236	10^{-3}




Figure 2: Numerical results on three real datasets. The y -axis is also primal sub-optimality.

begins to beat both SVRG and SAGA. Finally, when $d > n$, SPD1 becomes obviously faster than PSGD, and SPD1-VR is also significantly better than SVRG and SAGA. This indicates that our algorithms SPD1 and SPD1-VR are preferable in practice for high-dimensional problems.

5.2 Results on Real Datasets

In this part, we will demonstrate the efficiency of our proposed methods on real datasets. Here we only focus on the high-dimensional case where $d > n$ or $d \approx n$.

We will test all the algorithms on three real datasets: `colon-cancer`, `gisette` and `rcv1.binary`, downloaded from the LIBSVM website². The attributes of these data and λ used for each dataset are summarized in Table 1.

The experimental results on these real datasets are shown in Figure 2. For `colon-cancer` dataset, where d is much larger than n , the performance of SPD1-VR is really dominating over other methods, and SPD1 also performs better than PSGD. For `gisette` dataset, where n is slightly larger than d , SPD1-VR still outperforms all other competitors, but this time SPD1 is slower than PSGD. Besides, for `rcv1.binary`, both SPD1 and SPD1-VR are better than PSGD and SVRG/SAGA respectively.

These results on real datasets further confirm that our proposed methods, especially SPD1-VR, are faster than existing algorithms on high-dimensional problems.

6 Conclusion

In this paper, we developed two stochastic primal-dual algorithms, named SPD1 and SPD1-VR for solving regularized empirical risk minimization problems. Different from existing methods, our proposed algorithms have a brand-new updating style, which only need to use one coordinate of one sampled data in each iteration. As a result, the per-iteration cost is very low and the algorithms are very suitable for distributed systems. We proved that the overall convergence property of SPD1 and SPD1-VR resembles PSGD and SVRG respectively under certain condition, and empirically showed that they are faster than existing methods such as PSGD, SVRG and SAGA in high-dimensional settings. We believe that our new methods have great potential to be used in large-scale distributed optimization applications.

²www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets

References

- [1] Zeyuan Allen-Zhu and Elad Hazan. Variance reduction for faster non-convex optimization. In *International Conference on Machine Learning*, pages 699–707, 2016.
- [2] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- [3] Cong Dang and Guanghui Lan. Randomized first-order methods for saddle point optimization. <https://arxiv.org/abs/1409.8625>, 2014.
- [4] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in neural information processing systems*, pages 1646–1654, 2014.
- [5] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.
- [6] Jakub Konečný, Zheng Qu, and Peter Richtárik. Semi-stochastic coordinate descent. *Optimization Methods and Software*, 32(5):993–1005, 2017.
- [7] GM Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- [8] Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [9] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- [10] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*. Applied Optimization. Kluwer Academic Publishers, Boston, MA, 2004.
- [11] Yurii Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [12] Balamurugan Palaniappan and Francis Bach. Stochastic variance reduction methods for saddle-point problems. In *Advances in Neural Information Processing Systems*, pages 1416–1424, 2016.
- [13] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.
- [14] Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *International Conference on Machine Learning*, pages 64–72, 2014.
- [15] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, pages 71–79, 2013.
- [16] Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- [17] Adams Wei Yu, Qihang Lin, and Tianbao Yang. Doubly stochastic primal-dual coordinate method for empirical risk minimization and bilinear saddle-point problem. *arXiv preprint arXiv:1508.03390*, 2015.
- [18] Yuchen Zhang and Lin Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. *The Journal of Machine Learning Research*, 18(1):2939–2980, 2017.

Appendix A Proofs

Lemma 4. *If $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is L -Lipschitz continuous, then $\{y \in \mathbb{R} | \phi^*(y) < \infty\} \subseteq [-L, L]$.*

Proof. In this proof, we will first show that if $y > L$, then $\phi^*(y) = \infty$. By the Lipschitz continuity, we have

$$\phi(x) \leq \phi(0) + L|x|.$$

Plug it into the definition of convex conjugate (4), we can obtain

$$\phi^*(y) = \sup_{x \in \mathbb{R}} \{y \cdot x - \phi(x)\} \geq \sup_{x \in \mathbb{R}} \{y \cdot x - L|x|\} - \phi(0).$$

Since $y > L$, we can always let $(y \cdot x - L|x|)$ goes to infinity by set $x \rightarrow +\infty$, which means $\phi^*(y) = \infty$. And the proof for the case $y < -L$ is similar. \square

A.1 Proofs Concerning Algorithm 1

In this part, we will use

$$\mathcal{F}_t = \{i_0, j_0, \dots, i_{t-1}, j_{t-1}\}$$

to denote all the random variables generated before iteration t . And for simplicity, we will denote

$$\phi^*(y) = \sum_{i=1}^n \phi_i^*(y_i).$$

The following lemma is the key to prove both Theorem 1 and Theorem 2.

Lemma 5. *Consider t -th iteration of Algorithm 1. Assume $g(x)$ is μ -strongly convex ($\mu \geq 0$, and $\mu = 0$ means general convexity), and all ϕ_i are $(1/\gamma)$ -smooth ($\gamma \geq 0$). When conditioned on \mathcal{F}_t , it holds that*

$$\begin{aligned} & \frac{d}{2\eta_t} \cdot \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \mathbb{E} [\|x^{t+1} - x\|^2 | \mathcal{F}_t] \right] + \\ & \frac{d}{2\tau_t} \cdot \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \mathbb{E} [\|y^{t+1} - y\|^2 | \mathcal{F}_t] \right] \\ & \geq F(x^t, y) - F(x, y^t) - \eta_t L^2 R^2 - \frac{\eta_t D^2 R^2}{n} \\ & \quad + d \cdot [\mathbb{E} [g(x^{t+1}) | \mathcal{F}_t] - g(x^t)] + \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \phi^*(y^t) \end{aligned} \quad (13)$$

for any $x \in X$ and $y \in Y$.

Proof. Here we define two ‘‘imaginary’’ iterates x'^t and y'^t in the following way:

$$\begin{aligned} x'_j{}^t &= \text{prox}_{\eta_t g_j} (x_j^t - \eta_t a_{i_t j} y_{i_t}^t) \quad \forall j \in [d], \\ y'_i{}^t &= \text{prox}_{(\tau_t/d)\phi_i^*} (y_i^t - \tau_t a_{i_t i} x_{j_t}^t) \quad \forall i \in [n]. \end{aligned}$$

Note that $x_j^{t+1} = x'_j{}^t$ when $j = j_t$, and $x_j^{t+1} = x_j^t$ otherwise. And $y_i^{t+1} = y'_i{}^t$ when $i = i_t$, while $y_i^{t+1} = y_i^t$ when $i \neq i_t$.

