

# SECOND-ORDER MIN-MAX OPTIMIZATION WITH LAZY HESSIANS

**Lesi Chen** \*

IIS, Tsinghua University & Shanghai Qizhi Institute  
chenlc23@mails.tsinghua.edu.cn

**Chengchang Liu** \*

The Chinese University of Hong Kong  
7liuchengchang@gmail.com

**Jingzhao Zhang** †

IIS, Tsinghua University & Shanghai Qizhi Institute & Shanghai AI Lab  
jingzhaoz@mail.tsinghua.edu.cn

## ABSTRACT

This paper studies second-order methods for convex-concave minimax optimization. Monteiro & Svaiter (2012) proposed a method to solve the problem with an optimal iteration complexity of  $\mathcal{O}(\epsilon^{-3/2})$  to find an  $\epsilon$ -saddle point. However, it is unclear whether the computational complexity,  $\mathcal{O}((N + d^2)d\epsilon^{-2/3})$ , can be improved. In the above, we follow Doikov et al. (2023) and assume the complexity of obtaining a first-order oracle as  $N$  and the complexity of obtaining a second-order oracle as  $dN$ . In this paper, we show that the computation cost can be reduced by reusing Hessian across iterations. Our methods take the overall computational complexity of  $\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3}\epsilon^{-2/3}))$ , which improves those of the previous methods by a factor of  $d^{1/3}$ . Furthermore, we generalize our method to strongly-convex-strongly-concave minimax problems and establish the complexity of  $\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3}\kappa^{2/3}))$  when the condition number of the problem is  $\kappa$ , enjoying a similar speedup upon the state-of-the-art method. Numerical experiments on both real and synthetic datasets also verify the efficiency of our method.

## 1 INTRODUCTION

We consider the following minimax optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^{d_x}} \max_{\mathbf{y} \in \mathbb{R}^{d_y}} f(\mathbf{x}, \mathbf{y}), \quad (1)$$

where we suppose  $f(\mathbf{x}, \mathbf{y})$  is (strongly-)convex in  $\mathbf{x}$  and (strongly-)concave in  $\mathbf{y}$ . This setting covers many useful applications, including functionally constrained optimization (Xu, 2020), game theory (Von Neumann & Morgenstern, 1947), robust optimization (Ben-Tal et al., 2009), fairness-aware machine learning (Zhang et al., 2018), reinforcement learning (Du et al., 2017; Wang, 2017; Paternain et al., 2022; Wai et al., 2018), decentralized optimization (Kovalev et al., 2021; 2020), AUC maximization (Ying et al., 2016; Hanley & McNeil, 1982; Yuan et al., 2021).

First-order methods are widely studied for this problem. Classical algorithms include ExtraGradient (EG) (Korpelevich, 1976; Nemirovski, 2004), Optimistic Gradient Descent Ascent (OGDA) (Popov, 1980; Mokhtari et al., 2020a;b), Hybrid Proximal Extragradient (HPE) (Monteiro & Svaiter, 2010), and Dual Extrapolation (DE) (Nesterov & Scramali, 2006; Nesterov, 2007). When the gradient of  $f(\cdot, \cdot)$  is  $L$ -Lipschitz continuous, these methods achieve the rate of  $\mathcal{O}(\epsilon^{-1})$  under the convex-concave (C-C) setting and the rate of  $\mathcal{O}((L/\mu) \log(\epsilon^{-1}))$  when  $f(\cdot, \cdot)$  is  $\mu$ -strongly convex in  $\mathbf{x}$  and  $\mu$ -strongly-concave in  $\mathbf{y}$  (SC-SC) for  $\mu > 0$ . They are all optimal in C-C and SC-SC settings due to the lower bounds reported by (Nemirovskij & Yudin, 1983; Zhang et al., 2022a).

\* Equal contributions.

† The corresponding author.

Second-order methods usually lead to faster rates than first-order methods when the Hessian of  $f(\cdot, \cdot)$  is  $\rho$ -Lipschitz continuous. A line of works (Nesterov & Scramali, 2006; Huang et al., 2022) extended the celebrated Cubic Regularized Newton (CRN) (Nesterov & Polyak, 2006) method to minimax problems with local superlinear convergence rates and global convergence guarantee. However, the established global convergence rates of  $\mathcal{O}(\epsilon^{-1})$  by Nesterov & Scramali (2006) and  $\mathcal{O}((L\rho/\mu^2)\log(\epsilon^{-1}))$  by Huang et al. (2022) under C-C and SC-SC conditions are no better than the optimal first-order methods. Another line of work generalizes the optimal first-order methods to higher-order methods. Monteiro & Svaiter (2012) proposed the Newton Proximal Extragradient (NPE) method with a global convergence rate of  $\mathcal{O}(\epsilon^{-2/3}\log\log(\epsilon^{-1}))$  under the C-C conditions. This result nearly matches the lower bounds (Adil et al., 2022; Lin & Jordan, 2024), except an additional  $\mathcal{O}(\log\log(\epsilon^{-1}))$  factor which is caused by the implicit binary search at each iteration. Bullins & Lai (2022); Adil et al. (2022); Huang & Zhang (2022); Lin et al. (2022) provided a simple proof of NPE motivated by the EG analysis and showed that replacing the quadratic regularized Newton step with the cubic regularized Newton (CRN) step in NPE achieves the optimal second-order oracle complexity of  $\mathcal{O}(\epsilon^{-2/3})$ . Recently, Alves & Svaiter (2023) proposed a search-free NPE method to achieve the optimal second-order oracle complexity with pure quadratic regularized Newton step based on ideas from homotopy. Over the past decade, researchers also proposed various second-order methods, in addition to the NPE framework, that achieve the same convergence rate, such as the second-order extensions of OGD (Jiang & Mokhtari, 2022; Jiang et al., 2024) (which we refer to as OGD-2) and DE (Lin & Jordan, 2024) (they name the method Persesus). The results for C-C problems can also be extended to SC-SC problems, where Jiang & Mokhtari (2022) proved the OGD-2 can converge at the rate of  $\mathcal{O}((\rho/\mu)^{2/3} + \log\log(\epsilon^{-1}))$ , and Huang & Zhang (2022) proposed the ARE-restart with the rate of  $\mathcal{O}((\rho/\mu)^{2/3}\log\log(\epsilon^{-1}))$ .

Although the aforementioned second-order methods Adil et al. (2022); Lin & Jordan (2024); Lin et al. (2022); Jiang & Mokhtari (2022); Monteiro & Svaiter (2012) enjoy an improved convergence rate over the first-order methods and have achieved optimal iteration complexities, they require querying one new Hessian at each iteration and solving a matrix inversion problem at each Newton step, which leads to a  $\mathcal{O}(d^3)$  computational cost per iteration. This becomes the main bottleneck that limits the applicability of second-order methods. Liu & Luo (2022a) proposed quasi-Newton methods for saddle point problems that access one Hessian-vector product instead of the exact Hessian for each iteration. The iteration complexity is  $\mathcal{O}(d^2)$  for quasi-Newton methods. However, their methods do not have a global convergence guarantee under general (S)C-(S)C conditions. Jiang et al. (2023) proposed an online-learning guided Quasi-Newton Proximal Extragradient (QNPE) algorithm, but their method relies on more complicated subroutines than classical Newton methods. Although the oracle complexity of QNPE is strictly better than the optimal first-order method EG, their method is worse in terms of total computational complexity.

In this paper, we propose a computation-efficient second-order method, which we call LEN (Lazy Extra Newton method). In contrast to all existing second-order methods or quasi-Newton methods for minimax optimization problems that always access new second-order information for the coming iteration, LEN reuses the second-order information from past iterations. Specifically, LEN solves a cubic regularized sub-problem using the Hessian from the snapshot point that is updated every  $m$  iteration, then conducts an extra-gradient step by the gradient from the current iteration. We provide a rigorous theoretical analysis of LEN to show it maintains fast global convergence rates and improves the (near)-optimal second-order methods (Monteiro & Svaiter, 2012) in terms of the overall computational complexity. We summarize our contributions as follows (also see Table 1).

