

---

# Stochastic Optimization for Non-convex Inf-Projection Problems

---

Yan Yan<sup>1</sup> Yi Xu<sup>2</sup> Lijun Zhang<sup>3</sup> Xiaoyu Wang<sup>4</sup> Tianbao Yang<sup>1</sup>

## Abstract

In this paper, we study a family of non-convex and possibly non-smooth inf-projection minimization problems, where the target objective function is equal to minimization of a joint function over another variable. This problem include difference of convex (DC) functions and a family of bi-convex functions as special cases. We develop stochastic algorithms and establish their first-order convergence for finding a (nearly) stationary solution of the target non-convex function under different conditions of the component functions. To the best of our knowledge, this is the first work that comprehensively studies stochastic optimization of non-convex inf-projection minimization problems with provable convergence guarantee. Our algorithms enable efficient stochastic optimization of a family of non-decomposable DC functions and a family of bi-convex functions. To demonstrate the power of the proposed algorithms we consider an important application in variance-based regularization. Experiments verify the effectiveness of our inf-projection based formulation and the proposed stochastic algorithm in comparison with previous stochastic algorithms based on the min-max formulation for achieving the same effect.

## 1. Introduction

In this paper, we consider a family of non-convex and possibly non-smooth problems with the following structure

$$\min_{x \in X} F(x) := \{g(x) + \min_{y \in \text{dom}(h)} h(y) - \langle y, \ell(x) \rangle\}, \quad (1)$$

where  $X \subseteq \mathbb{R}^d$  is a closed convex set,  $g : X \rightarrow \mathbb{R}$  is lower-semicontinuous,  $h : \text{dom}(h) \rightarrow \mathbb{R}$  is uniformly convex,  $\ell :$

<sup>1</sup>University of Iowa <sup>2</sup>DAMO Academy, Alibaba Group <sup>3</sup>Nanjing University <sup>4</sup>The Chinese University of Hong Kong (Shenzhen). Correspondence to: Yan Yan <yan-yan-2@uiowa.edu>, Tianbao Yang <tianbao-yang@uiowa.edu>.

$X \rightarrow \mathbb{R}^m$  is a lower-semicontinuous differentiable mapping, and  $\langle \cdot, \cdot \rangle$  is the inner product. The requirement of uniform convexity on  $h$  is to ensure the inner minimization problem is well defined and its solution is unique (cf. Section 2). Define  $f(x, y) = g(x) + h(y) - \langle y, \ell(x) \rangle$ , the objective function  $F(x)$  is called the inf-projection of  $f(x, y)$  in the literature. When  $g$  is convex, depending on  $\text{dom}(h)$ , the two subfamilies of above problem (1) deserve more discussion: difference of convex (DC) and bi-convex functions.

**DC functions.** When  $g$  is convex and  $\text{dom}(h) \subseteq \mathbb{R}_+^m$  and  $\ell$  is convex<sup>1</sup>, the inf-projection minimization problem (1) is equivalent to the following DC functions,

$$\min_{x \in X} F(x) = \left\{ g(x) - h^*(\ell(x)) \right\}, \quad (2)$$

where  $h^*$  denotes the convex conjugate function of  $h$ , the convexity of the second component  $h^*(\ell(x))$  is following the composition rule of convexity (Boyd & Vandenberghe, 2004)<sup>2</sup>. Minimizing DC functions has wide applications in machine learning and statistics (Nitanda & Suzuki, 2017; Kiryo et al., 2017). Although stochastic algorithms for DC problems have been considered recently (Nitanda & Suzuki, 2017; Xu et al., 2018a; Thi et al., 2017), working with the inf-projection minimization (1) is preferred when  $h^*(\ell(x))$  is *non-decomposable* such that an unbiased stochastic gradient of  $h^*(\ell(x))$  is not easily accessible as that of  $\langle y, \ell(x) \rangle$  in (1). Inspired by this scenario, let us particularly consider an important instance variance-based regularization. It refers to a learning paradigm that minimizes the empirical loss and its variance simultaneously, by which a better bias-variance trade-off may be achieved (Maurer & Pontil, 2009). To give a condensed understanding of its connection to the inf-projection formulation, we can re-formulate the problem (cf. the details and comparison with a related convex objective of (Namkoong & Duchi, 2017) in Section 5):

$$\min_{x \in X} \frac{1}{n} \sum_{i=1}^n l_i(x) + \lambda \frac{1}{2n} \sum_{i=1}^n (l_i(x))^2 - \frac{\lambda}{2} \left( \frac{1}{n} \sum_{i=1}^n l_i(x) \right)^2, \quad (3)$$

where  $l_i(x) : X \rightarrow \mathbb{R}_+$  is the loss function of a model  $x$  on the  $i$ -th example and  $\lambda > 0$  is a regularization pa-

<sup>1</sup> $\text{dom}(h) \subseteq \mathbb{R}^m$  and  $\ell$  is concave can be transferred to the considered case by a variable change.

<sup>2</sup>Note that  $h^*$  is monotonically increasing iff  $\text{dom}(h) \subseteq \mathbb{R}_+^m$ .

Table 1. Summary of results for finding a (nearly)  $\epsilon$ -stationary solution in this work under different conditions of  $g$ ,  $h$  and  $\ell$ . SM means smoothness, Lip. means Lipschitz continuous, Diff means differentiable, MO means monotonically increasing or decreasing for  $h^*$ , CVX means convex, and UC means  $p$ -uniformly convex ( $p \geq 2$ ),  $v = 1/(p - 1)$ .

$g$	$h(h^*)$	$\ell$	Alg.	Mini-Batch	Compl.
SM	UC & simple	SM & Lip	MSPG (Section 3)	Yes	$O(1/\epsilon^{4/v})$
SM & CVX	UC (MO)	Diff & Lip & CVX	St-SPG (Section 4)	No	$O(1/\epsilon^{4/v})$
Lip & CVX	UC (MO)	SM & Lip & CVX	St-SPG (Section 4)	No	$O(1/\epsilon^{4/v})$

parameter. The above problem (3) is a special case of (2) by regarding  $g(x)$  as the sum of the first two terms,  $\ell(x) = \frac{1}{n} \sum_{i=1}^n l_i(x)$  and  $h^*(s) = \frac{\lambda}{2} s^2$ . By noting the convex conjugate  $-\frac{1}{2} \left( \frac{1}{n} \sum_{i=1}^n l_i(x) \right)^2 = \min_{y \geq 0} \frac{1}{2} y^2 - \langle y, \frac{1}{n} \sum_{i=1}^n l_i(x) \rangle$ , the above problem can be viewed as a special case of (1). In this way, computing a stochastic gradient of  $f(x, y)$  in terms of  $x$  can be done based on one sampled loss function  $l_i(x)$ . It is easier than computing an unbiased stochastic gradient of  $h^*(\ell(x))$  in (3) that requires at least two sampled loss functions.