For each $j \in [d]$, due to $(a+b)^2 - a^2 - b^2 = 2a \cdot b$, we have

$$(x_j^t - x_j)^2 - (x_j'^t - x_j)^2 - (x_j^t - x_j'^t)^2 = 2(x_j^t - x_j'^t) \cdot (x_j'^t - x_j)$$

for any $x_j \in X_j$. By the definition of $x_j'^t$ and the optimality condition of the proximal subproblem, there must exist a subgradient $s \in \partial g_j(x_j'^t) + \partial \mathbb{1}_{X_j}(x_j'^t)$ such that

$$x_j'^t = x_j^t - \eta_t a_{i_t j} y_{i_t}^t - \eta_t s,$$

where $\mathbb{1}_{X_j}$ is the indicator function of set X_j , i.e.,

$$\mathbb{1}_X(x) = \begin{cases} 0, & \text{if } x \in X, \\ +\infty, & \text{if } x \notin X. \end{cases}$$

Plug in this fact into the previous equality, we have

$$(x_j^t - x_j)^2 - (x_j^{t'} - x_j)^2 - (x_j^t - x_j^{t'})^2 = 2\eta_t(a_{i_t j}y_{i_t}^t + s) \cdot (x_j^{t'} - x_j).$$

Since $g(x)$ is μ -strongly convex, and also because of the separable assumption (2), we know that each $g_j(x_j)$ is also μ -strongly convex. Then we apply strong convexity to subgradient s :

$$\begin{aligned} & (x_j^t - x_j)^2 - (x_j^{t'} - x_j)^2 - (x_j^t - x_j^{t'})^2 \\ &= 2\eta_t(a_{i_t j}y_{i_t}^t + s) \cdot (x_j^{t'} - x_j) \\ &\geq 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^{t'} - x_j) + 2\eta_t \left[g_j(x_j^{t'}) - g_j(x_j) + \frac{\mu}{2}(x_j^{t'} - x_j)^2 + \mathbb{1}_{X_j}(x_j^{t'}) - \mathbb{1}_{X_j}(x_j) \right]. \end{aligned}$$

Since $x^{t'}$ and x are always feasible by definition, thus $\mathbb{1}_{X_j}(x_j^{t'}) = \mathbb{1}_{X_j}(x_j) = 0$. As a result,

$$\begin{aligned} & (x_j^t - x_j)^2 - (x_j^{t'} - x_j)^2 - (x_j^t - x_j^{t'})^2 \\ &\geq 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^{t'} - x_j) + 2\eta_t \left[g_j(x_j^{t'}) - g_j(x_j) + \frac{\mu}{2}(x_j^{t'} - x_j)^2 \right] \\ &= 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) + 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^{t'} - x_j^t) + 2\eta_t [g_j(x_j^{t'}) - g_j(x_j)] + \eta_t \mu (x_j^{t'} - x_j)^2 \\ &\geq 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) - 2\eta_t^2 (a_{i_t j} y_{i_t}^t)^2 - \frac{1}{2}(x_j^t - x_j^{t'})^2 + 2\eta_t [g_j(x_j^{t'}) - g_j(x_j)] + \eta_t \mu (x_j^{t'} - x_j)^2, \end{aligned}$$

where the last line is due to Cauchy-Schwarz inequality. By cancelling term $(1/2)(x_j^t - x_j^{t'})^2$ on both sides, and note that

$$\frac{1}{2}(x_j^t - x_j^{t'})^2 + \eta_t \mu (x_j^{t'} - x_j)^2 \geq \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} (x_j^t - x_j)^2,$$

we can further rewrite the above inequality as

$$\begin{aligned} & \left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) (x_j^t - x_j)^2 - (x_j^{t'} - x_j)^2 \\ &\geq 2\eta_t a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) - 2\eta_t^2 (a_{i_t j} y_{i_t}^t)^2 + 2\eta_t [g_j(x_j^{t'}) - g_j(x_j)]. \end{aligned}$$

We can bound the second term on the right-hand-side by $(a_{i_t j} y_{i_t}^t)^2 \leq L^2 a_{i_t j}^2$ because of Lemma 4, and we divide both sides by $2\eta_t$, then we obtain

$$\frac{1}{2\eta_t} \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) (x_j^t - x_j)^2 - (x_j^{t'} - x_j)^2 \right] \geq a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) + g_j(x_j^{t'}) - g_j(x_j) - \eta_t L^2 a_{i_t j}^2. \quad (14)$$

Even though $x_j^{t'}$ depends on i_t , it is independent of j_t . We observe that x_j^{t+1} takes value $x_j^{t'}$ with probability $1/d$, and x_j^t otherwise. Hence, by conditioning on $\mathcal{F}_t' \triangleq \mathcal{F}_t \cup \{i_t\}$, we always have

$$\begin{aligned} \mathbb{E} [g_j(x_j^{t+1}) | \mathcal{F}_t'] &= \frac{1}{d} g_j(x_j^{t'}) + \frac{d-1}{d} g_j(x_j^t), \\ \mathbb{E} [\|x_j^{t+1} - x_j\|^2 | \mathcal{F}_t'] &= \frac{1}{d} \|x_j^{t'} - x_j\|^2 + \frac{d-1}{d} \|x_j^t - x_j\|^2. \end{aligned}$$

Put these relationships into (14), when conditioned on \mathcal{F}_t' , we obtain:

$$\begin{aligned} & \frac{d}{2\eta_t} \cdot \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) (x_j^t - x_j)^2 - \mathbb{E} [(x_j^{t+1} - x_j)^2 | \mathcal{F}_t'] \right] \\ &\geq a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) + d \cdot \mathbb{E} [g_j(x_j^{t+1}) | \mathcal{F}_t'] - (d-1)g_j(x_j^t) - g_j(x_j) - \eta_t L^2 a_{i_t j}^2. \end{aligned}$$

We sum this inequality from $j = 1$ to $j = d$, and finally obtain:

$$\begin{aligned} & \frac{d}{2\eta_t} \cdot \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \mathbb{E} [\|x^{t+1} - x\|^2 | \mathcal{F}_t'] \right] \\ &\geq \sum_{j=1}^d a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) + d \cdot \mathbb{E} [g(x^{t+1}) | \mathcal{F}_t'] - (d-1)g(x^t) - g(x) - \eta_t L^2 \sum_{j=1}^d a_{i_t j}^2 \\ &\geq \sum_{j=1}^d a_{i_t j} y_{i_t}^t \cdot (x_j^t - x_j) + d \cdot \mathbb{E} [g(x^{t+1}) | \mathcal{F}_t'] - (d-1)g(x^t) - g(x) - \eta_t L^2 R^2, \end{aligned}$$

where the last inequality is due to the definition $R = \max_i \|a_i\|$. Now, we further take expectation with respect to i_t :

$$\begin{aligned} & \frac{d}{2\eta_t} \cdot \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \mathbb{E} [\|x^{t+1} - x\|^2 | \mathcal{F}_t] \right] \\ & \geq \frac{1}{n} y^{t\top} A(x^t - x) + d \cdot \mathbb{E} [g(x^{t+1}) | \mathcal{F}_t] - (d-1)g(x^t) - g(x) - \eta_t L^2 R^2. \end{aligned} \quad (15)$$

On the other hand, due to the symmetric nature of Algorithm 1, and noticing $(1/\gamma)$ -smoothness of ϕ_i implies γ -strong convexity of ϕ_i^* , we can derive similar inequality for y^{t+1} :

$$\begin{aligned} & \frac{n}{2\tau_t} \cdot \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \mathbb{E} [\|y^{t+1} - y\|^2 | \mathcal{F}_t] \right] \\ & \geq \frac{1}{d} (y - y^t) A x^t + \frac{n}{d} \cdot \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \frac{n-1}{d} \phi^*(y^t) - \frac{1}{d} \phi^*(y) - \frac{\tau_t D^2 R'^2}{d}, \end{aligned}$$

or equivalently,

$$\begin{aligned} & \frac{d}{2\tau_t} \cdot \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \mathbb{E} [\|y^{t+1} - y\|^2 | \mathcal{F}_t] \right] \\ & \geq \frac{1}{n} (y - y^t) A x^t + \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \frac{n-1}{n} \phi^*(y^t) - \frac{1}{n} \phi^*(y) - \frac{\tau_t D^2 R'^2}{n}, \end{aligned}$$