- When the object function  $f(\cdot, \cdot)$  is convex in  $\mathbf{x}$  and concave in  $\mathbf{y}$ , we propose LEN and prove that it finds an  $\epsilon$ -saddle point in  $\mathcal{O}(m^{2/3}\epsilon^{-2/3})$  iterations. Under Assumption 3.4, where the complexity of calculating  $\mathbf{F}(\mathbf{z})$  is  $N$  and the complexity of calculating  $\nabla\mathbf{F}(\mathbf{z})$  is  $dN$ , the optimal choice is  $m = \Theta(d)$ . In this case, LEN only requires a computational complexity of  $\tilde{\mathcal{O}}((N+d^2)(d+d^{2/3}\epsilon^{-2/3}))$ , which is strictly better than  $\mathcal{O}((N+d^2)d\epsilon^{-2/3})$  for the existing optimal second-order methods by a factor of  $d^{1/3}$ .
- When the object function  $f(\cdot, \cdot)$  is  $\mu$ -strongly-convex in  $\mathbf{x}$  and  $\mu$ -strongly-concave in  $\mathbf{y}$ , we apply the restart strategy on LEN and propose LEN-restart. We prove the algorithm can find an  $\epsilon$ -root with  $\tilde{\mathcal{O}}((N+d^2)(d+d^{2/3}(\rho/\mu)^{2/3}))$  computational complexity, where  $\rho$  means the Hessian of  $f(\cdot, \cdot)$  is  $\rho$  Lipschitz-continuous. Our result is strictly better than the  $\tilde{\mathcal{O}}((N+d^2)d(\rho/\mu)^{2/3})$  in prior works.

Table 1: We compare the required computational complexity to achieve an  $\epsilon$ -saddle point of the proposed LEN with the optimal choice  $m = \Theta(d)$  and other existing algorithms on both convex-concave (C-C) and strongly-convex-strongly-concave (SC-SC) problems. Here,  $d = d_x + d_y$  is the dimension of the problem. We assume the gradient is  $L$ -Lipschitz continuous for EG and the Hessian is  $\rho$ -Lipschitz continuous for others. We count each gradient oracle call with  $N$  computational complexity, and each Hessian oracle with  $dN$  computational complexity.

Setup	Method	Computational Cost
	EG (Korpelevich, 1976)	$\mathcal{O}((N + d)\epsilon^{-1})$
	NPE (Monteiro & Svaiter, 2012)	$\tilde{\mathcal{O}}((N + d^2)d\epsilon^{-2/3})$
C-C	search-free NPE (Alves & Svaiter, 2023)	$\mathcal{O}((N + d^2)d\epsilon^{-2/3})$
	OGDA-2 (Jiang & Mokhtari, 2022)	$\mathcal{O}((N + d^2)d\epsilon^{-2/3})$
	LEN (Theorem 4.3)	$\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3}\epsilon^{-2/3}))$
	EG (Korpelevich, 1976)	$\tilde{\mathcal{O}}((N + d)(L/\mu))$
	OGDA-2 (Jiang & Mokhtari, 2022)	$\mathcal{O}((N + d^2)d(\rho/\mu)^{2/3})$
SC-SC	ARE-restart (Huang & Zhang, 2022)	$\tilde{\mathcal{O}}((N + d^2)d(\rho/\mu)^{2/3})$
	Perseus-restart (Lin & Jordan, 2024)	$\tilde{\mathcal{O}}((N + d^2)d(\rho/\mu)^{2/3})$
	LEN-restart (Corollary 4.1)	$\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3}(\rho/\mu)^{2/3}))$

**Notations.** Throughout this paper,  $\log$  is base 2 and  $\log_+(\cdot) := 1 + \log(\cdot)$ . We use  $\|\cdot\|$  to denote the spectral norm and the Euclidean norm of matrices and vectors, respectively. We denote  $\pi(t) = t - (t \bmod m)$  where  $m \in \mathbb{N}_+$ .

## 2 RELATED WORKS AND TECHNICAL CHALLENGES

**Lazy Hessian in minimization problems.** The idea of reusing Hessian was initially presented by Shamanskii (1967) and later incorporated into the Levenberg-Marquardt method, the Damped Newton method, and the proximal Newton method (Fan, 2013; Lampariello & Sciandrone, 2001; Wang et al., 2006; Adler et al., 2020). However, the explicit advantage of lazy Hessian update over ordinary Newton (-type) update was not discovered until the recent work of (Doikov et al., 2023; Chayti et al., 2023). They applied the following lazy Hessian update on cubic regularized Newton (CRN) methods (Nesterov & Polyak, 2006):

$$z_{t+1} = \arg \min_{z \in \mathbb{R}^d} \left\{ \langle \mathbf{F}(z_t), z - z_t \rangle + \frac{1}{2} \langle \nabla \mathbf{F}(z_{\pi(t)}) (z - z_t), z - z_t \rangle + \frac{M}{6} \|z - z_t\|^3 \right\}, \quad (2)$$

where  $M \geq 0$  and  $\mathbf{F} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is the gradient field of a convex function. They establish the convergence rates of  $\mathcal{O}(\sqrt{m}\epsilon^{-3/2})$  for nonconvex optimization (Doikov et al., 2023), and  $\mathcal{O}(\sqrt{m}\epsilon^{-1/2})$  for convex optimization (Chayti et al., 2023) respectively. Such rates lead to the total computational cost of  $\tilde{\mathcal{O}}((N + d^2)(d + \sqrt{d}\epsilon^{-3/2}))$  and  $\tilde{\mathcal{O}}((N + d^2)(d + \sqrt{d}\epsilon^{-1/2}))$  by setting  $m = \Theta(d)$ , which strictly improve the result by classical CRN methods by a factor of  $\sqrt{d}$  in both setups.

We have also observed that the idea of the ‘‘lazy Hessian’’ is widely used in practical second-order methods. KFAC (Martens & Grosse, 2015; Grosse & Martens, 2016) approximates the Fisher information matrix and uses an exponential moving average (EMA) to update the estimate of the Fisher information matrix, which can be viewed as a soft version of lazy update. Sophia (Liu et al., 2024) estimates a diagonal Hessian matrix as a pre-conditioner, which is updated in a lazy manner to reduce the complexity. C2EDEN (Liu et al., 2023) atomizes the communication of local Hessian in several consecutive iterations, which also benefits from the idea of lazy updates.

**Challenge of using lazy Hessian updates in minimax problems.** In comparison to previous work on lazy Hessian, our LEN and LEN-restart methods demonstrate the advantage of using lazy Hessian

for a broader class of optimization problems, the *minimax* problems. Our analysis differs from the ones in Doikov et al. (2023); Chayti et al. (2023). Their methods only take a lazy CRN update (2) at each iteration, which makes it easy to bound the error of lazy Hessian updates using Assumption 3.1 and the triangle inequality in the following way:

$$\|\nabla \mathbf{F}(\mathbf{z}_t) - \nabla \mathbf{F}(\mathbf{z}_{\pi(t)})\| \leq \rho \|\mathbf{z}_{\pi(t)} - \mathbf{z}_t\| \leq \rho \sum_{i=\pi(t)}^{t-1} \|\mathbf{z}_i - \mathbf{z}_{i+1}\|.$$

Our method, on the other hand, not only takes a lazy (regularized) Newton update but also requires an extra gradient step (Line 4 in Algorithm 1). Thus, the progress of one Newton update  $\{\|\mathbf{z}_{i+1/2} - \mathbf{z}_i\|\}_{i=\pi(t)}^t$  cannot directly bound the error term  $\|\mathbf{z}_t - \mathbf{z}_{\pi(t)}\|$  introduced by the lazy Hessian update. Moreover, for minimax problems the matrix  $\nabla \mathbf{F}(\mathbf{z}_{\pi(t)})$  is no longer symmetric, which leads to different analysis and implementation of sub-problem solving (Section 4.3). We refer the readers to Section 4 for more detailed discussions.