**Bi-convex functions.** When  $g$  is convex and  $\text{dom}(h) \subseteq \mathbb{R}^m$  and  $\ell$  is convex, the inf-projection minimization problem (1) reduces to minimization of a bi-convex function. In particular,  $f(x, y)$  is convex in terms of  $x$  for every fixed  $y \in \text{dom}(h)$  and  $f(x, y)$  is convex in terms of  $y$  for every fixed  $x \in X$ . The concerned family of bi-convex functions also find some applications in machine learning and computer vision (Kumar et al., 2010; Shah et al., 2016). For example, the self-paced learning method proposed by (Kumar et al., 2010) needs to solve the following bi-convex problem  $\min_{\mathbf{w} \in \mathbb{R}^d, \mathbf{v} \in \{0, 1\}^n} r(\mathbf{w}) + \sum_{i=1}^n v_i f_i(\mathbf{w}) - \frac{1}{K} \sum_{i=1}^n v_i$ , where  $r$  and  $f_i$  are convex in  $\mathbf{w}$ ,  $v_i = 0$  if  $f_i(\mathbf{w}) > \frac{1}{K}$  and  $v_i = 1$  if  $f_i(\mathbf{w}) < \frac{1}{K}$ , which can be covered by (1). Although deterministic optimization methods (e.g., proximal alternating linearized minimization and its variants (Bolte et al., 2014; Davis et al., 2016)) and their convergence theory have been studied for minimizing a bi-convex function (Gorski et al., 2007), algorithms and convergence theory for stochastic optimization of a bi-convex function remains under-explored especially when we are interested in the convergence respect to the target function  $F(x)$ . A special case that belongs to both DC and Bi-convex functions is when  $\ell(x) = Ax$ , and  $\text{dom}(h)$  can be any convex set.

A naive idea to tackle (1) is by alternating minimization or block coordinate descent, i.e., alternatively solving the inner minimization problem over  $y$  given  $x$  and then updating  $x$  by certain approaches (e.g., stochastic gradient descent) (Bolte et al., 2014; Davis et al., 2016; Hong et al., 2015; Xu & Yin, 2013; Driggs et al., 2020). However, this approach suffers from two issues: (i) solving the inner minimization might

not be a trivial task (e.g., solving the inner minimization problem related to (3) requires passing  $n$  examples once); (ii) the target objective function  $F(x)$  is not necessarily a smooth function or a convex function, which makes the convergence analysis challenging. Additionally, their convergence analysis focus on  $f(x, y)$  instead of  $F(x)$ . In this paper, the main question that we tackle is: how to design efficient stochastic algorithms using simple updates for both  $x$  and  $y$  to enjoy a provable convergence guarantee in terms of finding a stationary point of  $F(x)$ ? **Our contributions** are summarized below:

- First, we consider the case when  $g$  and  $\ell$  are smooth but not necessarily convex and  $h$  is a simple function whose proximal mapping is easy to compute. Under the condition that  $\ell$  is Lipschitz continuous, we prove the convergence of mini-batch stochastic proximal gradient method (MSPG) with increasing mini-batch size that employ parallel stochastic gradient updates for  $x$  and  $y$ , and establish the convergence rate.
- Second, we consider the cases when  $g$  and  $\ell$  are not necessarily smooth but convex, and  $h$  is not necessarily a simple function (corresponding to DC and bi-convex functions). We develop an algorithmic framework that employs a suitable stochastic algorithm for solving strongly convex functions in a stagewise manner. We analyze the convergence rates for finding a (nearly) stationary point when employing the stochastic proximal gradient (SPG) method at each stage, resulting St-SPG. The complexity results of our algorithms under different conditions of  $g$ ,  $h$  and  $\ell$  are shown in Table 1.

The novelty and significance of our results are (i) this is the first work that comprehensively studies the stochastic optimization of a non-smooth non-convex inf-projection problem; (ii) the application of the inf-projection formulation to variance-based regularization demonstrates much faster convergence of our algorithms comparing with existing algorithms based on a min-max formulation.

## 2. Preliminaries

Let us first present some notations. We let  $\|\cdot\|$  denote the Euclidean norm of a vector and the spectral norm of a matrix. We use  $\xi$  to denote some random variable. Given a function  $g : \mathbb{R}^d \rightarrow \mathbb{R}$ , we denote the Fréchet subgradients and limiting Fréchet gradients by  $\hat{\partial}g$  and  $\partial g$  respectively, i.e., at  $x$ ,  $\hat{\partial}g(x) = \{y \in \mathbb{R}^d : \lim_{x' \rightarrow x} \inf \frac{g(x) - g(x') - y^\top(x - x')}{\|x - x'\|} \geq 0\}$ , and  $\partial g(x) = \{y \in \mathbb{R}^d : \exists x_k \xrightarrow{g} x, v_k \in \partial g(x_k), v_k \rightarrow y\}$ . Here  $x_k \xrightarrow{g} x$  represents  $x_k \rightarrow x$  and  $g(x_k) \rightarrow g(x)$ . When the function  $g$  is differentiable, the subgradients ( $\hat{\partial}g$  and  $\partial g$ ) reduce to the standard gradient  $\nabla g$ . It is known that  $\hat{\partial}g(x) \subset \partial g(x)$ ,  $\hat{\partial}g(x) = \{\nabla g(x)\}$  if  $g(x)$  is differential and  $\partial g(x) = \{\nabla g(x)\}$  if  $g(x)$  is continuously differential. We denote by  $\partial_x g(x, y)$  the partial derivative in the direction of  $x$  and  $\partial g(x, y) = (\partial_x g(x, y), \partial_y g(x, y))^\top$ . In this paper, we will prove the convergence in terms of the limiting gradient. But all results can be extended to the Fréchet subgradients.

Let  $\nabla \ell(x) \in \mathbb{R}^{m \times d}$  denote the Jacobian matrix of the differentiable mapping  $\ell(x)$ .  $\ell$  is said  $G_\ell$ -Lipschitz continuous if  $\|\nabla \ell(x)\| \leq G_\ell$ . A differentiable function  $f(\cdot)$  has  $(L, v)$ -Hölder continuous gradient if  $\|\nabla f(x_1) - \nabla f(x_2)\| \leq L\|x_1 - x_2\|^v$  holds for some  $v \in (0, 1]$  and  $L > 0$ . When  $v = 1$ , it is known as  $L$ -smooth function. If  $\nabla f$  is Hölder continuous, then it holds  $f(x_1) - f(x_2) \leq \langle \nabla f(x_2), x_1 - x_2 \rangle + \frac{L}{1+v}\|x_1 - x_2\|^{1+v}$ . A related condition is uniform convexity. A function  $f(\cdot)$  is  $(\varrho, p)$ -uniformly convex where  $p \geq 2$ , if  $f(x_1) - f(x_2) \geq \langle \partial f(x_2), x_1 - x_2 \rangle + \frac{\varrho}{2}\|x_1 - x_2\|_2^p$ . When  $p = 2$ , it is known as strong convexity. If  $f$  is  $(\varrho, p)$ -uniformly convex, then the following inequality holds

$$\begin{aligned} \varrho\|x_1 - x_2\|^p &\leq \langle \partial f(x_1) - \partial f(x_2), x_1 - x_2 \rangle \\ &\leq \|\partial f(x_1) - \partial f(x_2)\| \cdot \|x_1 - x_2\|. \end{aligned} \quad (4)$$

It is obvious that a uniformly convex function has a unique minimizer. If  $f$  is uniformly convex, then its convex conjugate  $f^*$  has Hölder continuous gradient and vice versa, which is summarized in the following lemma.