Let us add this inequality with (15):

$$\begin{aligned} & \frac{d}{2\eta_t} \cdot \left[\left(1 - \frac{\eta_t \mu}{2} \right) \|x^t - x\|^2 - \mathbb{E} [\|x^{t+1} - x\|^2 | \mathcal{F}_t] \right] + \\ & \frac{d}{2\tau_t} \cdot \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \mathbb{E} [\|y^{t+1} - y\|^2 | \mathcal{F}_t] \right] \\ & = \frac{1}{n} y^{t\top} A(x^t - x) + \frac{1}{n} (y - y^t) A x^t + d \cdot \mathbb{E} [g(x^{t+1}) | \mathcal{F}_t] - (d-1)g(x^t) - g(x) \\ & \quad + \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \frac{n-1}{n} \phi^*(y^t) - \frac{1}{n} \phi^*(y) - \eta_t L^2 R^2 - \frac{\eta_t D^2 R'^2}{n} \\ & = \left[\frac{1}{n} y^\top A x^t + g(x^t) - \frac{1}{n} \phi^*(y) \right] - \left[\frac{1}{n} y^{t\top} A x + g(x) - \frac{1}{n} \phi^*(y^t) \right] - \eta_t L^2 R^2 - \frac{\eta_t D^2 R'^2}{n} \\ & \quad + d \cdot [\mathbb{E} [g(x^{t+1}) | \mathcal{F}_t] - g(x^t)] + \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \phi^*(y^t) \\ & = F(x^t, y) - F(x, y^t) - \eta_t L^2 R^2 - \frac{\eta_t D^2 R'^2}{n} \\ & \quad + d \cdot [\mathbb{E} [g(x^{t+1}) | \mathcal{F}_t] - g(x^t)] + \mathbb{E} [\phi^*(y^{t+1}) | \mathcal{F}_t] - \phi^*(y^t), \end{aligned}$$

which is our desired result. \square

Lemma 6. *Under same assumption as Lemma 5, it holds that*

$$\begin{aligned} & \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 \right]}{\eta_t} + \\ & \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 \right]}{\tau_t} \\ & \geq T \cdot \mathbb{E} [F(\hat{x}^T, y) - F(x, \hat{y}^T)] - L^2 R^2 \sum_{t=0}^{T-1} \eta_t - \frac{D^2 R'^2}{n} \sum_{t=0}^{T-1} \tau_t. \end{aligned} \quad (16)$$

for Algorithm 1.

Proof. By summing up inequality (13) from $t = 0$ to $t = T - 1$, and applying the law of total expectation, we have

$$\begin{aligned}
& \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 \right]}{\eta_t} \\
& + \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 \right]}{\tau_t} \\
& \geq \sum_{t=0}^{T-1} \mathbb{E} [F(x^t, y) - F(x, y^t)] - L^2 R^2 \sum_{t=0}^{T-1} \eta_t - \frac{D^2 R'^2}{n} \sum_{t=0}^{T-1} \tau_t \\
& + d \sum_{t=0}^{T-1} \mathbb{E} [g(x^{t+1}) - g(x^t)] + \sum_{t=0}^{T-1} \mathbb{E} [\phi^*(y^{t+1}) - \phi^*(y^t)] \\
& = \sum_{t=0}^{T-1} \mathbb{E} [F(x^t, y) - F(x, y^t)] - L^2 R^2 \sum_{t=0}^{T-1} \eta_t - \frac{D^2 R'^2}{n} \sum_{t=0}^{T-1} \tau_t \\
& + d \cdot \mathbb{E} [g(x^T) - g(x^0)] + \mathbb{E} [\phi^*(y^T) - \phi^*(y^0)].
\end{aligned}$$

Recall that x^0 and y^0 satisfy $x^0 = \arg \min_{x \in X} g(x)$ and $y^0 = \arg \min_{y \in Y} \phi^*(y)$, which means $g(x^T) - g(x^0) \geq 0$ and $\phi^*(y^T) - \phi^*(y^0) \geq 0$. Hence,

$$\begin{aligned}
& \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\eta_t \mu}{2}, \frac{1}{4} \right\} \right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 \right]}{\eta_t} \\
& + \frac{d}{2} \cdot \sum_{t=0}^{T-1} \frac{\mathbb{E} \left[\left(1 - \min \left\{ \frac{\tau_t \gamma}{2d}, \frac{1}{4} \right\} \right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 \right]}{\tau_t} \\
& \geq \sum_{t=0}^{T-1} \mathbb{E} [F(x^t, y) - F(x, y^t)] - dL^2 R^2 \sum_{t=0}^{T-1} \eta_t - \frac{D^2 R'^2}{n} \sum_{t=0}^{T-1} \tau_t. \tag{17}
\end{aligned}$$

Besides, we use the definition of \bar{x}^T and \bar{y}^T , along with the convexity-concavity of $F(x, y)$, we can obtain

$$T \cdot \mathbb{E} [F(\hat{x}^T, y) - F(x, \hat{y}^T)] \leq \sum_{t=0}^{T-1} \mathbb{E} [F(x^t, y) - F(x, y^t)].$$

Now we can complete the proof by combining this fact with (17). \square

Now, we are ready to prove Theorem 1:

Proof of Theorem 1. Lemma 6 can be applied here with $\mu = 0$ and $\gamma = 0$. Let us first bound the term on the left-hand-side of (16):

$$\begin{aligned}
& \sum_{t=0}^{T-1} \frac{\mathbb{E} [\|x^t - x\|^2 - \|x^{t+1} - x\|^2]}{\eta_t} \\
& = \frac{1}{\eta_0} \cdot \|x^0 - x\|^2 + \sum_{t=1}^{T-1} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\|x^t - x\|^2] - \frac{1}{\eta_{T-1}} \cdot \mathbb{E} [\|x^T - x\|^2] \\
& \leq \frac{1}{\eta_0} \cdot 4D^2 + \sum_{t=1}^{T-1} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) 4D^2 \\
& = \frac{4D^2}{\eta_{T-1}}
\end{aligned}$$

by the boundness assumption of X , along with the fact sequence $\{\eta_t\}$ is non-increasing. By applying Lemma 4 again, we know that $\|y^t - y\|^2 \leq 2nL^2$, and hence we can similarly show that

$$\sum_{t=0}^{T-1} \frac{\mathbb{E} [\|y^t - y\|^2 - \|y^{t+1} - y\|^2]}{\tau_t} \leq \frac{4nL^2}{\tau_{T-1}}.$$

Combine all these facts into (16), one can obtain

$$T \cdot \mathbb{E} [F(\hat{x}^T, y) - F(x, \hat{y}^T)] \leq \frac{2dD^2}{\eta_{T-1}} + \frac{2ndL^2}{\tau_{T-1}} + L^2R^2 \sum_{t=0}^{T-1} \eta_t + \frac{D^2R'^2}{n} \sum_{t=0}^{T-1} \tau_t.$$

Divide both sides by T , and then take supremum with respect to $x \in X$ and $y \in Y$:

$$\begin{aligned} & \mathbb{E} [\mathcal{G}(\hat{x}^T, \hat{y}^T)] \\ &= \sup_{x \in X, y \in Y} \mathbb{E} [F(\hat{x}^T, y) - F(x, \hat{y}^T)] \\ &\leq \frac{2dD^2/\eta_{T-1} + 2ndL^2/\tau_{T-1} + L^2R^2 \sum_{t=0}^{T-1} \eta_t + (D^2R'^2/n) \sum_{t=0}^{T-1} \tau_t}{T}, \end{aligned}$$

and the theorem can be proven by plug the values of $\{\eta_t\}$ and $\{\tau_t\}$ into this inequality. \square