### 3 PRELIMINARIES

In this section, we introduce the notation and basic assumptions used in our work. We start with several standard definitions for Problem (1).

**Definition 3.1.** We call a function  $f(\mathbf{x}, \mathbf{y}) : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}$  has  $\rho$ -Lipschitz Hessians if we have

$$\|\nabla^2 f(\mathbf{x}, \mathbf{y}) - \nabla^2 f(\mathbf{x}', \mathbf{y}')\| \leq \rho \left\| \begin{bmatrix} \mathbf{x} - \mathbf{x}' \\ \mathbf{y} - \mathbf{y}' \end{bmatrix} \right\|, \quad \forall (\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}') \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y}.$$

**Definition 3.2.** A differentiable function  $f(\cdot, \cdot)$  is  $\mu$ -strongly-convex- $\mu$ -strongly-concave for some  $\mu > 0$  if

$$\begin{aligned} f(\mathbf{x}', \mathbf{y}) &\geq f(\mathbf{x}, \mathbf{y}) + (\mathbf{x}' - \mathbf{x})^\top \nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) + \frac{\mu}{2} \|\mathbf{x} - \mathbf{x}'\|^2, \quad \forall \mathbf{x}', \mathbf{x} \in \mathbb{R}^{d_x}, \mathbf{y} \in \mathbb{R}^{d_y}; \\ f(\mathbf{x}, \mathbf{y}') &\leq f(\mathbf{x}, \mathbf{y}) + (\mathbf{y}' - \mathbf{y})^\top \nabla_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) - \frac{\mu}{2} \|\mathbf{y} - \mathbf{y}'\|^2, \quad \forall \mathbf{y}', \mathbf{y} \in \mathbb{R}^{d_y}, \mathbf{x} \in \mathbb{R}^{d_x}. \end{aligned}$$

We say  $f$  is convex-concave if  $\mu = 0$ .

We are interested in finding a saddle point of Problem (1), formally defined as follows.

**Definition 3.3.** We call a point  $(\mathbf{x}^*, \mathbf{y}^*) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y}$  a saddle point of a function  $f(\cdot, \cdot)$  if we have

$$f(\mathbf{x}^*, \mathbf{y}) \leq f(\mathbf{x}^*, \mathbf{y}^*) \leq f(\mathbf{x}, \mathbf{y}^*), \quad \forall \mathbf{x} \in \mathbb{R}^{d_x}, \mathbf{y} \in \mathbb{R}^{d_y}.$$

Next, we introduce all the assumptions made in this work. In this paper, we focus on Problem (1) that satisfies the following assumptions.

**Assumption 3.1.** We assume the function  $f(\cdot, \cdot)$  is twice continuously differentiable, has  $\rho$ -Lipschitz continuous Hessians, and has at least one saddle point  $(\mathbf{x}^*, \mathbf{y}^*)$ .

We will study convex-concave problems and strongly-convex-strongly-concave problems.

**Assumption 3.2** (C-C setting). We assume the function  $f(\cdot, \cdot)$  is convex in  $\mathbf{x}$  and concave in  $\mathbf{y}$ .

**Assumption 3.3** (SC-SC setting). We assume the function  $f(\cdot, \cdot)$  is  $\mu$ -strongly-convex- $\mu$ -strongly-concave. We denote the condition number as  $\kappa := \rho/\mu$ .

We let  $d := d_x + d_y$  and denote the aggregated variable  $\mathbf{z} := (\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d$ . We also denote the GDA field of  $f$  and its Jacobian as

$$\mathbf{F}(\mathbf{z}) := \begin{bmatrix} \nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) \\ -\nabla_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) \end{bmatrix}, \quad \nabla \mathbf{F}(\mathbf{z}) := \begin{bmatrix} \nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{x}, \mathbf{y}) & \nabla_{\mathbf{x}\mathbf{y}}^2 f(\mathbf{x}, \mathbf{y}) \\ -\nabla_{\mathbf{y}\mathbf{x}}^2 f(\mathbf{x}, \mathbf{y}) & -\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}, \mathbf{y}) \end{bmatrix}. \quad (3)$$

The GDA field of  $f(\cdot, \cdot)$  has the following properties.

**Lemma 3.1** (Lemma 2.7 Lin et al. (2022)). Under Assumptions 3.1 and 3.2, we have

1.  $\mathbf{F}$  is monotone, i.e.  $\langle \mathbf{F}(\mathbf{z}) - \mathbf{F}(\mathbf{z}'), \mathbf{z} - \mathbf{z}' \rangle \geq 0, \forall \mathbf{z}, \mathbf{z}' \in \mathbb{R}^d$ .
2.  $\nabla \mathbf{F}$  is  $\rho$ -Lipschitz continuous, i.e.  $\|\nabla \mathbf{F}(\mathbf{z}) - \nabla \mathbf{F}(\mathbf{z}')\| \leq \rho \|\mathbf{z} - \mathbf{z}'\|, \forall \mathbf{z}, \mathbf{z}' \in \mathbb{R}^d$ .
3.  $\mathbf{F}(\mathbf{z}^*) = 0$  if and only if  $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$  is a saddle point of function  $f(\cdot, \cdot)$ .

Furthermore, if Assumption 3.3 holds, we have  $\mathbf{F}(\cdot)$  is  $\mu$ -strongly-monotone, i.e.

$$\langle \mathbf{F}(\mathbf{z}) - \mathbf{F}(\mathbf{z}'), \mathbf{z} - \mathbf{z}' \rangle \geq \mu \|\mathbf{z} - \mathbf{z}'\|^2, \forall \mathbf{z}, \mathbf{z}' \in \mathbb{R}^d.$$

For the C-C case, the commonly used optimality criterion is the following restricted gap.

**Definition 3.4** (Nesterov (2007)). Let  $\mathbb{B}_\beta(\mathbf{w})$  be the ball centered at  $\mathbf{w}$  with radius  $\beta$ . Let  $(\mathbf{x}^*, \mathbf{y}^*)$  be a saddle point of function  $f$ . For a given point  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ , we let  $\beta$  sufficiently large such that it holds

$$\max \{\|\hat{\mathbf{x}} - \mathbf{x}^*\|, \|\hat{\mathbf{y}} - \mathbf{y}^*\|\} \leq \beta,$$

we define the restricted gap function as

$$\text{Gap}(\hat{\mathbf{x}}, \hat{\mathbf{y}}; \beta) := \max_{\mathbf{y} \in \mathbb{B}_\beta(\mathbf{y}^*)} f(\hat{\mathbf{x}}, \mathbf{y}) - \min_{\mathbf{x} \in \mathbb{B}_\beta(\mathbf{x}^*)} f(\mathbf{x}, \hat{\mathbf{y}}),$$

We call  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  an  $\epsilon$ -saddle point if  $\text{Gap}(\hat{\mathbf{x}}, \hat{\mathbf{y}}; \beta) \leq \epsilon$  and  $\beta = \Omega(\max\{\|\mathbf{x}_0 - \mathbf{x}^*\|, \|\mathbf{y}_0 - \mathbf{y}^*\|\})$ .

For the SC-SC case, we use the following stronger notion.

**Definition 3.5.** Suppose that Assumption 3.3 holds. Let  $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$  be the unique saddle point of function  $f$ . We call  $\hat{\mathbf{z}} = (\hat{\mathbf{x}}, \hat{\mathbf{y}})$  an  $\epsilon$ -root if  $\|\hat{\mathbf{z}} - \mathbf{z}^*\| \leq \epsilon$ .