**Lemma 1.** *Let  $f$  be differentiable and  $\nabla f$  be  $(L, v)$ -Hölder continuous where  $v \in (0, 1]$ . Then  $f^*$  is  $(\varrho, p)$ -uniformly convex with  $p = 1 + \frac{1}{v}$  and  $\varrho = \frac{2v}{1+v}(1/L)^{\frac{1}{v}}$ . If  $f$  is  $(\varrho, p)$ -uniformly convex, then  $f^*$  has  $(L, v)$ -Hölder continuous gradient with  $L = (1/\varrho)^{1/(p-1)}$  and  $v = 1/(p-1)$ .*

Next we discuss the convergence measure for the considered inf-projection problem. Let  $f_h(x) = \min_y h(y) - y^\top \ell(x)$ . If  $h$  is uniformly convex, let  $y^*(x) = \arg \min_y h(y) - y^\top \ell(x)$  denote the unique minimizer. In this way, under a regularity condition that  $h(y) - y^\top \ell(x)$  is level-bounded in  $y$  uniformly in  $x$ , then Theorem 10.58 of (Rockafellar & Wets, 2009) implies  $\partial f_h(x) = -\nabla \ell(x)^\top y^*(x)$ . Then  $\partial F(x) = \partial g(x) - \nabla \ell(x)^\top y^*(x)$  under the smoothness or

---

### Algorithm 1 MSPG

---

- 1: **Input:** initialized  $x_1, y_1$ .
  - 2: **for**  $t = 1, \dots, T$  **do**
  - 3:   Compute mini-batch stochastic partial gradients  $\tilde{\nabla}_x f_0^{(t)} = \frac{1}{m_t} \sum_{i=1}^{m_t} \nabla_x f_0(x_t, y_t; \xi_i)$  and  $\tilde{\nabla}_y f_0^{(t)} = \frac{1}{m_t} \sum_{i=1}^{m_t} \nabla_y f_0(x_t, y_t; \xi_i)$
  - 4:    $x_{t+1} = \Pi_X[x_t - \eta \tilde{\nabla}_x f_0^{(t)}]$
  - 5:    $y_{t+1} = P_{\eta h}[y_t - \eta \tilde{\nabla}_y f_0^{(t)}]$
  - 6: **end for**
  - 7: **Output:**  $w_\tau = (x_\tau, y_\tau)$ , where  $\tau \in \{1, \dots, T\}$  is randomly sampled
- 

convexity condition of  $g$ , which allows us to connect  $\partial F(x)$  by  $\partial g(x) - \nabla \ell(x)^\top y^*(x)$ . In this paper, we aim to find a solution that is  $\epsilon$ -stationary or nearly  $\epsilon$ -stationary of  $F$ , which are defined as follows.

**Definition 1.** *A solution  $x$  satisfying  $\text{dist}(0, \partial F(x)) \leq \epsilon$  is called an  $\epsilon$ -stationary point of  $F$ . A solution  $x$  is called a nearly  $\epsilon$ -stationary if there exists  $z$  and a constant  $c > 0$  such that  $\|z - x\| \leq c\epsilon$  and  $\text{dist}(0, \partial F(z)) \leq \epsilon$ .*

Particularly, nearly stationarity has been used to measure the convergence for non-smooth non-convex optimization in the literature (Davis & Grimmer, 2017; Davis & Drusvyatskiy, 2018b;a; Chen et al., 2018; Xu et al., 2018a).

Before ending this section, we state basic assumptions below. For simplicity, here all variance bounds are denoted by  $\sigma^2$ . Additional conditions regarding  $g, h$  and  $\ell$  are presented in individual theorems.

**Assumption 1.** *For the problem (1) we assume:*

- (i)  $h^*$  has  $(L_{h^*}, v)$ -Hölder continuous gradient, and  $\ell$  is continuously differentiable;
- (ii) Let  $\partial g(x; \xi_g)$  denote a stochastic gradient of  $g(x)$ . If  $g(x)$  is smooth, assume  $\mathbb{E}[\|\partial g(x; \xi_g) - \partial g(x)\|^2] \leq \sigma^2$ , otherwise assume  $\mathbb{E}[\|\partial g(x; \xi)\|^2] \leq \sigma^2$  for  $x \in X$ ;
- (iii) Let  $\ell(x; \xi_\ell)$  denote a stochastic version of  $\ell(x)$  and assume  $\mathbb{E}[\|\ell(x; \xi_\ell) - \ell(x)\|^2] \leq \sigma^2$ . If  $\ell(x)$  is smooth, assume  $\mathbb{E}[\|\nabla \ell(x; \xi_\ell) - \nabla \ell(x)\|^2] \leq \sigma^2$ , otherwise assume  $\mathbb{E}[\|\nabla \ell(x; \xi_\ell)\|^2] \leq \sigma^2$  for  $x \in X$ ;
- (iv) Let  $\partial h(y; \xi_h)$  denote a stochastic gradient of  $h(y)$  and assume  $\mathbb{E}[\|\partial h(y; \xi_h)\|^2] \leq \sigma^2$  for  $y \in \text{dom}(h)$ ;
- (v)  $\max_{x \in X, y \in \text{dom}(h)} f(x, y) - \min_{x \in X, y \in \text{dom}(h)} f(x, y) \leq M$ .

## 3. Mini-batch Stochastic Gradient Methods For Smooth Functions

In this section, we consider the case when  $g$  and  $\ell$  are smooth functions but not necessarily convex. Please note that the target function  $F$  is still not necessarily smooth and is non-convex. We assume  $h$  is simple such that its proximal map-

ping defined by  $P_{\eta h}[\hat{y}] = \arg \min_y h(y) + \frac{1}{2\eta} \|y - \hat{y}\|^2$  is easy to compute. Let  $f_0(x, y) = g(x) - y^\top \ell(x)$ . The key idea of our first algorithm is that we treat  $f(x, y) = f_0(x, y) + h(y) + I_X(x)$  as a function of the joint variable  $w = (x, y)$ , which consists of a smooth component  $f_0$  and a non-smooth component  $h(y) + I_X(x)$ . Hence, we can employ mini-batch stochastic proximal gradient (MSPG) method to minimize  $f(x, y)$  based on stochastic gradients of  $f_0(w)$  denoted by  $\nabla f_0(w; \xi)$  for a random variable  $\xi$ . The detailed steps of MSPG are shown in Algorithm 1. At each iteration, stochastic partial gradients w.r.t.  $x$  and  $y$  are computed and used for updating.

Although the convergence of MSPG for  $f(w)$  has been considered in literature of composite optimization (Ghadimi et al., 2016) or alternating minimization (Hong et al., 2015; Xu & Yin, 2013; Driggs et al., 2020), there is still a gap when applying the existing convergence result, since we are interested in the convergence analysis of  $\text{dist}(0, \partial F(x))$ , rather than  $f_0(w)$ . In the following, we fill this gap by four main steps. In brief, first, we establish the joint smoothness of  $f_0$  in  $(x, y)$  by Lemma 2. Then, based on Lemma 2, we derive the convergence of  $\text{dist}(0, \partial f(w))$  in Proposition 1. Next, Lemma 5 connects  $\text{dist}(0, \partial f(w))$  to  $\text{dist}(0, \partial F(x))$ . Finally, the convergence of  $\text{dist}(0, \partial F(x))$  is achieved (Theorem 2).

**Lemma 2.** *Suppose  $g(x)$  is  $L_g$  smooth,  $\ell$  is  $G_\ell$ -Lipschitz continuous and  $L_\ell$ -smooth, and  $\max_{y \in \text{dom}(h)} \|y\| \leq D_y$ . Then  $f_0(x, y)$  is smooth over  $(x, y) \in X \times \text{dom}(h)$ , i.e.,*

$$\begin{aligned} & \|\nabla_x f_0(x, y) - \nabla_x f_0(x', y')\|_2^2 \\ & \quad + \|\nabla_y f_0(x, y) - \nabla_y f_0(x', y')\|_2^2 \\ & \leq L^2 (\|x - x'\|_2^2 + \|y - y'\|_2^2), \end{aligned}$$

where  $L = \sqrt{\max(2L_g^2 + 4L_\ell^2 D_y^2 + G_\ell^2, 4G_\ell^2)}$ .