We are ready to prove Theorem 2:

Proof of Theorem 2. Again, we need to apply Lemma 6 here. The first term on the left-hand-side of (16) can be bounded in the following way:

$$\begin{aligned} & \sum_{t=0}^{T-1} \frac{\mathbb{E} [(1 - \min\{\frac{\eta_t\mu}{2}, \frac{1}{4}\}) \|x^t - x\|^2 - \|x^{t+1} - x\|^2]}{\eta_t} \\ &= \frac{1 - \min\{\frac{\eta_0\mu}{2}, \frac{1}{4}\}}{\eta_0} \cdot \|x^0 - x\|^2 + \sum_{t=1}^{T-1} \left(\frac{1 - \min\{\frac{\eta_t\mu}{2}, \frac{1}{4}\}}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \mathbb{E} [\|x^t - x\|^2] - \frac{1}{\eta_{T-1}} \cdot \mathbb{E} [\|x^T - x\|^2] \\ &\leq \frac{1}{\eta_0} \cdot \|x^0 - x\|^2 + \sum_{t=1}^{T-1} \max\left\{ \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} - \frac{\mu}{2}, \frac{3}{4\eta_t} - \frac{1}{\eta_{t-1}} \right\} \mathbb{E} [\|x^t - x\|^2] \\ &\leq \frac{1}{\eta_0} \cdot \|x^0 - x\|^2 + \sum_{t=1}^{T-1} 0 \cdot \mathbb{E} [\|x^t - x\|^2] \\ &\leq \frac{4}{\eta_0} \cdot D^2, \end{aligned}$$

where the third equality follows from our choices of $\{\eta_t\}$, and the last inequality is due to boundness assumption again. Similarly, we can derive the bound

$$\sum_{t=0}^{T-1} \frac{\mathbb{E} [(1 - \min\{\frac{\tau_t\gamma}{2d}, \frac{1}{4}\}) \|y^t - y\|^2 - \|y^{t+1} - y\|^2]}{\tau_t} \leq \frac{4nL^2}{\tau_0}.$$

Combine these two bounds into (16) and then take supremum with respect to x and y , we can finally obtain:

$$\mathbb{E} [\mathcal{G}(\hat{x}^T, \hat{y}^T)] \leq \frac{2dD^2/\eta_0 + 2ndL^2/\tau_0 + L^2R^2 \sum_{t=0}^{T-1} \eta_t + (D^2R'^2/n) \sum_{t=0}^{T-1} \tau_t}{T}.$$

We can finish the proof by plug the values of $\{\eta_t\}$ and $\{\tau_t\}$. \square

A.2 Proofs Concerning Algorithm 2

Again, we use

$$\mathcal{F}_t = \{i_0, i'_0, j_0, j'_0, \dots, i_{t-1}, i'_{t-1}, j_{t-1}, j'_{t-1}\}$$

to denote all the random variables generated before iteration t . We use the following notation to denote the variance-reduced gradients used in our algorithms:

$$\begin{aligned}\nabla_{x,j}^t &= a_{i'_t j} (y_{i'_t}^t - \tilde{y}_{i'_t}^k) + G_{x,j}^k, \\ \nabla_{y,i}^t &= a_{i j'_t} (x_{j'_t}^t - \tilde{x}_{j'_t}^k) + G_{y,i}^k, \\ \nabla_{x,j}^{\prime t} &= a_{i_t j} (\bar{y}_{i_t}^t - \tilde{y}_{i_t}^k) + G_{x,j}^k, \\ \nabla_{y,i}^{\prime t} &= a_{i j_t} (\bar{x}_{j_t}^t - \tilde{x}_{j_t}^k) + G_{y,i}^k.\end{aligned}$$

Besides, we will also define two ‘‘imaginary’’ iterates $x^{\prime t}$ and $y^{\prime t}$:

$$\begin{aligned}x_j^{\prime t} &= \text{prox}_{\eta g_j} (x_j^t - \eta \nabla_{x,j}^t) \quad \forall j \in [d], \\ y_i^{\prime t} &= \text{prox}_{(\tau/d)\phi_i^*} (y_i^t - \tau \nabla_{y,i}^t) \quad \forall i \in [n].\end{aligned}$$

Obviously, $\bar{x}_j^t = x_j^{\prime t}$ when $j = j_t$ and $\bar{y}_i^t = y_i^{\prime t}$ when $i = i_t$. A key observation here is that each $x_j^{\prime t}$ only depends on i'_t when conditioned on \mathcal{F}_t , and it is independent of i_t, j_t and j'_t . Similarly, $y_i^{\prime t}$ is also independent of i_t, j_t and i'_t .

First, let us develop a bound for the gradient variance:

Lemma 7. *Algorithm 2 has the following bounds for its gradient variance:*

$$\mathbb{E} [(\nabla_{x,j_t}^t - \nabla_{x,j_t}^{\prime t})^2 | \mathcal{F}_t] \leq \frac{8R^2}{nd} \|y - \tilde{y}^k\|^2 + \frac{12R^2}{nd} \|y^t - y\|^2 + \frac{8R^2}{d} \mathbb{E} [\|y^t - \bar{y}^t\|^2 | \mathcal{F}_t] \quad (18)$$

$$\mathbb{E} [(\nabla_{y,i_t}^t - \nabla_{y,i_t}^{\prime t})^2 | \mathcal{F}_t] \leq \frac{8R'^2}{nd} \|x - \tilde{x}^k\|^2 + \frac{12R'^2}{nd} \|x^t - x\|^2 + \frac{8R'^2}{n} \mathbb{E} [\|x^t - \bar{x}^t\|^2 | \mathcal{F}_t] \quad (19)$$

for any $x \in X$ and $y \in Y$.

Proof. Just by definition of ∇_{x,j_t}^t and $\nabla_{x,j_t}^{\prime t}$, and repeatedly using $(a + b)^2 \leq 2a^2 + 2b^2$, we have

$$\begin{aligned}(\nabla_{x,j_t}^t - \nabla_{x,j_t}^{\prime t})^2 &= \left[a_{i'_t j_t} (y_{i'_t}^t - \tilde{y}_{i'_t}^k) - a_{i_t j_t} (\bar{y}_{i_t}^t - \tilde{y}_{i_t}^k) \right]^2 \\ &\leq 2a_{i'_t j_t}^2 (y_{i'_t}^t - \tilde{y}_{i'_t}^k)^2 + 2a_{i_t j_t}^2 (\bar{y}_{i_t}^t - \tilde{y}_{i_t}^k)^2 \\ &\leq 2a_{i'_t j_t}^2 \left[2(y_{i'_t}^t - y_{i'_t}^t)^2 + 2(y_{i'_t}^t - \tilde{y}_{i'_t}^k)^2 \right] \\ &\quad + 2a_{i_t j_t}^2 \left[2(\tilde{y}_{i_t}^k - y_{i_t}^t)^2 + 2(y_{i_t}^t - \bar{y}_{i_t}^t)^2 \right] \\ &\leq 2a_{i'_t j_t}^2 \left[2(y_{i'_t}^t - y_{i'_t}^t)^2 + 2(y_{i'_t}^t - \tilde{y}_{i'_t}^k)^2 \right] \\ &\quad + 2a_{i_t j_t}^2 \left[2(\tilde{y}_{i_t}^k - y_{i_t}^t)^2 + 4(y_{i_t}^t - y_{i_t}^t)^2 + 4(y_{i_t}^t - \bar{y}_{i_t}^t)^2 \right] \\ &= 2a_{i'_t j_t}^2 \left[2(y_{i'_t}^t - y_{i'_t}^t)^2 + 2(y_{i'_t}^t - \tilde{y}_{i'_t}^k)^2 \right] \\ &\quad + 2a_{i_t j_t}^2 \left[2(\tilde{y}_{i_t}^k - y_{i_t}^t)^2 + 4(y_{i_t}^t - y_{i_t}^t)^2 + 4\|y^t - \bar{y}^t\|^2 \right], \quad (20)\end{aligned}$$