Most previous works only consider the complexity metric as the number of oracle calls, where an oracle takes a point  $\mathbf{z} \in \mathbb{R}^d$  as the input and returns a tuple  $(\mathbf{F}(\mathbf{z}), \nabla \mathbf{F}(\mathbf{z}))$  as the output. The existing algorithms (Monteiro & Svaiter, 2012; Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022) have achieved optimal complexity regarding the number of oracle calls. In this work, we focus on the computational complexity of the oracle. More specifically, we distinguish between the different computational complexities of calculating the Hessian matrix  $\nabla^2 \mathbf{F}(\mathbf{z})$  and the gradient  $\mathbf{F}(\mathbf{z})$ . Formally, we make the following assumption as Doikov et al. (2023).

**Assumption 3.4.** We count the computational complexity of computing  $\mathbf{F}(\cdot)$  as  $N$  and the computational complexity of  $\nabla \mathbf{F}(\cdot)$  as  $Nd$ .

**Remark 3.1.** Assumption 3.4 supposes the cost of computing  $\nabla \mathbf{F}(\cdot)$  is  $d$  times that of computing  $\mathbf{F}(\cdot)$ . It holds in many practical scenarios as one Hessian oracle can be computed via  $d$  Hessian-vector products on standard basis vectors  $\mathbf{e}_1, \dots, \mathbf{e}_d$ , and one Hessian-vector product oracle is typically as expensive as one gradient oracle (Wright, 2006):

1. When the computational graph of  $f$  is obtainable, both  $\mathbf{F}(\mathbf{z})$  and  $\nabla \mathbf{F}(\mathbf{z})\mathbf{v}$  can be computed using automatic differentiation with the same cost for any  $\mathbf{z}, \mathbf{v} \in \mathbb{R}^d$ .
2. When  $f$  is a black box function, we can estimate the Hessian-vector  $\nabla^2 \mathbf{F}(\mathbf{z})\mathbf{v}$  via the finite-difference  $\mathbf{u}_\delta(\mathbf{z}; \mathbf{v}) = \frac{1}{\delta}(\mathbf{F}(\mathbf{z} + \delta\mathbf{v}) - \mathbf{F}(\mathbf{z} - \delta\mathbf{v}))$  and we have  $\lim_{\delta \rightarrow 0} \mathbf{u}_\delta(\mathbf{z}; \mathbf{v}) = \nabla^2 \mathbf{F}(\mathbf{z})\mathbf{v}$  under mild conditions on  $\mathbf{F}(\cdot)$ .

## 4 ALGORITHMS AND CONVERGENCE ANALYSIS

In this section, we present novel second-order methods for solving minimax optimization problems (1). We present LEN and its convergence analysis for convex-concave minimax problems in Section 4.1. We generalize LEN for strongly-convex-strongly-concave minimax problems by presenting LEN-restart in Section 4.2. We discuss the details of solving minimax cubic-regularized sub-problem, present detailed implementation of LEN, and give the total computational complexity of proposed methods in Section 4.3.

### 4.1 THE LEN ALGORITHM FOR CONVEX-CONCAVE PROBLEMS

We propose LEN for convex-concave problems in Algorithm 1. Our method builds on the optimal Newton Proximal Extragradient (NPE) method (Monteiro & Svaiter, 2012; Bullins & Lai, 2022;

**Algorithm 1** LEN( $z_0, T, m, M$ )

- 1: **for**  $t = 0, \dots, T - 1$  **do**
- 2:   Compute lazy cubic step, *i.e.* find  $z_{t+1/2}$  that satisfies

$$\mathbf{F}(z_t) = (\nabla \mathbf{F}(z_{\pi(t)}) + M \|z_t - z_{t+1/2}\| \mathbf{I}_d)(z_t - z_{t+1/2}).$$

- 3:   Compute  $\gamma_t = M \|z_t - z_{t+1/2}\|$ .
- 4:   Compute extra-gradient step  $z_{t+1} = z_t - \gamma_t^{-1} \mathbf{F}(z_{t+1/2})$ .
- 5: **end for**
- 6: **return**  $\bar{z}_T = \frac{1}{\sum_{t=0}^{T-1} \gamma_t^{-1}} \sum_{t=0}^{T-1} \gamma_t^{-1} z_{t+1/2}$ .

Adil et al., 2022; Lin et al., 2022). The only change is that we reuse the Hessian from previous iterates, as colored in blue. Each iteration of LEN contains the following two steps:

$$\begin{cases} \mathbf{F}(z_t) + \nabla \mathbf{F}(z_{\pi(t)})(z_{t+1/2} - z_t) + M \|z_{t+1/2} - z_t\| (z_{t+1/2} - z_t) = \mathbf{0}, & \text{(Implicit Step)} \\ z_{t+1} = z_t - \frac{\mathbf{F}(z_{t+1/2})}{M \|z_{t+1/2} - z_t\|}. & \text{(Explicit Step)} \end{cases} \quad (4)$$

The first step (implicit step) solves a cubic regularized sub-problem based on the  $\nabla \mathbf{F}(z_{\pi(t)})$  computed at the latest snapshot point and  $\mathbf{F}(z_t)$  at the current iteration point. This step is often viewed as an oracle (Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022) as there exists efficient solvers, which will also be discussed in Section 4.3. The second one (explicit step) conducts an extra gradient step based on  $\mathbf{F}(z_{t+1/2})$ .

Reusing the Hessian in the implicit step makes each iteration much cheaper, but would cause additional errors compared to previous methods (Monteiro & Svaiter, 2012; Huang & Zhang, 2022; Adil et al., 2022; Lin et al., 2022). The error resulting from the lazy Hessian updates is formally characterized by the following theorem.

**Lemma 4.1.** *Suppose that Assumption 3.1 and 3.2 hold. For any  $z \in \mathbb{R}^d$ , Algorithm 1 ensures*

$$\begin{aligned} \gamma_t^{-1} \langle \mathbf{F}(z_{t+1/2}), z_{t+1/2} - z \rangle &\leq \frac{1}{2} \|z_t - z\|^2 - \frac{1}{2} \|z_{t+1} - z\|^2 - \frac{1}{2} \|z_{t+1/2} - z_{t+1}\|^2 \\ &\quad - \frac{1}{2} \|z_t - z_{t+1/2}\|^2 + \frac{\rho^2}{2M^2} \|z_t - z_{t+1/2}\|^2 + \underbrace{\frac{2\rho^2}{M^2} \|z_{\pi(t)} - z_t\|^2}_{(*)}. \end{aligned}$$

Above, (\*) is the error from lazy Hessian updates. Note that (\*) vanishes when the current Hessian is used. For lazy Hessian updates, the error would accumulate in the epoch.

The key step in our analysis shows that we can use the negative terms in the right-hand side of the inequality in Lemma 4.1 to bound the accumulated error by choosing  $M$  sufficiently large, with the help of the following technical lemma.

**Lemma 4.2.** *For any sequence of positive numbers  $\{r_t\}_{t \geq 0}$ , it holds for any  $m \geq 2$  that*

$$\sum_{t=1}^{m-1} \left( \sum_{i=0}^{t-1} r_i \right)^2 \leq \frac{m^2}{2} \sum_{t=0}^{m-1} r_t^2.$$

When  $m = 1$ , the algorithm reduces to the NPE algorithm (Monteiro & Svaiter, 2012; Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022) without using lazy Hessian updates. When  $m \geq 2$ , we use Lemma 4.2 to upper bound the error that arises from lazy Hessian updates. Finally, we prove the following guarantee for our proposed algorithm.