Based on the joint smoothness of  $f_0$  in  $(x, y)$ , we can establish the convergence of MSPG in terms of  $\text{dist}(0, \partial f(x_\tau, y_\tau))$  in the following proposition. Note that this convergence result in terms of  $\text{dist}(0, \partial f(x_\tau, y_\tau))$  is stronger than that in (Ghadimi et al., 2016) in terms of proximal gradient, which follows the analysis in (Xu et al., 2019).

**Proposition 1.** *Under the same conditions as in Lemma 2 and suppose the stochastic gradient has bounded variance  $\mathbb{E}[\|\nabla f_0(w; \xi) - \nabla f_0(w)\|^2] \leq \sigma_0^2$ , run MSPG with  $\eta = \frac{c}{L}$  ( $0 < c < \frac{1}{2}$ ) and a sequence of mini-batch sizes  $m_t = b(t + 1)$  for  $t = 0, \dots, T - 1$ , where  $b > 0$  is a constant, then the output  $w_\tau$  of Algorithm 1 satisfies*

$$\mathbb{E}[\text{dist}(0, \partial f(w_\tau))^2] \leq \frac{c_1 \sigma_0^2 (\log(T) + 1)}{bT} + \frac{c_2 \Delta}{\eta T},$$

where  $c_1 = \frac{2c(1-2c)+2}{c(1-2c)}$  and  $c_2 = \frac{6-4c}{1-2c}$ .

The next lemma establishes the relation between  $\text{dist}(0, \partial f(w_\tau))$  and  $\text{dist}(0, \partial F(x_\tau))$ , allowing us to bridge the convergence of  $\text{dist}(0, \partial F(x_\tau))$  by employing that of  $\text{dist}(0, \partial f(w_\tau))$ .

**Lemma 3.** *Under the same conditions as in Lemma 2 and  $h^*$  has  $(L_{h^*}, v)$ -Hölder continuous gradient. Then for any  $(\tilde{x}, \tilde{y}) \in X \times \text{dom}(h)$ , we have*

$$\begin{aligned} \text{dist}(0, \partial F(\tilde{x})) & \leq \|\nabla_x f(\tilde{x}, \tilde{y})\|_2 \\ & \quad + G_\ell \left( \frac{(1+v)}{2v} \right)^v L_{h^*} \text{dist}(0, \partial_y f(\tilde{x}, \tilde{y}))^v. \end{aligned}$$

Finally, combining the above results, we can state the main result in this section regarding the convergence of MSPG in terms of the concerned  $\text{dist}(0, \partial F(x_\tau))$  as follows.

**Theorem 2.** *Suppose the same conditions as in Lemma 2 and Assumption 1 hold. Algorithm 1 guarantees that  $\mathbb{E}[\text{dist}(0, \partial F(x_\tau))^2] \leq O(1/T^v)$ . To ensure  $\mathbb{E}[\text{dist}(0, \partial F(x_\tau))] \leq \epsilon$ , we can set  $T = O(1/\epsilon^{2/v})$ . The total complexity is  $O(1/\epsilon^{4/v})$ .*

## 4. Stochastic Algorithms for Non-Smooth Functions

In this section, we consider the case when  $g$  or  $\ell$  are not necessarily smooth but are convex. We also assume  $h^*$  is monotonic, i.e.,  $\text{dom}(h) \subseteq \mathbb{R}_+^m$  or  $\text{dom}(h) \subseteq \mathbb{R}_-^m$ . In the former case, the objective function belongs to **DC functions**, and in the latter case the objective function belongs to **Bi-Convex** functions. Please note that the target function  $F$  is still not necessarily convex and is non-smooth. The proposed algorithm is inspired by the stagewise stochastic DC algorithm proposed in (Xu et al., 2018a) but with some major changes. Let us first briefly discuss the main idea and logic behind the proposed algorithm. There are two difficulties that we need to tackle: (i) non-smoothness and non-convexity in terms of  $x$ , (ii) minimization over  $y$ .

To tackle the first issue, let us assume the optimal solution  $y^*(x) = \arg \min_y h(y) - y^\top \ell(x)$  given  $x$  is available. Then the problem regarding  $x$  becomes:

$$\min_{x \in X} g(x) - y^*(x)^\top \ell(x) \quad (5)$$

When  $\text{dom}(h) \subseteq \mathbb{R}_+^m$  (corresponding to a DC function), the above problem is still non-convex. In order to obtain a provable convergence guarantee, we consider the following strongly convex problem from some  $\gamma > 0$  and  $x_0 \in X$ , whose objective function is an upper bound of the function in (5) at  $x_0$ :

$$\begin{aligned} P(x_0) = \arg \min_{x \in X} \left\{ g(x) - y_{x_0}^*{}^\top (\ell(x_0) \right. \\ \left. + \nabla \ell(x_0)(x - x_0)) + \frac{\gamma}{2} \|x - x_0\|^2 \right\}. \quad (6) \end{aligned}$$

Note  $P(x_0)$  is uniquely defined due to strong convexity. If  $x_0 = P(x_0)$  it can be shown that  $x_0$  is the critical point of  $F(x)$ , i.e.,  $0 \in \partial F(x_0) = \partial g(x_0) - \nabla \ell(x_0)^\top y^*(x_0)$ . Then we can iteratively solve the fixed-point problem  $x = P(x)$  until it converges.

When  $\text{dom}(h) \subseteq \mathbb{R}^m$  (corresponding to a Bi-convex function), we can simply consider the following strongly convex problem:

$$P(x_0) = \arg \min_{x \in X} g(x) - y^*(x_0)^\top \ell(x) + \frac{\gamma}{2} \|x - x_0\|^2.$$

A remaining issue in the above approach is that  $y^*(x_0)$  is assumed available, which is related to the second issue mentioned above. It may not be easy to obtain an exact minimizer  $y^*(x_0)$  given a  $x_0$ . To this end, we can employ an iterative stochastic algorithm to optimize  $\min_y h(y) - y^\top \ell(x_0)$  approximately given  $x_0$ , and obtain an inexact solution  $\hat{y}(x_0)$  such that  $h(\hat{y}(x_0)) - \hat{y}(x_0)^\top \ell(x_0) - h(y^*(x_0)) - y^*(x_0)^\top \ell(x_0) \leq \varepsilon$  for some approximation error  $\varepsilon$ . Then, we combine these two pieces together, i.e., replacing  $y^*(x_0)$  in the definition of  $P(x_0)$  with  $\hat{y}(x_0)$ , and employing a stochastic algorithm to solve the fixed-point equation by  $x \leftarrow \hat{P}(x)$ , where  $\hat{P}(x)$  is an approximation of  $P(x)$ . Therefore, we have two sources of approximation error — one from using  $\hat{y}_x$  instead of  $y^*(x)$  and another one from solving the minimization problem of  $x$  inexactly. Our analysis is to show that with well-controlled approximation error, we can still achieve provable convergence guarantee.