where the last line is due to the fact that \bar{x}^t and x^t only differ in coordinate i_t . In the next, we will take expectation. Since i'_t and j_t are independent, then

$$\begin{aligned}\mathbb{E} \left[a_{i'_t j_t}^2 (y_{i'_t}^t - y_{i'_t}^t)^2 | \mathcal{F}_t \right] &= \mathbb{E} \left[\mathbb{E}_{j_t} \left[a_{i'_t j_t}^2 \right] \cdot (y_{i'_t}^t - y_{i'_t}^t)^2 | \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\frac{\|a_{i_t}\|^2}{d} \cdot (y_{i'_t}^t - y_{i'_t}^t)^2 | \mathcal{F}_t \right] \\ &\leq \frac{R^2}{d} \cdot \mathbb{E} \left[(y_{i'_t}^t - y_{i'_t}^t)^2 | \mathcal{F}_t \right] \\ &= \frac{R^2}{nd} \|y - y^t\|^2.\end{aligned}$$

And we can bound other terms in (20) similarly, then we obtain (18). The proof for variance bound (19) is analogous. \square

The following lemma is the key to prove the convergence of SPD1-VR:

Lemma 8. *Assume $g(x)$ is μ -strongly convex, and all ϕ_i is $(1/\gamma)$ -smooth, then by conditioning on \mathcal{F}_t , it holds that*

$$\begin{aligned} & \frac{1}{2\eta} \mathbb{E} \left[\left(1 - \frac{\eta\mu}{4d} + \frac{12R'^2\eta\tau}{nd} \right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{n} \right) \|x^t - \bar{x}^t\|^2 \middle| \mathcal{F}_t \right] \\ & + \frac{1}{2\tau} \mathbb{E} \left[\left(1 - \frac{\tau\gamma}{4nd} + \frac{12R'^2\eta\tau}{nd} \right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{d} \right) \|y^t - \bar{y}^t\|^2 \middle| \mathcal{F}_t \right] \\ & \geq \frac{1}{d} \cdot \mathbb{E} [F(x^t, y) - F(x, y^t) | \mathcal{F}_t] - \frac{4\eta R'^2}{nd} \|\bar{y}^t - y\|^2 - \frac{4\tau R'^2}{nd} \|\bar{x}^t - x\|^2, \end{aligned}$$

for any $x \in X$ and $y \in Y$.

Proof. First, by $2a \cdot b = (a + b)^2 - a^2 - b^2$, we have the following two inequalities:

$$2(x^t - \bar{x}^t)^\top (\bar{x}^t - x^{t+1}) = \|x^t - x^{t+1}\|^2 - \|x^{t+1} - \bar{x}^t\|^2 - \|x^t - \bar{x}^t\|^2$$

and

$$2(x^t - x^{t+1})^\top (x^{t+1} - x) = \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \|x^t - x^{t+1}\|^2$$

for any $x \in X$. By adding these two inequality together, we have

$$\begin{aligned} & \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \|x^{t+1} - \bar{x}^t\|^2 - \|x^t - \bar{x}^t\|^2 \\ & = 2(x^t - \bar{x}^t)^\top (\bar{x}^t - x^{t+1}) + 2(x^t - x^{t+1})^\top (x^{t+1} - x) \\ & = 2(x_{j_t}^t - \bar{x}_{j_t}^t) \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2(x_{j_t}^t - x_{j_t}^{t+1}) \cdot (x_{j_t}^{t+1} - x_{j_t}), \end{aligned} \quad (21)$$

where the last line is because x^t and \bar{x}^t only differs in coordinate j_t , and similarly for x^t and x^{t+1} . According to the updating rules

$$\begin{aligned} \bar{x}_{j_t}^t &= \text{prox}_{\eta g_{j_t}} (x_{j_t}^t - \eta \nabla_{x, j_t}^t), \\ x_{j_t}^{t+1} &= \text{prox}_{\eta g_{j_t}} (x_{j_t}^t - \eta \nabla_{x, j_t}^t), \end{aligned}$$

and the optimality condition of the proximal mapping subproblem, there must exist $s \in \partial g_{j_t}(\bar{x}_{j_t}^t) + \partial \mathbb{1}_{X_{j_t}}(\bar{x}_{j_t}^t)$ and $s' \in \partial g_{j_t}(x_{j_t}^{t+1}) + \partial \mathbb{1}_{X_{j_t}}(x_{j_t}^{t+1})$ such that

$$\begin{aligned} \bar{x}_{j_t}^t &= x_{j_t}^t - \eta \nabla_{x, j_t}^t - \eta s, \\ x_{j_t}^{t+1} &= x_{j_t}^t - \eta \nabla_{x, j_t}^t - \eta s', \end{aligned}$$

where $\mathbb{1}_Z(z)$ is the indicator function of convex set Z , which takes value 0 when $z \in Z$, while $\mathbb{1}_Z(z) = +\infty$ when $z \notin Z$. Combine these facts into (21), one can obtain:

$$\begin{aligned} & \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \|x^{t+1} - \bar{x}^t\|^2 - \|x^t - \bar{x}^t\|^2 \\ & = 2\eta s \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta s' \cdot (x_{j_t}^{t+1} - x_{j_t}) + 2\eta \nabla_{x, j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta \nabla_{x, j_t}^t \cdot (x_{j_t}^{t+1} - x_{j_t}). \end{aligned}$$

Because of the separable assumption (2), μ -strongly convex of $g(x)$ implies μ -strongly convex of every component function $g_j(x_j)$. Now we apply the convexity of indicator function $\mathbb{1}_{X_{j_t}}(x_{j_t})$ and the strong convexity of $g_{j_t}(x_{j_t})$, and further observe that all of \bar{x}^t , x^{t+1} and x are always feasible, which mean $\mathbb{1}_{X_{j_t}}(\bar{x}_{j_t}^t) = \mathbb{1}_{X_{j_t}}(x_{j_t}^{t+1}) = \mathbb{1}_{X_{j_t}}(x_{j_t}) = 0$, and have

$$\begin{aligned} & \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \|x^{t+1} - \bar{x}^t\|^2 - \|x^t - \bar{x}^t\|^2 \\ & \geq 2\eta \left[g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t}^{t+1}) + \frac{\mu}{2} (\bar{x}_{j_t}^t - x_{j_t}^{t+1})^2 \right] + 2\eta \left[g_{j_t}(x_{j_t}^{t+1}) - g_{j_t}(x_{j_t}) + \frac{\mu}{2} (x_{j_t}^{t+1} - x_{j_t})^2 \right] \\ & \quad + 2\eta \nabla_{x, j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta \nabla_{x, j_t}^t \cdot (x_{j_t}^{t+1} - x_{j_t}) \\ & = 2\eta \left[g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t}) + \frac{\mu}{2} (\bar{x}_{j_t}^t - x_{j_t}^{t+1})^2 + \frac{\mu}{2} (x_{j_t}^{t+1} - x_{j_t})^2 \right] \\ & \quad + 2\eta \nabla_{x, j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta \nabla_{x, j_t}^t \cdot (x_{j_t}^{t+1} - x_{j_t}) \\ & \geq 2\eta \left[g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t}) + \frac{\mu}{4} (\bar{x}_{j_t}^t - x_{j_t})^2 \right] + 2\eta \nabla_{x, j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta \nabla_{x, j_t}^t \cdot (x_{j_t}^{t+1} - x_{j_t}), \end{aligned}$$

where inequality $a^2 + b^2 \geq (1/2)(a + b)^2$ is used in the second inequality. When $\eta\mu \leq 1$, it holds that

$$\frac{1}{2}\|x^t - \bar{x}^t\|^2 + \frac{\eta\mu}{2}(\bar{x}_{j_t}^t - x_{j_t})^2 \geq \frac{1}{2}(x_{j_t}^t - \bar{x}_{j_t}^t)^2 + \frac{\eta\mu}{2}(\bar{x}_{j_t}^t - x_{j_t})^2 \geq \frac{\eta\mu}{4}(x_{j_t}^t - x_{j_t})^2.$$