**Theorem 4.1** (C-C setting). *Suppose that Assumption 3.1 and 3.2 hold. Let  $z^* = (x^*, y^*)$  be a saddle point and  $\beta = \|z_0 - z^*\|$ . Set  $M \geq 3\rho m$ . The sequence of iterates generated by Algorithm 1 is bounded  $z_t \in \mathbb{B}_\beta(z^*)$ ,  $z_{t+1/2} \in \mathbb{B}_{3\beta}(z^*)$ ,  $\forall t = 0, \dots, T - 1$ , and satisfies the following ergodic convergence:*

$$\text{Gap}(\bar{x}_T, \bar{y}_T; 3\beta) \leq \frac{32M \|z_0 - z^*\|^3}{T^{3/2}}.$$

Let  $M = 3\rho m$ . Algorithm 1 finds an  $\epsilon$ -saddle point within  $\mathcal{O}(m^{2/3} \epsilon^{-2/3})$  iterations.

**Algorithm 2** LEN-restart( $z_0, T, m, M, S$ )

---

```

1:  $z^{(0)} = z_0$ 
2: for  $s = 0, \dots, S - 1$ 
3:    $z^{(s+1)} = \text{LEN}(z^{(s)}, T, m, M)$ 
end for

```

---

**Discussion on the computational complexity of the oracles.** Theorem 4.1 indicates that LEN requires  $\mathcal{O}(m^{2/3}\epsilon^{-2/3})$  calls to  $\mathbf{F}(\cdot)$  and  $\mathcal{O}(m^{2/3}\epsilon^{-2/3}/m + 1)$  calls to  $\nabla\mathbf{F}(\cdot)$  to find the  $\epsilon$ -saddle point. Under Assumption 3.4, the computational cost to call the oracles  $\mathbf{F}(\cdot)$  and  $\nabla\mathbf{F}(\cdot)$  is

$$\text{Oracle Computational Cost} = \mathcal{O}\left(N \cdot m^{2/3}\epsilon^{-2/3} + (Nd) \cdot \left(\epsilon^{-2/3}m^{-1/3} + 1\right)\right). \quad (5)$$

Taking  $m = \Theta(d)$  minimizes (5) to  $\mathcal{O}(Nd + Nd^{2/3}\epsilon^{-2/3})$ . Compared to state-of-the-art second-order methods (Monteiro & Svaiter, 2012; Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022), whose computational cost in terms of the oracles is  $\mathcal{O}(Nde^{-2/3})$  since they require to query  $\nabla\mathbf{F}(\cdot)$  at each iteration, our methods significantly improve their results by a factor of  $d^{1/3}$ .

It is worth noticing that the computational cost of an algorithm includes both the computational cost of calling the oracles, which we have discussed above, and the computational cost of performing the updates (*i.e.* solving auxiliary problems) after accessing the required oracles. We will give an efficient implementation of LEN and analyze the total computational cost later in Section 4.3.

#### 4.2 THE LEN-RESTART ALGORITHM FOR STRONGLY-CONVEX-STRONGLY-CONCAVE PROBLEMS

We generalize LEN to solve the strongly-convex-strongly-concave minimax optimization by incorporating the restart strategy introduced by Huang & Zhang (2022); Lin & Jordan (2024). We propose the LEN-restart in Algorithm 2, which works in epochs. Each epoch of LEN-restart invokes LEN (Algorithm 1), which gets  $z^{(s)}$  as inputs and outputs  $z^{(s+1)}$ .

The following theorem shows that the sequence  $\{z^{(s)}\}$  enjoys a superlinear convergence in epochs. Furthermore, the required number of iterations in each epoch to achieve such a superlinear rate is only a constant.

**Theorem 4.2** (SC-SC setting). *Suppose that Assumptions 3.1 and 3.3 hold. Let  $z^* = (x^*, y^*)$  be the unique saddle point. Set  $M = 3\rho m$  as Theorem 4.1 and  $T = \left(\frac{2M\|z_0 - z^*\|}{\mu}\right)^{2/3}$ . Then the sequence of iterates generated by Algorithm 2 converge to  $z^*$  superlinearly as  $\|z^{(s)} - z^*\|^2 \leq \left(\frac{1}{2}\right)^{(3/2)^s} \|z_0 - z^*\|^2$ . In other words, Algorithm 2 finds a point  $z^{(s)}$  such that  $\|z^{(s)} - z^*\| \leq \epsilon$  within  $S = \log_{3/2} \log_2(1/\epsilon)$  epochs. The total number of inner loop iterations is given by*

$$TS = \mathcal{O}\left(m^{2/3}\kappa^{2/3} \log \log(1/\epsilon)\right).$$

Under Assumption 3.4, Algorithm 2 with  $m = \Theta(d)$  takes the computational complexity of  $\mathcal{O}((Nd + Nd^{2/3}\kappa^{2/3}) \log \log(1/\epsilon))$  to call the oracles  $\mathbf{F}(\cdot)$  and  $\nabla\mathbf{F}(\cdot)$ .

#### 4.3 IMPLEMENTATION DETAILS AND COMPUTATIONAL COMPLEXITY ANALYSIS

We provide details of implementing the cubic regularized Newton oracle (Implicit Step, (4)). Inspired by Monteiro & Svaiter (2012); Bullins & Lai (2022); Adil et al. (2022); Lin et al. (2022), we transform the sub-problem into a root-finding problem for a univariate function.

**Lemma 4.3** (Section 4.3 Lin et al. (2022)). *Suppose Assumption 3.1 and 3.2 hold for function  $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}$  and let  $\mathbf{F}$  be its GDA field. Define  $\gamma_t = M\|z_{t+1/2} - z_t\|$ . The cubic regularized Newton oracle (Implicit Step, (4)) can be rewritten as:*

$$z_{t+1/2} = z_t - (\nabla\mathbf{F}(z_{\pi(t)}) + \gamma_t \mathbf{I}_d)^{-1} \mathbf{F}(z_t),$$

which can be implemented by finding the root of the following univariate function:

$$\phi(\gamma) := M \|(\nabla\mathbf{F}(z_{\pi(t)}) + \gamma \mathbf{I}_d)^{-1} \mathbf{F}(z_t)\| - \gamma. \quad (6)$$

Furthermore, the function  $\phi(\gamma)$  is strictly decreasing when  $\lambda > 0$ .

**Algorithm 3** Implementation of LEN ( $\mathbf{z}_0, T, m, M$ )

---

```

1: for  $t = 0, \dots, T - 1$  do
2:   if  $t \bmod m = 0$  do
3:     Compute the Schur decomposition such that  $\nabla \mathbf{F}(\mathbf{z}_t) = \mathbf{Q}\mathbf{U}\mathbf{Q}^{-1}$ .
4:   end if
5:   Let  $\phi(\cdot)$  defined as 6 and compute  $\gamma_t$  as its root by a binary search.
6:   Compute lazy cubic step  $\mathbf{z}_{t+1/2} = \mathbf{Q}(\mathbf{U} + \gamma_t \mathbf{I}_d)^{-1} \mathbf{Q}^{-1} \mathbf{F}(\mathbf{z}_t)$ .
7:   Compute extra-gradient step  $\mathbf{z}_{t+1} = \mathbf{z}_t - \gamma_t^{-1} \mathbf{F}(\mathbf{z}_{t+1/2})$ .
8: end for
9: return  $\bar{\mathbf{z}}_T = \frac{1}{\sum_{t=0}^{T-1} \gamma_t} \sum_{t=0}^{T-1} \gamma_t^{-1} \mathbf{z}_{t+1/2}$ .

```

---

From the above lemma, to implement the cubic regularized Newton oracle, it suffices to find the root of a strictly decreasing function  $\phi(\gamma)$ , which can be solved within  $\tilde{\mathcal{O}}(1)$  iteration. The main operation is to solve the following linear system:

$$(\nabla \mathbf{F}(\mathbf{z}_{\pi(t)}) + \gamma \mathbf{I}_d) \mathbf{h} = \mathbf{F}(\mathbf{z}_t). \quad (7)$$

Naively solving this linear system at every iteration still results in an expensive computational complexity of  $\mathcal{O}(d^3)$  per iteration.