For the sake of presentation, let us first introduce some important notations by considering different conditions of DC and bi-convex functions. For the  $k$ -th stage of St-SPG, define

$$f_x^k(x) = g(x) - y_k^\top (\ell(x_k) + \nabla \ell(x_k)(x - x_k)),$$

for  $\text{dom}(h) \subseteq \mathbb{R}_+^m$ ,

and

$$f_x^k(x) = g(x) - y_k^\top \ell(x), \text{ for } \text{dom}(h) \subseteq \mathbb{R}^m.$$

A stochastic gradient of  $f_x^k(x)$  can be computed by  $\partial g(x; \xi_g) - \nabla \ell(x_k; \xi_\ell)^\top y_k$  for  $\text{dom}(h) \subseteq \mathbb{R}_+^m$  or  $\partial g(x; \xi_g) - \nabla \ell(x; \xi_\ell)^\top y_k$  for  $\text{dom}(h) \subseteq \mathbb{R}^m$ . For both conditions, let

$$f_y^k(y) = h(y) - y^\top \ell(x_{k+1}),$$

$$R_x^k(x) = \frac{\gamma}{2} \|x - x_k\|^2, \quad R_y^k(y) = \frac{\mu}{2} \|y - y_k\|^2.$$

A stochastic gradient of  $f_y^k(y)$  can be computed by  $\partial h(y; \xi_h) - \ell(x_{k+1}; \xi'_\ell)$ , where  $\xi_g, \xi_\ell, \xi_h, \xi'_\ell$  denote independent random variables.

The proposed algorithm is shown in Algorithm 2 named St-SPG, which employs SPG in Algorithm 3 to solve the

---

**Algorithm 2** St-SPG
 

---

- 1: Initialize  $x_1 \in X, y_1 \in \text{dom}(h)$
  - 2: Set a sequence of integers  $T_x^k, T_y^k$  and numbers  $\gamma, \mu$
  - 3: **for**  $k = 1, \dots, K$  **do**
  - 4:    $x_{k+1} = \text{SPG}(f_x^k, R_x^k, x_k, X, T_x^k, \gamma)$
  - 5:    $y_{k+1} = \text{SPG}(f_y^k, R_y^k, y_k, \text{dom}(h), T_y^k, \mu)$
  - 6: **end for**
- 

---

**Algorithm 3** SPG( $f, R, z_1, \Omega, T, \gamma$ )
 

---

- 1: Set  $\eta_t$  according to  $\gamma$
  - 2: **for**  $t = 1, \dots, T$  **do**
  - 3:    $z_{t+1} = \arg \min_{z \in \Omega} \partial f(z_t; \xi_t)^\top z + R(z) + \frac{1}{2\eta_t} \|z - z_t\|^2$
  - 4: **end for**
  - 5: **Output**  $\hat{z}_T = \frac{\sum_{t=1}^T z_t}{\sum_{t=1}^T 1}$
- 

subproblems of  $x$  and  $y$  in a stagewise manner.  $x$  and  $y$  share the same update method SPG, so we can summarize it in general notations. To this end, let us consider the convergence of SPG for solving  $H(z) = f(z) + R(z)$ , where  $f(z)$  is a convex function and  $R(z) = \frac{\gamma}{2} \|z - z_1\|^2$  is a strongly convex function. Its convergence has been considered in many previous works. Here, we adopt the results derived in (Xu et al., 2018a) to establish the convergence of St-SPG under different conditions of  $g$  and  $\ell$  as follows.

**Proposition 2.** *Let  $H(z) = f(z) + R(z)$  where  $R(z) = \frac{\gamma}{2} \|z - z_1\|^2$  is  $\gamma$ -strongly convex. If  $f(z)$  is  $L$ -smooth and  $\mathbb{E}[|\nabla f(z; \xi) - \nabla f(z)|^2] \leq \sigma^2$  and  $\gamma \geq 3L$ , then by setting  $\eta_t = 3/(\gamma(t+1))$  SPG guarantees that*

$$\mathbb{E}[H(\hat{z}_T) - H(z_*)] \leq \frac{4\gamma \|z_* - z_1\|^2}{3T(T+3)} + \frac{6\sigma^2}{(T+3)\gamma}.$$

*If  $f$  is non-smooth with  $\mathbb{E}[|\nabla f(z; \xi)|^2] \leq \sigma^2$ , then by setting  $\eta_t = 4/(\gamma t)$  SPG guarantees that*

$$\mathbb{E}\left[H(\hat{z}_T) - H(z_*)\right] \leq \frac{\gamma \|z_* - z_1\|^2}{4T(T+1)} + \frac{17\sigma^2}{\gamma(T+1)},$$

where  $z_* = \arg \min_{z \in \Omega} H(z)$ .

With the above proposition, we can apply the above convergence guarantee of SPG for  $f_x^k(x) + R_x^k(x)$  and  $f_y^k(y) + R_y^k(y)$ . Then define  $v_k$  and  $u_k$  as the optimal solutions to the subproblems of  $x$  and  $y$  at the  $k$ -th stage, respectively:

$$v_k = \arg \min_{x \in X} f_x^k(x) + R_x^k(x),$$

$$u_k = \arg \min_y f_y^k(y) + R_y^k(y).$$

We can establish the following result regarding the convergence of St-SPG related to fixed-point convergence ( $x_{\tau+1} - x_\tau$ ), and also the minimization error of  $P(x)$ ,

i.e.,  $\|x_{\tau+1} - v_\tau\|$ , for a randomly sampled index  $\tau \in \{1, \dots, K\}$ . We have boundedness assumptions on  $y$  and  $\ell$  below to guarantee the boundedness of the second moment of stochastic gradients, which can be implied by assuming the domain  $X$  is a compact set and  $\text{dom}(h)$  is bounded.

**Theorem 3.** *Suppose Assumption 1 holds, and  $\max(\|y_k\|^2, \mathbb{E}[\|\ell(x_{k+1}; \xi)\|^2]) \leq D^2$  for  $k \in \{1, \dots, K\}$ . There exists a constant  $G = 17 \max\{2\sigma^2 + 2D^2\sigma^2, 2\sigma^2 + 2D^2\}$ , and for any constants  $\gamma > 0, \mu > 0, \alpha \geq 1$  Algorithm 2 with  $T_k^y = k/\gamma + 1, T_k^x = k/\mu + 1$  guarantees that the following inequalities hold:*

$$\begin{aligned} \frac{1}{2} \mathbb{E}[\|x_{\tau+1} - v_\tau\|^2] &\leq \mathbb{E}[\|x_\tau - v_\tau\|^2 + \|x_{\tau+1} - x_\tau\|^2] \\ &\leq \frac{4(M + 2G^2)(\alpha + 1)}{\gamma K} \\ \frac{1}{2} \mathbb{E}[\|y_{\tau+1} - u_\tau\|^2] &\leq \mathbb{E}[\|y_\tau - u_\tau\|^2 + \|y_{\tau+1} - y_\tau\|^2] \\ &\leq \frac{4(M + 2G^2)(\alpha + 1)}{\mu K} \end{aligned}$$

for  $\tau$  sampled by  $P(\tau = k) = \frac{k^\alpha}{\sum_{s=1}^K s^\alpha}$ .

The lemma below connects  $\|\nabla F(x_k)\|$  (or  $\text{dist}(0, F(x_k))$ ) to the quantities in Theorem 3, by which we can derive the convergence of (nearly) stationary point.

