

Near-Optimal Algorithms for Making the Gradient Small in Stochastic Minimax Optimization

Lesi Chen

CHENLC23@MAILS.EDU.CN

*Institute for Interdisciplinary Information Sciences, Tsinghua University
Beijing, China*

Luo Luo*

LUOLUO@FUDAN.EDU.CN

*School of Data Science, Fudan University
Shanghai, China*

and

*Shanghai Key Laboratory for Contemporary Applied Mathematics
Shanghai, China*

Editor: Lam Nguyen

Abstract

We study the problem of finding a near-stationary point for smooth minimax optimization. The recently proposed extra anchored gradient (EAG) methods achieve the optimal convergence rate for the convex-concave minimax problem in the deterministic setting. However, the direct extension of EAG to stochastic optimization is not efficient. In this paper, we design a novel stochastic algorithm called Recursive Anchored IterationN (RAIN). We show that the RAIN achieves near-optimal stochastic first-order oracle (SFO) complexity for stochastic minimax optimization in both convex-concave and strongly-convex-strongly-concave cases. In addition, we extend the idea of RAIN to solve structured nonconvex-nonconcave minimax problem and it also achieves near-optimal SFO complexity.

Keywords: Stochastic minimax optimization, ϵ -stationary point, recursive anchoring

1. Introduction

This work studies the stochastic minimax problem of the form:

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f(x, y) \triangleq \mathbb{E}[f(x, y; \xi)], \quad (1)$$

where the stochastic component $f(x, y; \xi)$ is indexed by some random variable ξ ; and the objective function $f(x, y)$ is L -smooth. This formulation has aroused great interest in machine learning community (Lin et al., 2020a; Liu and Orabona, 2022; Zhang et al., 2021; Yang et al., 2020a; Xu et al., 2020; Xian et al., 2021; Zhang et al., 2022a,b) due to its wide applications, including generative adversarial networks (GANs) (Goodfellow et al., 2014a,b; Liu et al., 2020), AUC maximization (Guo et al., 2020; Yuan et al., 2021; Liu et al., 2019; Yang and Ying, 2022), adversarial training (Madry et al., 2017) and multi-agent reinforcement learning (Omidshafiei et al., 2017; Dai et al., 2018; Wai et al., 2018).

*. The corresponding author

We focus on using stochastic first-order oracle (SFO) algorithms to find an ϵ -stationary point of problem (1), that is, a point (x, y) where the gradient satisfies

$$\|\nabla f(x, y)\| \leq \epsilon.$$

The norm of gradient is a generic metric to quantify the sub-optimality for solving smooth optimization problem. It is always well-defined for differentiable objective functions and helps us understand general minimax problems without convex-concave assumption. Some recent works consider making the gradient small in minimax optimization (Yoon and Ryu, 2021; Cai et al., 2022; Diakonikolas et al., 2021; Lee and Kim, 2021), but the understanding of this task in stochastic setting is still limited.

For convex-concave minimax problems, a popular metric to quantify an approximate solution (x, y) on the domain $\mathcal{X} \times \mathcal{Y}$ is the duality gap, which is defined as

$$\text{Gap}(x, y) \triangleq \max_{y' \in \mathcal{Y}} f(x, y') - \min_{x' \in \mathcal{X}} f(x', y).$$

The optimal SFO algorithms in terms of duality gap have been established (Zhao, 2022; Alacaoglu and Malitsky, 2022), while the duality gap is difficult to be measured and it is not always well-defined. In contrast, the norm of gradient is a more general metric in optimization problems and it is easy to be measured in practice. Yoon and Ryu (2021) proposed the extra anchor gradient (EAG) method for finding ϵ -stationary point of deterministic (offline) convex-concave minimax problem. They showed that EAG is an optimal deterministic first-order algorithm to find an ϵ -stationary point and its exact first-order oracle upper bound complexity is $\mathcal{O}(L\epsilon^{-1})$. For the stochastic convex-concave minimax problem, Lee and Kim (2021) proved the stochastic EAG (SEAG) can achieve an SFO complexity of $\mathcal{O}(\sigma^2 L^2 \epsilon^{-4} + L\epsilon^{-1})$, where σ^2 is a bound on the variance of gradient estimates. Recently, Cai et al. (2022) proposed a variant of stochastic Halpern iteration, which has an SFO complexity of $\tilde{\mathcal{O}}(\sigma^2 L\epsilon^{-3} + L^3 \epsilon^{-3})$ under the additional average-smooth assumption. However, the tightness of SFO complexity for finding ϵ -stationary points in stochastic convex-concave minimax optimization was hitherto unclear.

For nonconvex-nonconcave minimax problems, there exists the counter-example that all first-order algorithms will diverge (Lee and Kim, 2021). Hence, we are required to introduce additional assumptions. Grimmer et al. (2023) proposed the intersection dominant condition to relax the convex-concave assumption and established the convergence guarantee of the exact proximal point method in such a setting. Later, Diakonikolas et al. (2021) considered the negative comonotonicity condition as a relaxation of the common monotonicity property in the convex-concave setting. They proposed a variant of stochastic extragradient named as EG^+ , which has SFO upper bound of $\mathcal{O}(\sigma^2 L^2 \epsilon^{-4} + L^2 \epsilon^{-2})$ for finding an ϵ -stationary point, matching the complexity of ordinary stochastic extragradient (SEG) in convex-concave case. In addition, Lee and Kim (2021) extend the idea of EAG to solve minimax problem under the negative comonotonicity condition, but their analysis does not contain the stochastic algorithms.