We present a computationally efficient way to implement LEN by leveraging the Schur factorization at the snapshot point  $\nabla \mathbf{F}(\mathbf{z}_{\pi(t)}) = \mathbf{Q}\mathbf{U}\mathbf{Q}^{-1}$ , where  $\mathbf{Q} \in \mathbb{C}^d$  is a unitary matrix and  $\mathbf{U} \in \mathbb{C}^d$  is an upper-triangular matrix. Then solving the linear system (7) is equivalent to

$$\mathbf{h} = \mathbf{Q}(\mathbf{U} + \gamma \mathbf{I}_d)^{-1} \mathbf{Q}^{-1} \mathbf{F}(\mathbf{z}_t). \quad (8)$$

The final implementable algorithm is presented in Algorithm 3.

Now, we are ready to analyze the total computational complexity of LEN, which can be divided into the following two parts:

$$\text{Computational Cost} = \text{Oracle Computational Cost} + \text{Update Computational Cost},$$

where the first part has been discussed in Section 4.1. Regarding the update computational cost, the Schur decomposition with an computational complexity  $\mathcal{O}(d^3)$  is required once every  $m$  iterations. After Schur's decomposition has been given at the snapshot point, the dominant part of the update computational complexity is solving the upper-triangular linear system (8) with the back substitution algorithm within the computational complexity  $\mathcal{O}(d^2)$ . Thus, we have

$$\text{Update Computational Cost} = \tilde{\mathcal{O}} \left( d^2 \cdot m^{2/3} \epsilon^{-2/3} + d^3 \cdot \left( \epsilon^{-2/3} m^{-1/3} + 1 \right) \right), \quad (9)$$

and the total computational cost of LEN is

$$\text{Computational Cost} \stackrel{(5),(9)}{=} \tilde{\mathcal{O}} \left( (d^2 + N) \cdot m^{2/3} \epsilon^{-2/3} + (d^3 + Nd) \cdot \left( \epsilon^{-2/3} m^{-1/3} + 1 \right) \right). \quad (10)$$

By taking  $m = \Theta(d)$ , we obtain the best computational complexity in (10) of LEN, which is formally stated in the following theorem.

**Theorem 4.3** (C-C setting). *Under the same setting of Theorem 4.1, Algorithm 3 with  $m = \Theta(d)$  finds an  $\epsilon$ -saddle point with computational complexity  $\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3} \epsilon^{-2/3}))$ .*

We also present the total computational complexity of LEN-restart for SC-SC setting.

**Corollary 4.1** (SC-SC setting). *Under the same setting as in Theorem 4.2, Algorithm 2 implemented in the same way as Algorithm 3 with  $m = \Theta(d)$  finds an  $\epsilon$ -root with computational complexity  $\tilde{\mathcal{O}}((N + d^2)(d + d^{2/3} \kappa^{2/3}))$ .*

In both cases, our proposed algorithms improve the total computational cost of the optimal second-order methods (Monteiro & Svaiter, 2012; Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022) by a factor of  $d^{1/3}$ .



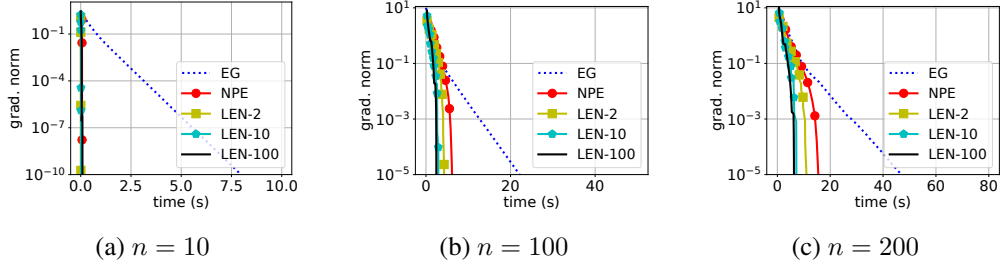


Figure 1: We demonstrate running time v.s. gradient norm  $\|F(z)\|$  for Problem (11) with different sizes:  $n \in \{10, 100, 200\}$ .

**Remark 4.1.** In the main text, we assume the use of the classical algorithm for matrix inversion/decomposition, which has a computational complexity of  $\mathcal{O}(d^3)$ . The fast matrix multiplication proposed by researchers in the field of theoretical computer science only requires a complexity of  $d^\omega$ , where the best known  $\omega$  is currently around 2.371552 (Williams et al., 2024). This also implies faster standard linear algebra operators including Schur decomposition and matrix inversion (Demmel et al., 2007). However, the large hidden constant factors in these fast matrix multiplication algorithms mean that the matrix dimensions necessary for these algorithms to be superior to classical algorithms are much larger than what current computers can effectively handle. Consequently, these algorithms are not always used in practice. We present the computational complexity of using fast matrix operations in Appendix G.

In Appendix H, we extend our algorithms to allow inexact auxiliary CRN sub-problem solving and analyze the total complexity. Specifically, we design an efficient sub-procedure (Algorithm 5) to solve the CRN sub-problem to desired accuracy in only  $\mathcal{O}(\log \log(1/\epsilon))$  number of linear system solving. It tightens the  $\mathcal{O}(\log(1/\epsilon))$  iteration complexity in (Bullins & Lai, 2022; Adil et al., 2022). Additionally, (Bullins & Lai, 2022; Adil et al., 2022) assume  $\sigma_{\min}(\nabla F(z)) \geq \mu$  for some positive constant  $\mu$ , which makes the problem similar to strongly-convex(-strongly-concave) problems, while our analysis does not require such an assumption.

## 5 NUMERICAL EXPERIMENTS

We conduct our algorithms on a regularized bilinear min-max problem and fairness-aware machine learning tasks. We include EG (Korpelevich, 1976) and NPE (Monteiro & Svaiter, 2012; Bullins & Lai, 2022; Adil et al., 2022; Lin et al., 2022) (which is our algorithm with  $m = 1$ ) as baselines, since they are the optimal first- and second-order methods for convex-concave minimax problems, respectively. We run the programs on an AMD EPYC 7H12 64-Core Processor.<sup>1</sup>

### 5.1 REGULARIZED BILINEAR MIN-MAX PROBLEM

We first conduct numerical experiments on the cubic regularized bilinear min-max problem considered in the literature (Lin et al., 2022; Jiang et al., 2024):

$$\min_{\mathbf{x} \in \mathbb{R}^n} \max_{\mathbf{y} \in \mathbb{R}^n} f(\mathbf{x}, \mathbf{y}) = \frac{\rho}{6} \|\mathbf{x}\|^3 + \mathbf{y}^\top (\mathbf{A}\mathbf{x} - \mathbf{b}). \quad (11)$$

The function  $f(\mathbf{x}, \mathbf{y})$  is convex-concave and has  $\rho$ -Lipschitz continuous Hessians. The unique saddle point  $\mathbf{z}^* = (\mathbf{x}^*, \mathbf{y}^*)$  of  $f(\mathbf{x}, \mathbf{y})$  can be explicitly calculated as  $\mathbf{x}^* = \mathbf{A}^{-1}\mathbf{b}$  and  $\mathbf{y}^* = -\rho\|\mathbf{x}^*\|^2(\mathbf{A}^\top)^{-1}\mathbf{x}^*/2$ , so we can compute the distance to  $\mathbf{z}^*$  to measure the performance of the algorithms. Following Lin et al. (2022), we generate each element in  $\mathbf{b}$  as independent Rademacher

variables in  $\{-1, +1\}$ , set  $\rho = 1/(20n)$  and the matrix  $\mathbf{A} = \begin{bmatrix} 1 & -1 & & & \\ & \ddots & \ddots & & \\ & & 1 & -1 & \\ & & & & 1 \end{bmatrix}$ .