**Lemma 4.** *Suppose  $g$  is  $L_g$ -smooth, and  $\ell$  is  $G_\ell$ -Lipschitz continuous. Then for any  $k$  we have*

$$\begin{aligned} \|\nabla F(x_k)\| &\leq (\gamma + L_g)\|x_k - v_k\| + G_\ell\|y_k - u_k\| \\ &\quad + G_\ell \mu^v \left(\frac{1+v}{2v}\right)^v L_{h^*}\|u_k - y_k\|^v \\ &\quad + G_\ell^{v+1} \left(\frac{1+v}{2v}\right)^v L_{h^*}\|x_{k+1} - x_k\|^v. \end{aligned}$$

Suppose  $g$  is non-smooth, and  $\ell$  is  $G_\ell$ -Lipschitz continuous and  $L_\ell$ -smooth and  $\max_{y \in \text{dom}(h)} \|y\| \leq D$ , then for any  $k$  we have

$$\begin{aligned} \text{dist}(0, \partial F(v_k)) &\leq (\gamma + DL_\ell)\|x_k - v_k\| + G_\ell\|y_k - u_k\| \\ &\quad + G_\ell \left(\frac{1+v}{2v}\right)^v L_{h^*} \left(\mu\|y_k - u_k\| + G_\ell\|x_{k+1} - z_k\|\right)^v. \end{aligned}$$

Combining Lemma 4 and Theorem 3, we have the following corollaries regarding the convergence of St-SPG under different conditions of  $g$  and  $\ell$ .

**Corollary 4.** *Suppose  $g$  is  $L_g$ -smooth and  $\ell$  is  $G_\ell$ -Lipschitz continuous and both are convex. Under the same conditions as in Theorem 3, we have  $\mathbb{E}[\text{dist}(0, \nabla F(x_\tau))] \leq \epsilon$  after  $K = O(\epsilon^{-\frac{2}{v}})$  stages. Therefore, the total iteration complexity is  $\sum_{k=1}^K (T_k^x + T_k^y) = O(\epsilon^{-\frac{4}{v}})$ .*

**Corollary 5.** *Suppose  $g$  is non-smooth and convex,  $\ell$  is  $G_\ell$ -Lipschitz continuous and  $L_\ell$ -smooth and convex, and*

$\max_{y \in \text{dom}(h)} \|y\| \leq D$ . Under the same conditions as in Theorem 3, we have  $\mathbb{E}[\text{dist}(0, \nabla F(v_\tau))] \leq \epsilon$  and  $\mathbb{E}[\|x_\tau - v_\tau\|] \leq O(\epsilon^{1/v})$  after  $K = O(\epsilon^{-\frac{2}{v}})$  stages. Therefore, the total iteration complexity is  $\sum_{k=1}^K (T_k^x + T_k^y) = O(\epsilon^{-\frac{4}{v}})$ .

**Remark:** Our algorithms enjoy the same iteration complexity of that in (Xu et al., 2018a) for DC functions when  $v$  is unknown or  $v = 1$ , but we do not assume a stochastic gradient of  $h^*(\ell(x))$  is easily computed. It is also notable that St-SPG does not need the knowledge of  $v$  to run.

Finally, we would like to mention that the SPG algorithm for solving subproblems in Algorithm 2 can be replaced by other suitable stochastic optimization algorithms for solving a strongly convex problem similar to the developments in (Xu et al., 2018a) for minimizing DC functions. For example, one can use adaptive stochastic gradient methods in order to enjoy an adaptive convergence, and one can use variance reduction methods if the involved functions are smooth and have a finite-sum structure to achieve an improved convergence.

## 5. Application for Variance Regularization

Table 2. Data statistics.

Datasets	#Examples	#Features	#pos:#neg
a9a	32,561	123	0.3172:1
covtype	581,012	54	1.0509:1
RCV1	697,641	47,236	1.1033:1
URL	2,396,130	3,231,961	0.4939:1

In this section, we consider the application of the proposed algorithms for variance-based regularization in machine learning. Let  $l(\theta, \mathbf{z}) \in \mathbb{R}^+$  denote a loss of model  $\theta \in \Theta$  on a random data  $\mathbf{z}$ . A fundamental task in machine learning is to minimize the expected risk  $R(\theta) = \mathbb{E}_{\mathbf{z}}[l(\theta, \mathbf{z})]$ . However, in practice one has to find an approximate model based on sampled data  $\mathcal{S}_n = \{\mathbf{z}_1, \dots, \mathbf{z}_n\}$ . An advanced learning theory according to Bennett's inequality bounds the expected risk by (Maurer & Pontil, 2009):

$$R(\theta) \leq \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + c_1 \sqrt{\frac{\text{Var}(\ell(\theta, \mathbf{z}))}{n}} + \frac{c_2}{n}, \quad (7)$$

where  $c_1$  and  $c_2$  are constants. This motivates the variance-based regularization approach (Maurer & Pontil, 2009):

$$\min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + \lambda \sqrt{\frac{\text{Var}_n(\theta, \mathcal{S}_n)}{n}}, \quad (8)$$

where  $\text{Var}_n(\theta, \mathcal{S}_n) = \frac{1}{n} \sum_{i=1}^n (\ell(\theta, \mathbf{z}_i) - \bar{l}_n(\theta))^2$  is the empirical variance of loss,  $\bar{l}_n(\theta)$  is the average of empirical loss, and  $\lambda > 0$  is a regularization parameter.

However, the above formulation does not favor efficient stochastic algorithms. To tackle the optimization problem

for variance-based regularization, (Namkoong & Duchi, 2017) proposed a min-max formulation based on *distributionally robust optimization*, given below and proposed stochastic algorithms for solving the resulting min-max formulation when the loss function is convex (Namkoong & Duchi, 2016),

$$\min_{\theta \in \Theta} \max_{P \in \Delta_n} \left\{ \sum_{i=1}^n P_i \ell(\theta, X_i) : D_\phi(P || \hat{P}_n) \leq \rho \right\}, \quad (9)$$

where  $\rho > 0$  is a hyper-parameter,  $\Delta_n = \{P \in \mathbb{R}^n; P \geq 0, \sum_{i=1}^n P_i = 1\}$ ,  $\hat{P}_n = (1/n, \dots, 1/n)$ , and  $D_\phi(P || Q) = \int \phi(\frac{dP}{dQ}) dQ$  is called the  $\phi$ -divergence based on  $\phi(t) = \frac{1}{2}(t-1)^2$ . The min-max formulation is convex and concave when the loss function is convex. Nevertheless, the stochastic optimization algorithms proposed for solving the min-max formulation are not scalable. The reason is that it introduces an  $n$ -dimensional dual variable  $P$  that is restricted on a probability simplex. As a result, the per-iteration cost could be dominated by updating the dual variable that scales as  $O(n)$ , which is prohibitive when the training set is large. Although one can use a special structure and a stochastic coordinate update on  $P$  to reduce the per-iteration cost to  $O(\log(n))$  (Namkoong & Duchi, 2016), the iteration complexity could be still blown up by a factor up to  $n$  due to the variance in the stochastic gradient on  $P$ .