As a result, we have

$$\begin{aligned} & \|x^t - x\|^2 - \frac{\eta\mu}{4}(x_{j_t}^t - x_{j_t})^2 - \|x^{t+1} - x\|^2 - \|x^{t+1} - \bar{x}^t\|^2 - \frac{1}{2}\|x^t - \bar{x}^t\|^2 \\ & \geq 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta \nabla_{x,j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) + 2\eta \nabla_{x,j_t}^{t'} \cdot (x_{j_t}^{t+1} - x_{j_t}) \\ & = 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta \nabla_{x,j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}) + 2\eta (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'}) \cdot (\bar{x}_{j_t}^t - x_{j_t}^{t+1}) \\ & \geq 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta \nabla_{x,j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}) - \eta^2 (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2 - (\bar{x}_{j_t}^t - x_{j_t}^{t+1})^2 \\ & \geq 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta \nabla_{x,j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}) - \eta^2 (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2 - \|\bar{x}^t - x^{t+1}\|^2, \end{aligned}$$

where Cauchy-Schwarz inequality is used in the second inequality. After cancelling term $\|x^{t+1} - \bar{x}^t\|^2$ on both sides, we finally obtain

$$\begin{aligned} & \|x^t - x\|^2 - \frac{\eta\mu}{4}(x_{j_t}^t - x_{j_t})^2 - \|x^{t+1} - x\|^2 - \frac{1}{2}\|x^t - \bar{x}^t\|^2 \\ & \geq 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta \nabla_{x,j_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}) - \eta^2 (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2 \\ & = 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta a_{i_t,j_t} \bar{y}_{i_t}^t \cdot (\bar{x}_{j_t}^t - x_{j_t}) - 2\eta (a_{i_t,j_t} \bar{y}_{i_t}^t - \nabla_{y,i_t}^{t'}) \cdot (\bar{x}_{j_t}^t - x_{j_t}) - \eta^2 (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2 \\ & = 2\eta [g_{j_t}(\bar{x}_{j_t}^t) - g_{j_t}(x_{j_t})] + 2\eta a_{i_t,j_t} y_{i_t}^{t'} \cdot (x_{j_t}^t - x_{j_t}) - 2\eta (a_{i_t,j_t} y_{i_t}^{t'} - \nabla_{y,i_t}^{t'}) \cdot (x_{j_t}^t - x_{j_t}) - \eta^2 (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2, \end{aligned}$$

where in the last line we replaced $\bar{x}_{j_t}^t$ by $x_{j_t}^t$, and $\bar{y}_{i_t}^t$ by $y_{i_t}^{t'}$.

Similarly, we can derive an analogous bound for dual variable y :

$$\begin{aligned} & \|y^t - x\|^2 - \frac{\tau\gamma}{4d}(y_{i_t}^t - y_{i_t})^2 - \|y^{t+1} - y\|^2 - \frac{1}{2}\|y^t - \bar{y}^t\|^2 \\ & = \frac{2\tau}{d} [\phi_{i_t}^*(y_{i_t}^{t'}) - \phi_{i_t}^*(y_{i_t})] + 2\tau a_{i_t,j_t} x_{j_t}^{t'} \cdot (y_{i_t} - y_{i_t}^{t'}) - 2\tau (a_{i_t,j_t} x_{j_t}^{t'} - \nabla_{y,i_t}^{t'}) \cdot (y_{i_t} - y_{i_t}^{t'}) - \tau^2 (\nabla_{y,i_t}^{t'} - \nabla_{y,i_t}^t)^2. \end{aligned}$$

After diving them by 2η and 2τ respectively, we add the above two inequalities together:

$$\begin{aligned} & \frac{1}{2\eta} \left[\|x^t - x\|^2 - \frac{\eta\mu}{4}(x_{j_t}^t - x_{j_t})^2 - \|x^{t+1} - x\|^2 - \frac{1}{2}\|x^t - \bar{x}^t\|^2 \right] \\ & + \frac{1}{2\tau} \left[\|y^t - x\|^2 - \frac{\tau\gamma}{4d}(y_{i_t}^t - y_{i_t})^2 - \|y^{t+1} - y\|^2 - \frac{1}{2}\|y^t - \bar{y}^t\|^2 \right] \\ & = g_{j_t}(x_{j_t}^{t'}) - g_{j_t}(x_{j_t}) - a_{i_t,j_t} y_{i_t}^{t'} \cdot x_{j_t} + a_{i_t,j_t} x_{j_t}^{t'} \cdot y_{i_t} + \frac{1}{d} [\phi_{i_t}^*(y_{i_t}^{t'}) - \phi_{i_t}^*(y_{i_t})] \\ & - (a_{i_t,j_t} y_{i_t}^{t'} - \nabla_{x,j_t}^{t'}) \cdot (x_{j_t}^t - x_{j_t}) - (a_{i_t,j_t} x_{j_t}^{t'} - \nabla_{y,i_t}^{t'}) \cdot (y_{i_t} - y_{i_t}^{t'}) \\ & - \frac{\eta}{2} (\nabla_{x,j_t}^t - \nabla_{x,j_t}^{t'})^2 - \frac{\tau}{2} (\nabla_{y,i_t}^{t'} - \nabla_{y,i_t}^t)^2. \end{aligned} \tag{22}$$

Now, we need to take conditional expectation of this inequality, by conditioning on \mathcal{F}_t . First, since both $x_{j_t}^{t'}$ and $y_{i_t}^{t'}$ are independent of both i_t and j_t for any $i \in [n]$ and $j \in [d]$. Thus,

$$\begin{aligned} \mathbb{E} [g_{j_t}(x_{j_t}^{t'}) - g_{j_t}(x_{j_t}) | \mathcal{F}_t] &= \frac{1}{d} \sum_{j=1}^d \mathbb{E} [g_j(x_j^{t'}) - g_j(x_j) | \mathcal{F}_t] = \frac{1}{d} \mathbb{E} [g(x^{t'}) - g(x) | \mathcal{F}_t], \\ \mathbb{E} [\phi_{i_t}^*(y_{i_t}^{t'}) - \phi_{i_t}^*(y_{i_t}) | \mathcal{F}_t] &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\phi_i^*(y_i^{t'}) - \phi_i^*(y_i) | \mathcal{F}_t] = \frac{1}{n} \mathbb{E} [\phi^*(y^{t'}) - \phi^*(y) | \mathcal{F}_t], \\ \mathbb{E} [a_{i_t,j_t} y_{i_t}^{t'} \cdot x_{j_t} | \mathcal{F}_t] &= \frac{1}{nd} \sum_{i=1}^n \sum_{j=1}^d \mathbb{E} [a_{ij} x_j y_i^{t'}] = \frac{1}{nd} \mathbb{E} [y^{t'\top} Ax], \\ \mathbb{E} [a_{i_t,j_t} x_{j_t}^{t'} \cdot y_{i_t} | \mathcal{F}_t] &= \frac{1}{nd} \sum_{i=1}^n \sum_{j=1}^d \mathbb{E} [a_{ij} x_j^{t'} y_i] = \frac{1}{nd} \mathbb{E} [y^\top Ax^{t'}]. \end{aligned}$$