<sup>1</sup>The source codes are available at <https://github.com/TrueNobility303/LEN>.

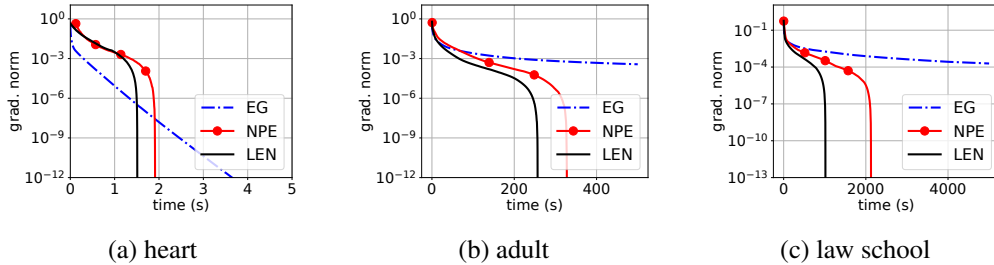


Figure 2: We demonstrate running time v.s. gradient norm  $\|\mathbf{F}(\mathbf{z})\|$  for fairness-aware machine learning task (Problem (12)) on datasets “heart”, “adult”, and “law school”.

We compare our methods with the baselines on different sizes of the problem:  $n \in \{10, 100, 200\}$ . For EG, we tune the stepsize in  $\{1, 0.1, 0.01, 0.001\}$ . For LEN, we vary  $m$  in  $\{1, 2, 10, 100\}$ . The results of the running time against  $\|\mathbf{F}(\mathbf{z})\|$  are presented in Figure 1.

## 5.2 FAIRNESS-AWARE MACHINE LEARNING

We then examine our algorithm for the task of fairness-aware machine learning. Let  $\{\mathbf{a}_i, b_i, c_i\}_{i=1}^n$  be the training set, where  $\mathbf{a}_i \in \mathbb{R}^{d_x}$  denotes the features of the  $i$ -th sample,  $b_i \in \{-1, +1\}$  is the corresponding label, and  $c_i \in \{-1, +1\}$  is an additional feature that is required to be protected and debiased. For example,  $c_i$  can denote gender. Zhang et al. (2018) proposed to solve the following minimax problem to mitigate unwanted bias of  $c_i$  by adversarial learning:

$$\min_{\mathbf{x} \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \ell(b_i \mathbf{a}_i^\top \mathbf{x}) - \beta \ell(c_i y \mathbf{a}_i^\top \mathbf{x}) + \lambda \|\mathbf{x}\|^2 - \gamma y^2, \quad (12)$$

where  $\ell$  is the logit function such that  $\ell(t) = \log(1 + \exp(-t))$ . We set  $\lambda = \gamma = 10^{-4}$  and  $\beta = 0.5$ . We conduct the experiments on datasets “heart” ( $n = 270, d_x = 13$ ) (Chang & Lin, 2011), “adult” ( $n = 32,561, d_x = 123$ ) (Chang & Lin, 2011) and “law school” ( $n = 20,798, d_x = 380$ ) (Le Quy et al., 2022; Liu et al., 2022). For all the datasets, we choose “gender” as the protected feature. For EG, we tune the stepsize in  $\{0.1, 0.01, 0.001\}$ . For second-order methods (NPE and LEN), as we do not know the value of  $\rho$  in advance, we view it as a hyperparameter and tune it in  $\{1, 10, 100\}$ . We set  $m = 10$  for LEN and we find that this simple choice performs well in all the datasets we test. We show the results of the running time against the gradient norm  $\|\mathbf{F}(\mathbf{z})\|$  in Figure 2.

## 6 CONCLUSION AND FUTURE WORKS

In this paper, we propose LEN and LEN-restart for C-C and SC-SC minimax problems, respectively. Using lazy Hessian updates, our methods improve the computational complexity of the current best-known second-order methods by a factor of  $d^{1/3}$ . Numerical experiments on both real and synthetic datasets also verify the efficiency of our method.

For future work, it will be interesting to extend our idea to adaptive second-order methods (Wang et al., 2024a; Doikov et al., 2024; Carmon et al., 2022; Antonakopoulos et al., 2022; Liu & Luo, 2022b) and stochastic problems with sub-sampled Newton methods (Lin et al., 2022; Chayti et al., 2023; Zhou et al., 2019; Tripuraneni et al., 2018; Wang et al., 2019). Besides, our methods only focus on the convex-concave case, it is also possible to reduce the Hessian oracle for nonconvex-(strongly)-concave problems (Luo et al., 2022; Lin et al., 2020; Yang et al., 2023; Zhang et al., 2022b; Wang et al., 2024b) or study structured nonconvex-nonconcave problems (Zheng et al., 2024; Diakonikolas et al., 2021; Yang et al., 2020; Lee & Kim, 2021; Chen & Luo, 2024).

## ACKNOWLEDGMENTS

Lesi Chen and Jingzhao Zhang are supported by the Shanghai Qi Zhi Institute Innovation Program. Chengchang Liu is supported by the National Natural Science Foundation of China (624B2125).

## REFERENCES

- Deeksha Adil, Brian Bullins, Arun Jambulapati, and Sushant Sachdeva. Optimal methods for higher-order smooth monotone variational inequalities. *arXiv preprint arXiv:2205.06167*, 2022.
- Ilan Adler, Zhiyue T. Hu, and Tianyi Lin. New proximal newton-type methods for convex optimization. In *CDC*, 2020.
- M. Marques Alves and Benar F. Svaiter. A search-free  $\mathcal{O}(1/k^{3/2})$  homotopy inexact proximal-newton extragradient algorithm for monotone variational inequalities. *arXiv preprint arXiv:2308.05887*, 2023.
- Kimon Antonakopoulos, Ali Kavis, and Volkan Cevher. Extra-newton: A first approach to noise-adaptive accelerated second-order methods. In *NeurIPS*, 2022.
- Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009.
- Brian Bullins and Kevin A. Lai. Higher-order methods for convex-concave min-max optimization and monotone variational inequalities. *SIAM Journal on Optimization*, 32(3):2208–2229, 2022.
- Yair Carmon, Danielle Hausler, Arun Jambulapati, Yujia Jin, and Aaron Sidford. Optimal and adaptive Monteiro-Svaiter acceleration. In *NeurIPS*, 2022.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- El Mahdi Chayti, Nikita Doikov, and Martin Jaggi. Unified convergence theory of stochastic and variance-reduced cubic newton methods. *arXiv preprint arXiv:2302.11962*, 2023.
- Lesi Chen and Luo Luo. Near-optimal algorithms for making the gradient small in stochastic min-max optimization. *JMLR*, 2024.
- James Demmel, Ioana Dumitriu, and Olga Holtz. Fast linear algebra is stable. *Numerische Mathematik*, 108(1):59–91, 2007.
- Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS*, 2021.
- Nikita Doikov, El Mahdi Chayti, and Martin Jaggi. Second-order optimization with lazy Hessians. In *ICML*, 2023.
- Nikita Doikov, Konstantin Mishchenko, and Yurii Nesterov. Super-universal regularized newton method. *SIAM Journal on Optimization*, 34(1):27–56, 2024.
- Simon S Du, Jianshu Chen, Lihong Li, Lin Xiao, and Dengyong Zhou. Stochastic variance reduction methods for policy evaluation. In *ICML*, 2017.
- Jinyan Fan. A shamanskii-like levenberg-marquardt method for nonlinear equations. *Computational Optimization and Applications*, 56(1):63–80, 2013.
- Roger Grosse and James Martens. A kronecker-factored approximate fisher matrix for convolution layers. In *ICML*, 2016.
- James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- Kevin Huang and Shuzhong Zhang. An approximation-based regularized extra-gradient method for monotone variational inequalities. *arXiv preprint arXiv:2210.04440*, 2022.
- Kevin Huang, Junyu Zhang, and Shuzhong Zhang. Cubic regularized newton method for the saddle point models: A global and local convergence analysis. *Journal of Scientific Computing*, 91(2):60, 2022.
- Ruichen Jiang and Aryan Mokhtari. Generalized optimistic methods for convex-concave saddle point problems. *arXiv preprint arXiv:2202.09674*, 2022.