As a potential solution to addressing the scalability issue, we consider the following reformulation:

$$\begin{aligned} F(\theta) &= \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + \lambda \sqrt{\frac{\text{Var}_n(\theta, \mathcal{S}_n)}{n}} \\ &= \min_{\alpha > 0} \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + \lambda \left( \frac{\text{Var}_n(\theta, \mathcal{S}_n)}{2\alpha} + \frac{\alpha}{2n} \right) \\ &= \min_{\alpha > 0} \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + \lambda \left( \frac{\alpha}{2n} \right. \\ &\quad \left. + \frac{\frac{1}{n} \sum_{i=1}^n (\ell(\theta, \mathbf{z}_i))^2 - (\mathbb{E}_i[\ell(\theta, \mathbf{z}_i)])^2}{2\alpha} \right). \end{aligned} \quad (10)$$

In practice, one usually needs to tune the regularization parameter  $\lambda$  in order to achieve the best performance. As a result, we can further simplify the problem by absorbing  $\alpha$  into the regularization parameter  $\lambda$  and end up with the following formulation by noting  $-\frac{1}{2}s^2 = \max_{y \geq 0} \frac{1}{2}y^2 - ys$  for  $s \geq 0$ :

$$\begin{aligned} \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) + \lambda \frac{1}{2n} \sum_{i=1}^n (l(\theta, \mathbf{z}_i))^2 \\ + \lambda \left\{ \min_{y \geq 0} \frac{1}{2}y^2 - y \frac{1}{n} \sum_{i=1}^n l(\theta, \mathbf{z}_i) \right\}. \end{aligned} \quad (11)$$

It is notable that the above formulation only introduces one additional scalable variable  $y \in \mathbb{R}^+$ , though the problem

might become a non-convex problem of  $\theta$ . However, when the loss function  $l(\theta, \mathbf{z})$  itself is a non-convex function, the min-max formulation (9) also loses its convexity, which makes our inf-projection formulation more favorable.

We conduct experiments to verify the efficacy of the inf-projection formulation and proposed stochastic algorithms in comparison to the stochastic algorithms for solving min-max formulation (9). We perform two experiments on four datasets, i.e., a9a, RCV1, covtype and URL from the libsvm website, whose number of examples are  $n = 32561, 581012, 697641$  and  $2396130$ , respectively (Table 2). For each dataset, we randomly sample 80% as training data and the rest as testing data. We evaluate training error and testing error of our algorithms and baselines versus cpu time.

In the first experiment, we use (convex) logistic loss for  $l(\theta, \mathbf{z}_i)$  in our inf-projection formulation (11) and min-max formulation (9). We compare our St-SPG with the stochastic algorithm Bandit Mirror Descend (BMD) proposed in (Namkoong & Duchi, 2016). We implement two versions of BMD, one using the standard mirror descent method to update the dual variable  $P$  and the other (denoted by BMD-eff) exploiting binary search tree (BST) to update the  $P$ . To this end, it needs to use a modified constraint on  $P$ , i.e.,  $P \in \{p \in \mathbb{R}_+^n | p_i \geq \delta/n, n^2/2 \|p - \mathbf{1}/n\|^2 \leq \rho\}$  (see Sec. 4 in (Namkoong & Duchi, 2016)). We tune hyper-parameters from a reasonable range, i.e., for St-SPG,  $\lambda \in \{10^{-5:2}\}$ ,  $\gamma, \mu \in \{10^{-3:3}\}$ . For BMD and BMD-eff, we tune step size  $\eta_P \in \{10^{-8:-15}\}$  for updating  $P$ , step size  $\eta_\theta \in \{10^{-5:3}\}$  for updating  $\theta$ ,  $\rho \in \{n \times 10^{-3:3}\}$  and fix  $\delta = 10^{-5}$ . Training and testing errors against cpu time (s) of the three algorithms on four datasets are reported in Figure 1.

In the second experiment, we use (non-convex) truncated logistic loss in (11) and (9). In particular, the truncated loss function is given by  $\phi(l(\theta, \mathbf{z}_i)) = \alpha \log(1 + l(\theta, \mathbf{z}_i)/\alpha)$ , where  $l$  is logistic loss and we set  $\alpha = \sqrt{10n}$  as suggested in (Xu et al., 2018b). Since the loss is non-convex, we compare MSPG with proximally guided stochastic mirror descent (PGSMD) (Rafique et al., 2018) and its efficient variant (denoted by PGSMD-eff) for solving the min-max formulation that is non-convex and concave, where the efficient variant is implemented with the same modified constraint on  $P$  and BST as BMD-eff. For MSPG, we tune  $\lambda \in \{10^{-5:2}\}$ , the step size parameter  $c$  in Proposition 1 from  $\{10^{-5:2}\}$ . Hyper-parameters of PGSMD and PGSMD-eff including  $\eta_P, \eta_\theta, \rho$  and  $\delta$  are selected in the same range as in the first experiment. The weak convexity parameter  $\rho_{wc}$  are chosen from  $\{10^{-5:5}\}$ . Training and testing errors against cpu time (s) of the three algorithms on four datasets are reported in Figure 2.

We can observe two conclusions from the results of both experiments. First, the training and testing errors from solving the inf-projection formulation (11) converge to a

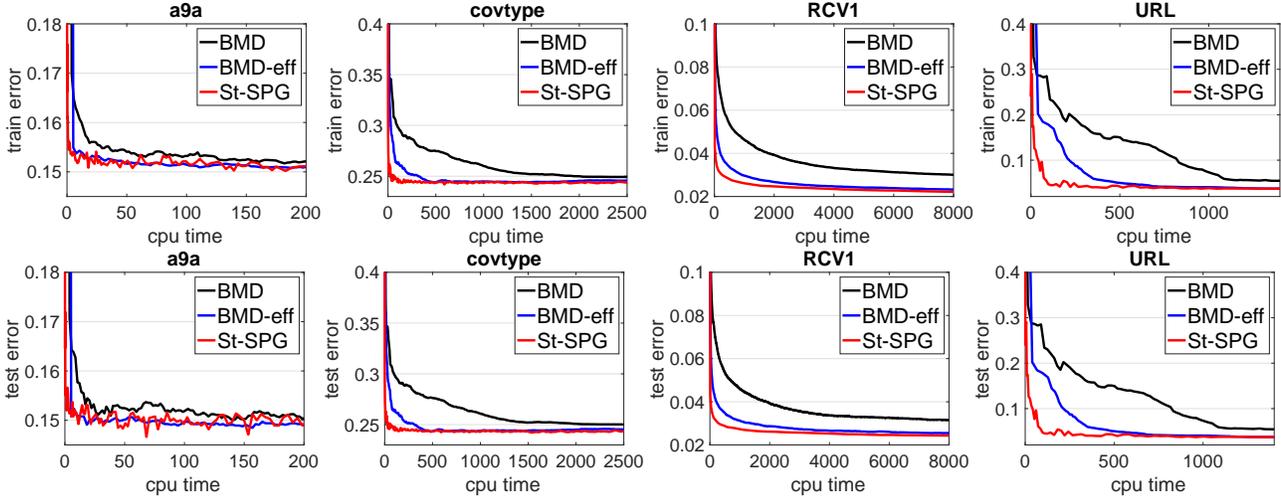


Figure 1. Results of variance-based regularization with (convex) logistic loss.