Second, by using the definition of $\nabla_{x,j_t}^{t_t}$ and $\nabla_{y,i_t}^{t_t}$, we have

$$\begin{aligned} & \mathbb{E} \left[(a_{i_t,j_t} y_{i_t}^{t_t} - \nabla_{x,j_t}^{t_t}) \cdot (x_{j_t}^{t_t} - x_{j_t}) | \mathcal{F}_t \right] \\ &= \mathbb{E} \left[(a_{i_t,j_t} \tilde{y}_{i_t}^k - G_{x,j_t}^k) \cdot (x_{j_t}^{t_t} - x_{j_t}) | \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\mathbb{E}_{i_t} [a_{i_t,j_t} \tilde{y}_{i_t}^k - G_{x,j_t}^k] \cdot (x_{j_t}^{t_t} - x_{j_t}) | \mathcal{F}_t \right] \\ &= \mathbb{E} [0 \cdot (x_{j_t}^{t_t} - x_{j_t}) | \mathcal{F}_t] = 0, \end{aligned}$$

and similarly

$$\mathbb{E} \left[(a_{i_t,j_t} x_{j_t}^{t_t} - \nabla_{y,i_t}^{t_t}) \cdot (y_{i_t} - y_{i_t}^{t_t}) | \mathcal{F}_t \right] = 0.$$

Finally, observe that

$$\begin{aligned} \mathbb{E} \left[(x_{j_t}^t - x_{j_t})^2 | \mathcal{F}_t \right] &= \frac{1}{d} \|x^t - x\|^2, \\ \mathbb{E} \left[(y_{i_t}^t - y_{i_t})^2 | \mathcal{F}_t \right] &= \frac{1}{n} \|y^t - y\|^2, \end{aligned}$$

because both x^t and x is deterministic when conditioned \mathcal{F}_t . By putting all these facts back into (22), we can get:

$$\begin{aligned} & \frac{1}{2\eta} \mathbb{E} \left[\left(1 - \frac{\eta\mu}{4d}\right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \frac{1}{2} \|x^t - \bar{x}^t\|^2 | \mathcal{F}_t \right] \\ &+ \frac{1}{2\tau} \mathbb{E} \left[\left(1 - \frac{\tau\gamma}{4nd}\right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 - \frac{1}{2} \|y^t - \bar{y}^t\|^2 | \mathcal{F}_t \right] \\ &\geq \frac{1}{d} \cdot \mathbb{E} \left[g(x^{t_t}) - g(x) - \frac{1}{n} y^{t_t \top} Ax + \frac{1}{n} y^\top Ax^{t_t} + \frac{1}{n} [\phi^*(y^{t_t}) - \phi^*(y)] | \mathcal{F}_t \right] \\ &\quad - \frac{\eta}{2} \mathbb{E} \left[(\nabla_{x,j_t}^{t_t} - \nabla_{x,j_t}^t)^2 | \mathcal{F}_t \right] - \frac{\tau}{2} \mathbb{E} \left[(\nabla_{y,i_t}^{t_t} - \nabla_{y,i_t}^t)^2 | \mathcal{F}_t \right] \\ &= \frac{1}{d} \cdot \mathbb{E} [F(x^{t_t}, y) - F(x, y^{t_t}) | \mathcal{F}_t] - \frac{\eta}{2} \mathbb{E} \left[(\nabla_{x,j_t}^{t_t} - \nabla_{x,j_t}^t)^2 | \mathcal{F}_t \right] - \frac{\tau}{2} \mathbb{E} \left[(\nabla_{y,i_t}^{t_t} - \nabla_{y,i_t}^t)^2 | \mathcal{F}_t \right]. \end{aligned}$$

Now, we apply Lemma 7 to bound the variance terms, and rearrange terms, then get

$$\begin{aligned} & \frac{1}{2\eta} \mathbb{E} \left[\left(1 - \frac{\eta\mu}{4d} + \frac{12R'^2\eta\tau}{nd}\right) \|x^t - x\|^2 - \|x^{t+1} - x\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{n}\right) \|x^t - \bar{x}^t\|^2 | \mathcal{F}_t \right] \\ &+ \frac{1}{2\tau} \mathbb{E} \left[\left(1 - \frac{\tau\gamma}{4nd} + \frac{12R'^2\eta\tau}{nd}\right) \|y^t - y\|^2 - \|y^{t+1} - y\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{d}\right) \|y^t - \bar{y}^t\|^2 | \mathcal{F}_t \right] \\ &\geq \frac{1}{d} \cdot \mathbb{E} [F(x^{t_t}, y) - F(x, y^{t_t}) | \mathcal{F}_t] - \frac{4\eta R'^2}{nd} \|\tilde{y}^k - y\|^2 - \frac{4\tau R'^2}{nd} \|\tilde{x}^k - x\|^2, \end{aligned}$$

which is the desired result. \square

Now, we are ready to prove Theorem 3:

Proof of Theorem 3. First, we apply Lemma 8, and set $(x, y) = (x^*, y^*)$, we can get

$$\begin{aligned} & \frac{1}{2\eta} \mathbb{E} \left[\left(1 - \frac{\eta\mu}{4d} + \frac{12R'^2\eta\tau}{nd}\right) \|x^t - x^*\|^2 - \|x^{t+1} - x^*\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{n}\right) \|x^t - \bar{x}^t\|^2 | \mathcal{F}_t \right] \\ &+ \frac{1}{2\tau} \mathbb{E} \left[\left(1 - \frac{\tau\gamma}{4nd} + \frac{12R'^2\eta\tau}{nd}\right) \|y^t - y^*\|^2 - \|y^{t+1} - y^*\|^2 - \left(\frac{1}{2} - \frac{8\eta\tau R'^2}{d}\right) \|y^t - \bar{y}^t\|^2 | \mathcal{F}_t \right] \\ &\geq \frac{1}{d} \cdot \mathbb{E} [F(x^{t_t}, y^*) - F(x^*, y^{t_t}) | \mathcal{F}_t] - \frac{4\eta R'^2}{nd} \|\tilde{y}^k - y^*\|^2 - \frac{4\tau R'^2}{nd} \|\tilde{x}^k - x^*\|^2 \\ &\geq -\frac{4\eta R'^2}{nd} \|\tilde{y}^k - y^*\|^2 - \frac{4\tau R'^2}{nd} \|\tilde{x}^k - x^*\|^2 \end{aligned} \tag{23}$$

where $F(x, y^*) - F(x^*, y) \geq 0$ for any $x \in X$ and $y \in Y$ is a property of the optimal solution (x^*, y^*) of saddle point problems.