- Ruichen Jiang, Qiujiang Jin, and Aryan Mokhtari. Online learning guided curvature approximation: A quasi-newton method with global non-asymptotic superlinear convergence. In *COLT*, 2023.
- Ruichen Jiang, Ali Kavis, Qiujiang Jin, Sujay Sanghavi, and Aryan Mokhtari. Adaptive and optimal second-order optimistic methods for minimax optimization. In *NeurIPS*, 2024.
- Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- Dmitry Kovalev, Adil Salim, and Peter Richtárik. Optimal and practical algorithms for smooth and strongly convex decentralized optimization. In *NeurIPS*, 2020.
- Dmitry Kovalev, Elnur Gasanov, Alexander Gasnikov, and Peter Richtarik. Lower bounds and optimal algorithms for smooth and strongly convex decentralized optimization over time-varying networks. In *NeurIPS*, 2021.
- Francesco Lampariello and Marco Sciandrone. Global convergence technique for the newton method with periodic hessian evaluation. *Journal of optimization theory and applications*, 111: 341–358, 2001.
- Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsis. A survey on datasets for fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(3):e1452, 2022.
- Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. In *NeurIPS*, 2021.
- Tianyi Lin and Michael I Jordan. Perseus: A simple high-order regularization method for variational inequalities. *Mathematical Programming*, pp. 1–42, 2024.
- Tianyi Lin, Chi Jin, and Michael I Jordan. Near-optimal algorithms for minimax optimization. In *COLT*, 2020.
- Tianyi Lin, Panayotis Mertikopoulos, and Michael I. Jordan. Explicit second-order min-max optimization methods with optimal convergence guarantee. *arXiv preprint arXiv:2210.12860*, 2022.
- Chengchang Liu and Luo Luo. Quasi-newton methods for saddle point problems. In *NeurIPS*, 2022a.
- Chengchang Liu and Luo Luo. Regularized newton methods for monotone variational inequalities with Holders continuous jacobians. *arXiv preprint arXiv:2212.07824*, 2022b.
- Chengchang Liu, Shuxian Bi, Luo Luo, and John CS Lui. Partial-quasi-newton methods: efficient algorithms for minimax optimization problems with unbalanced dimensionality. In *SIGKDD*, 2022.
- Chengchang Liu, Lesi Chen, Luo Luo, and John CS Lui. Communication efficient distributed newton method with fast convergence rates. In *SIGKDD*, 2023.
- Hong Liu, Zhiyuan Li, David Hall, Percy Liang, and Tengyu Ma. Sophia: A scalable stochastic second-order optimizer for language model pre-training. In *ICLR*, 2024.
- Luo Luo, Yujun Li, and Cheng Chen. Finding second-order stationary points in nonconvex-strongly-concave minimax optimization. In *NeurIPS*, 2022.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *ICML*, 2015.
- Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *AISTATS*, 2020a.
- Aryan Mokhtari, Asuman E. Ozdaglar, and Sarath Pattathil. Convergence rate of  $\mathcal{O}(1/k)$  for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251, 2020b.

- Renato DC Monteiro and Benar F Svaiter. Iteration-complexity of a newton proximal extragradient method for monotone variational inequalities and inclusion problems. *SIAM Journal on Optimization*, 22(3):914–935, 2012.
- Renato DC Monteiro and Benar Fux Svaiter. On the complexity of the hybrid proximal extragradient method for the iterates and the ergodic mean. *SIAM Journal on Optimization*, 20(6):2755–2787, 2010.
- Arkadi Nemirovski. Prox-method with rate of convergence  $\mathcal{O}(1/t)$  for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2-3):319–344, 2007.
- Yurii Nesterov and Boris T Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- Yurii Nesterov and Laura Scrimali. Solving strongly monotone variational and quasi-variational inequalities. 2006.
- Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *IEEE Transactions on Automatic Control*, 68(3):1321–1336, 2022.
- Leonid Denisovich Popov. A modification of the arrow-hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28:845–848, 1980.
- R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, 14(5):877–898, 1976.
- VE Shamanskii. A modification of newton’s method. *Ukrainian Mathematical Journal*, 19(1): 118–122, 1967.
- Nilesh Tripuraneni, Mitchell Stern, Chi Jin, Jeffrey Regier, and Michael I Jordan. Stochastic cubic regularization for fast nonconvex optimization. In *NeurIPS*, 2018.
- John Von Neumann and Oskar Morgenstern. Theory of games and economic behavior, 2nd rev. 1947.
- Hoi-To Wai, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong. Multi-agent reinforcement learning via double averaging primal-dual optimization. In *NeurIPS*, 2018.
- Chang-yu Wang, Yuan-yuan Chen, and Shou-qiang Du. Further insight into the shamanskii modification of newton method. *Applied mathematics and computation*, 180(1):46–52, 2006.
- Junlin Wang, Junnan Yang, and Zi Xu. A fully parameter-free second-order algorithm for convex-concave minimax problems with optimal iteration complexity. *arXiv preprint arXiv:2407.03571*, 2024a.
- Mengdi Wang. Primal-dual  $\pi$ -learning: Sample complexity and sublinear run time for ergodic markov decision problems. *arXiv preprint arXiv:1710.06100*, 2017.
- Nuozhou Wang, Junyu Zhang, and Shuzhong Zhang. Efficient first order method for saddle point problems with higher order smoothness. *SIAM Journal on Optimization*, 34(4):3342–3370, 2024b.
- Zhe Wang, Yi Zhou, Yingbin Liang, and Guanghui Lan. Stochastic variance-reduced cubic regularization for nonconvex optimization. In *AISTATS*, 2019.
- Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. New bounds for matrix multiplication: from alpha to omega. In *SODA*, 2024.

- Stephen J. Wright. Numerical optimization, 2006.
- Yangyang Xu. Primal-dual stochastic gradient method for convex programs with many functional constraints. *SIAM Journal on Optimization*, 30(2):1664–1692, 2020.
- Haikuo Yang, Luo Luo, Chris Junchi Li, Michael Jordan, and Maryam Fazel. Accelerating inexact hypergradient descent for bilevel optimization. In *OPT 2023: Optimization for Machine Learning*, 2023.
- Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In *NeurIPS*, 2020.
- Yiming Ying, Longyin Wen, and Siwei Lyu. Stochastic online auc maximization. In *NIPS*, 2016.
- Zhuoning Yuan, Yan Yan, Milan Sonka, and Tianbao Yang. Large-scale robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification. In *ICCV*, 2021.
- Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.
- Junyu Zhang, Mingyi Hong, and Shuzhong Zhang. On lower iteration complexity bounds for the convex concave saddle point problems. *Mathematical Programming*, 194(1-2):901–935, 2022a.
- Xuan Zhang, Necdet Serhat Aybat, and Mert Gurbuzbalaban. SAPD+: An accelerated stochastic method for nonconvex-concave minimax problems. In *NeurIPS*, 2022b.
- Taoli Zheng, Linglingzhi Zhu, Anthony Man-Cho So, José Blanchet, and Jiajin Li. Universal gradient descent ascent method for nonconvex-nonconcave minimax optimization. In *NeurIPS*, 2024.
- Dongruo Zhou, Pan Xu, and Quanquan Gu. Stochastic variance-reduced cubic regularization methods. *JMLR*, 20(134):1–47, 2019.