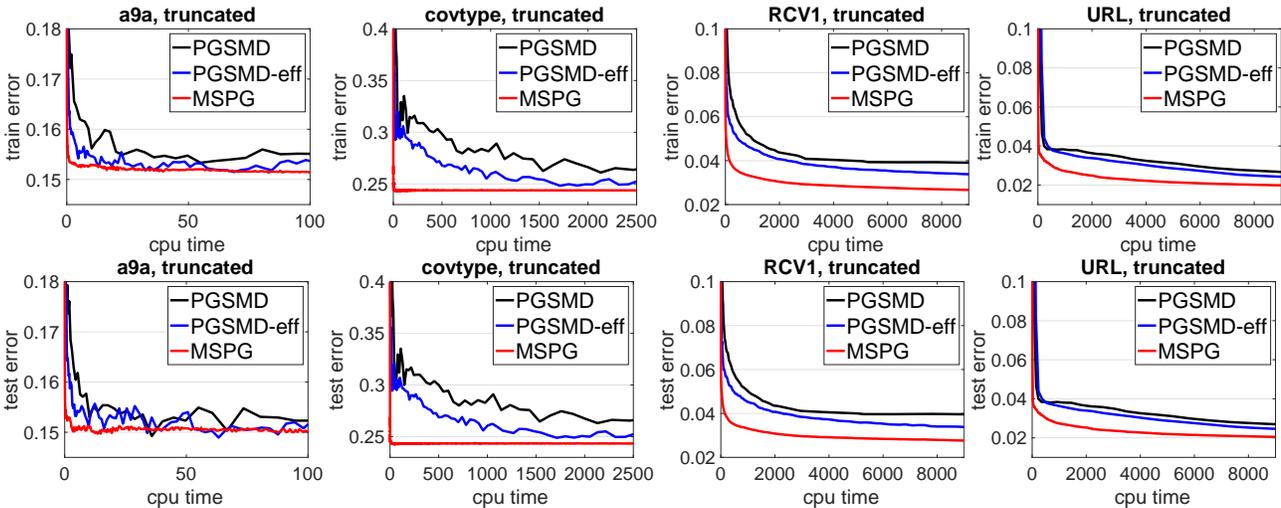


Figure 2. Results of variance-based regularization with for (non-convex) truncated logistic loss.

close or even a lower level compared to that from solving the min-max formulation (9), which verifies the efficacy of the inf-projection formulation. Second, the proposed stochastic algorithms have significant improvement in the convergence time of training/testing errors, especially on large datasets, covtype, RCV1 and URL, which can be verified by comparing convergence of training/testing errors against cpu time.

## 6. Conclusion

In this paper, we design and analyze stochastic optimization algorithms for a family of inf-projection minimization problems. We show that the concerned inf-projection structure covers a variety of special cases, including DC functions and bi-convex functions as special cases (non-smooth functions

in Section 4) and another family of inf-projection formulations (smooth functions in Section 3). We develop stochastic optimization algorithms for those problems with theoretical guarantees of their first-order convergence for finding a (nearly)  $\epsilon$ -stationary solution at  $O(1/\epsilon^{4/\nu})$ . To the best of our knowledge, this is the first work to provide comprehensive convergence analysis for stochastic optimization of non-convex inf-projection minimization problems. Additionally, to verify the significance of our inf-projection formulation, we investigate an important machine learning problem, variance-based regularization, and compare our algorithms with baselines for min-max formulation (distributionally robust optimization). Empirical results demonstrate the significance and effectiveness of our proposed algorithms.

## References

- Bolte, J., Sabach, S., and Teboulle, M. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Math. Program.*, 146(1-2):459–494, August 2014. ISSN 0025-5610.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.
- Chen, Z., Yang, T., Yi, J., Zhou, B., and Chen, E. Universal stagewise learning for non-convex problems with convergence on averaged solutions. *CoRR*, /abs/1808.06296, 2018.
- Davis, D. and Drusvyatskiy, D. Stochastic model-based minimization of weakly convex functions. *CoRR*, abs/1803.06523, 2018a.
- Davis, D. and Drusvyatskiy, D. Stochastic subgradient method converges at the rate  $o(k^{-1/4})$  on weakly convex functions. *CoRR*, /abs/1802.02988, 2018b.
- Davis, D. and Grimmer, B. Proximally guided stochastic subgradient method for nonsmooth, nonconvex problems. *arXiv preprint arXiv:1707.03505*, 2017.
- Davis, D., Edmunds, B., and Udell, M. The sound of apalm clapping: Faster nonsmooth nonconvex optimization with stochastic asynchronous palm. In *Advances in Neural Information Processing Systems*, pp. 226–234, 2016.
- Driggs, D., Tang, J., Davies, M., and Schönlieb, C.-B. Spring: A fast stochastic proximal alternating method for non-smooth non-convex optimization. *arXiv preprint arXiv:2002.12266*, 2020.
- Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Math. Program.*, 155(1-2):267–305, 2016.
- Gorski, J., Pfeuffer, F., and Klamroth, K. Biconvex sets and optimization with biconvex functions: a survey and extensions. *Mathematical Methods of Operations Research*, 66(3):373–407, Dec 2007.
- Hong, M., Razaviyayn, M., Luo, Z.-Q., and Pang, J.-S. A unified algorithmic framework for block-structured optimization involving big data: With applications in machine learning and signal processing. *IEEE Signal Processing Magazine*, 33(1):57–77, 2015.
- Kiryo, R., Niu, G., du Plessis, M. C., and Sugiyama, M. Positive-unlabeled learning with non-negative risk estimator. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 1675–1685. Curran Associates, Inc., 2017.
- Kumar, M. P., Packer, B., and Koller, D. Self-paced learning for latent variable models. In *Neural Information Processing Systems 23*, pp. 1189–1197, 2010.
- Maurer, A. and Pontil, M. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- Namkoong, H. and Duchi, J. C. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2208–2216, 2016.
- Namkoong, H. and Duchi, J. C. Variance-based regularization with convex objectives. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2975–2984, 2017.
- Nesterov, Y. Universal gradient methods for convex optimization problems. *Mathematical Programming*, 152(1): 381–404, Aug 2015. ISSN 1436-4646. doi: 10.1007/s10107-014-0790-0. URL <https://doi.org/10.1007/s10107-014-0790-0>.
- Nitanda, A. and Suzuki, T. Stochastic difference of convex algorithm and its application to training deep boltzmann machines. In *Artificial Intelligence and Statistics*, pp. 470–478, 2017.
- Rafique, H., Liu, M., Lin, Q., and Yang, T. Non-convex min-max optimization: Provable algorithms and applications in machine learning. *CoRR*, abs/1810.02060, 2018.
- Rockafellar, R. and Wets, R. J.-B. *Variational Analysis*. Springer Verlag, Heidelberg, Berlin, New York, 1998.
- Rockafellar, R. T. and Wets, R. J.-B. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- Shah, S., Yadav, A. K., Castillo, C. D., Jacobs, D. W., Studer, C., and Goldstein, T. Biconvex relaxation for semidefinite programming in computer vision. In *ECCV (6)*, volume 9910 of *Lecture Notes in Computer Science*, pp. 717–735. Springer, 2016.
- Shalev-Shwartz, S. and Singer, Y. On the equivalence of weak learnability and linear separability: New relaxations and efficient boosting algorithms. *Machine learning*, 80 (2-3):141–163, 2010.
- Thi, H. A. L., Le, H. M., Phan, D. N., and Tran, B. Stochastic DCA for the large-sum of non-convex functions problem and its application to group variable selection in classification. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3394–3403, International Convention Centre, Sydney, Australia, 06–11 Aug

2017. PMLR. URL <http://proceedings.mlr.press/v70/thi17a.html>.

Xu, Y. and Yin, W. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013.

Xu, Y., Qi, Q., Lin, Q., Jin, R., and Yang, T. Stochastic optimization for dc functions and non-smooth non-convex regularizers with non-asymptotic convergence. *arXiv preprint arXiv:1811.11829*, 2018a.

Xu, Y., Zhu, S., Yang, S., Zhang, C., Jin, R., and Yang, T. Learning with non-convex truncated losses by SGD. *CoRR*, abs/1805.07880, 2018b. URL <http://arxiv.org/abs/1805.07880>.

Xu, Y., Jin, R., and Yang, T. Stochastic proximal gradient methods for non-smooth non-convex regularized problems. *CoRR*, abs/1902.07672, 2019.

Zhao, P. and Zhang, T. Stochastic optimization with importance sampling for regularized loss minimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 1–9, 2015.