In the next, we will bound the coefficients in (23) respectively. Because of our choice of step sizes, we have

$$\eta \cdot \tau = \frac{\gamma}{128R^2} \min \left\{ \frac{d\kappa}{n\kappa'}, 1 \right\} \cdot \frac{n\mu}{128R'^2} \min \left\{ \frac{n\kappa'}{d\kappa}, 1 \right\} = \frac{n\gamma\mu}{2^{14}R^2R'^2} \min \left\{ \frac{d\kappa}{n\kappa'}, \frac{n\kappa'}{d\kappa} \right\}. \quad (24)$$

Since $\min\{\alpha, \alpha^{-1}\} \leq 1$ for any $\alpha > 0$, this equality implies

$$\frac{1}{2} - \frac{8\eta\tau R^2}{d} \geq \frac{1}{2} - \frac{n\gamma\mu}{2^{11}dR'^2} = \frac{1}{2} - \frac{1}{2^{11}\kappa'} \geq 0 \quad \text{and} \quad \frac{1}{2} - \frac{8\eta\tau R'^2}{n} \geq \frac{1}{2} - \frac{1}{2^{11}\kappa} \geq 0,$$

where the definition of κ and κ' along with the assumptions $\kappa \geq 1$ and $\kappa' \geq 1$ are used. Besides, also from (24) we can know:

$$\begin{aligned} \frac{4\eta R^2}{nd} &= \frac{\gamma\mu}{2^{12}dR'^2 \cdot \tau} \cdot \min \left\{ \frac{d\kappa}{n\kappa'}, \frac{n\kappa'}{d\kappa} \right\} = \frac{1}{2^{12}n\kappa' \cdot \tau} \cdot \min \left\{ \frac{d\kappa}{n\kappa'}, \frac{n\kappa'}{d\kappa} \right\} \leq \frac{1}{2^{12} \max\{n\kappa', d\kappa\} \cdot \tau} \\ &= \frac{1}{2^{12}m \cdot \tau}, \end{aligned}$$

where we denote $m \triangleq \max\{n\kappa', d\kappa\}$ for simplicity, and similarly we can show that

$$\frac{4\tau R'^2}{nd} \leq \frac{1}{2^{12}m \cdot \eta}.$$

Furthermore, observe that

$$\begin{aligned} \frac{\eta\mu}{4d} - \frac{12R'^2\eta\tau}{nd} &= \frac{\mu\eta}{4d} \left(1 - \frac{48R'^2\tau}{n\mu} \right) \\ &= \frac{\mu\eta}{4d} \left(1 - \frac{48}{128} \cdot \min \left\{ \frac{n\kappa'}{d\kappa}, 1 \right\} \right) \\ &\geq \frac{\eta\mu}{4d} \cdot \frac{5}{8} \\ &= \frac{5\mu\gamma}{2^{12}dR^2} \cdot \min \left\{ \frac{d\kappa}{n\kappa'}, 1 \right\} \\ &= \frac{5}{2^{12}d\kappa} \cdot \min \left\{ \frac{d\kappa}{n\kappa'}, 1 \right\} \\ &= \frac{5}{2^{12}m}, \end{aligned}$$

and

$$\frac{\tau\gamma}{4nd} - \frac{12R^2\eta\tau}{nd} \geq \frac{5}{2^{12}m}.$$

Combining all these facts into (23), and rearranging terms, we can get the recursive relationship:

$$\begin{aligned} &\mathbb{E} \left[\frac{\|x^{t+1} - x^*\|^2}{\eta} + \frac{\|y^{t+1} - y^*\|^2}{\tau} \middle| \mathcal{F}_t \right] \\ &\leq \left(1 - \frac{5}{2^{12}m} \right) \left[\frac{\|x^t - x^*\|^2}{\eta} + \frac{\|y^t - y^*\|^2}{\tau} \right] + \frac{1}{2^{11}m} \left[\frac{\|\tilde{x}^k - x^*\|^2}{\eta} + \frac{\|\tilde{y}^k - y^*\|^2}{\tau} \right] \\ &= \left(1 - \frac{5}{2^{12}m} \right) \left[\frac{\|x^t - x^*\|^2}{\eta} + \frac{\|y^t - y^*\|^2}{\tau} \right] + \frac{1}{2^{11}m} \Delta_k. \end{aligned}$$

Now, we recursively apply this inequality, and use the law of total expectation, we can finally obtain:

$$\begin{aligned}
& \mathbb{E}[\Delta_{k+1}] \\
&= \mathbb{E} \left[\frac{\|x^T - x^*\|^2}{\eta} + \frac{\|y^T - y^*\|^2}{\tau} \right] \\
&\leq \left(1 - \frac{5}{2^{12}m}\right) \mathbb{E} \left[\frac{\|x^{T-1} - x^*\|^2}{\eta} + \frac{\|y^{T-1} - y^*\|^2}{\tau} \right] + \frac{1}{2^{11}m} \cdot \Delta_k \\
&\leq \dots \\
&\leq \left(1 - \frac{5}{2^{12}m}\right)^T \cdot \left[\frac{\|x^0 - x^*\|^2}{\eta} + \frac{\|y^0 - y^*\|^2}{\tau} \right] + \frac{1}{2^{11}m} \sum_{t=0}^{T-1} \left(1 - \frac{5}{2^{12}m}\right)^t \Delta_k \\
&= \left(1 - \frac{5}{2^{12}m}\right)^T \cdot \Delta_k + \frac{1}{2^{11}m} \cdot \frac{1 - [1 - 5/(2^{12}m)]^T}{1 - [1 - 5/(2^{12}m)]} \Delta_k \\
&\leq \left(1 - \frac{5}{2^{12}m}\right)^T \cdot \Delta_k + \frac{2}{5} \Delta_k.
\end{aligned}$$

Therefore, to finish the proof, we only need to make sure $(1 - 5/(4096m))^T \leq 1/5$, which can be guaranteed if choosing some large enough $T \geq \Theta(m) = \Theta(\max\{d\kappa, n\kappa'\})$. \square




Figure 3: Numerical results on the problem of squared-hinge loss SVM. The y -axis is primal sub-optimality.

Appendix B Extra Numerical Experiments

In this part, we will show some extra experiment results. The experiment setting is basically same to Section 5. The only difference is that we change the model to support vector machine (SVM) with squared-hinge loss, i.e.,

$$\phi_i(u) = \max\{0, 1 - b_i u\}^2$$

where $b_i \in \{\pm 1\}$ is the class label. Note that the corresponding conjugate function is

$$\phi_i^*(y) = \begin{cases} b_i y + \frac{y^2}{4}, & \text{if } b_i y \leq 0, \\ +\infty, & \text{if } b_i y > 0, \end{cases}$$

whose proximal mapping has simple closed-form solution, and does not need to be solved iteratively.

The results are shown in Figure 3. The performance of all methods are similar to that in Section 5.

Besides, we further report the running time of all methods in Table 2. We can observe that SPD1 and SPD1-VR takes longer time than SGD and SVRG for each pass of data. We think there are mainly two reasons which result in such phenomenon: 1) SPD1 and SPD1-VR involves much more loops, which may incur computation overhead (our codes are written in Julia), while the updates of SGD and SVRG can be implemented in vector forms, and the vector operations are conducted by some highly-optimized computation libraries like OpenBLAS; 2) sampling only one scalar from the data matrix each time is cache-unfriendly for computers. However, we believe these two issues can be solved. For example, the former issue can be tackled by using faster programming languages like C/C++. While for the latter one, one possible way to solve it is adopting mini-batch versions of SPD1 and SPD1-VR, which sample a batch of continuous coordinates instead of just one in each iteration.

Table 2: Running time required by different methods for one pass of data (in seconds)

Methods	colon-cancer	gisette	rcv1.binary
SPD1	0.012	2.70	128.4
SPD1-VR	0.013	3.62	105.5
SGD	0.006	1.07	34.8
SVRG	0.006	1.37	45.1