In this paper, we propose a novel stochastic algorithm called Recursive Anchored IterationN (RAIN) to make the gradient small in stochastic minimax optimization. The algorithm solves a sequence of anchored sub-problems by the variant of the stochastic extragradient method, resulting the SFO upper bounds complexity of $\tilde{\mathcal{O}}(\sigma^2 \epsilon^{-2} + L\epsilon^{-1})$ and $\tilde{\mathcal{O}}(\sigma^2 \epsilon^{-2} + \kappa)$

Table 1: We summarize the SFO complexities for finding an ϵ -stationary point in convex-concave setting. We denote SCSC and CC as λ -strongly-convex- λ -strongly-concave case and the general convex-concave case respectively. In the case of SCSC, we define the condition number as $\kappa \triangleq L/\lambda$. We also define $D \triangleq \|z_0 - z^*\|$, where z_0 is the initial point of the algorithm. The dependency on D in the complexity of PDHG and RAIN (the case of SCSC) is only contained in the logarithmic term, which is omitted by using the notation $\tilde{\mathcal{O}}(\cdot)$. The results denoted by * indicate the analysis Cai et al. (2022) requires the additional assumption of average smoothness.

Setting	Algorithm	Complexity	Reference
SCSC	Halpern	$\tilde{\mathcal{O}}(\lambda^{-2}\kappa\sigma^2\epsilon^{-2} + \kappa^3D^2\epsilon^{-2})^*$	Cai et al. (2022)
	PDHG	$\tilde{\mathcal{O}}(\kappa\sigma^2\epsilon^{-2} + \kappa)$	Zhao (2022)
	RAIN	$\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + \kappa)$	Theorem 4.1
	Lower Bound	$\tilde{\Omega}(\sigma^2\epsilon^{-2} + \kappa)$	Theorem 6.1
CC	SEG	$\mathcal{O}(\sigma^2L^2\epsilon^{-4} + L^2D^2\epsilon^{-2})$	Diakonikolas et al. (2021)
	SEAG	$\mathcal{O}(\sigma^2L^2\epsilon^{-4} + LD\epsilon^{-1})$	Lee and Kim (2021)
	Halpern	$\tilde{\mathcal{O}}(\sigma^2LD\epsilon^{-3} + L^3D^3\epsilon^{-3})^*$	Cai et al. (2022)
	RAIN	$\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + LD\epsilon^{-1})$	Theorem 4.2
	Lower Bound	$\tilde{\Omega}(\sigma^2\epsilon^{-2} + LD\epsilon^{-1})$	Theorem 6.2

for finding an ϵ -stationary point in convex-concave and strongly-convex-strongly-concave settings respectively, where κ is the condition number. We also provide lower bounds to show the optimality of RAIN. We summarize our results for the convex-concave problem and compare them with prior work in Table 1. Additionally, we extend the idea of RAIN and propose the algorithm for solving stochastic nonconvex-nonconcave minimax problems. We prove that the SFO upper bound complexity of RAIN⁺⁺ is near-optimal under both intersection dominant (Grimmer et al., 2023) and negative comonotonic conditions (Diakonikolas et al., 2021). We present the results for these structured nonconvex-nonconcave minimax in Table 2. To the best of our knowledge, RAIN is the first near-optimal SFO algorithm for finding near-stationary points of the stochastic convex-concave minimax problem.

2. Notation and Preliminaries

We use $\|\cdot\|$ to present the Euclidean norm of matrix and vector respectively. For differentiable function $f(x, y)$, we denote $\nabla_x f(x, y)$ and $\nabla_y f(x, y)$ as its partial gradients with respect to x and y respectively, and introduce the gradient operator

$$F(x, y) \triangleq \begin{bmatrix} \nabla_x f(x, y) \\ -\nabla_y f(x, y) \end{bmatrix}.$$

For ease of presentation, we define $z = (x, y)$ and also write the gradient operator at $z = (x, y)$ as $F(z)$.

We consider the following assumptions for our stochastic minimax problems.

Table 2: We summarize the SFO complexities for finding an ϵ -stationary point in two specific nonconvex-nonconcave settings. We denote NC as the ρ -comonotonicity assumption with $\rho \in [-\frac{1}{2L}, 0)$ and denote ID as the (τ, α) -intersection-dominant assumption with $\alpha > 0$ and $\tau \geq 2L$. We also define $D \triangleq \|z_0 - z^*\|$, where z_0 is the initial point of the algorithm. The dependency on D in the complexity of RAIN (the case of ID) is only contained in the logarithmic term, which is omitted by using the notation $\tilde{O}(\cdot)$.

Setting	Algorithm	Complexity	Reference
NC	SEG ⁺	$\mathcal{O}(\sigma^2 L^2 \epsilon^{-4} + L^2 D^2 \epsilon^{-2})$	Diakonikolas et al. (2021)
	RAIN ⁺⁺	$\tilde{O}(\sigma^2 \epsilon^{-2} + LD \epsilon^{-1})$	Theorem 5.4
	Lower Bound	$\tilde{\Omega}(\sigma^2 \epsilon^{-2} + LD \epsilon^{-1})$	Corollary 6.1
ID	RAIN ⁺⁺	$\tilde{O}(\sigma^2 \epsilon^{-2} + L/\alpha)$	Theorem 5.3
	Lower Bound	$\tilde{\Omega}(\sigma^2 \epsilon^{-2} + L/\alpha)$	Corollary 6.2

Assumption 2.1 (stochastic first-order oracle) *We suppose the stochastic first-order oracle $F(z; \xi)$ is unbiased and has bounded variance such that $\mathbb{E}[F(z; \xi)] = F(z)$ and $\mathbb{E}\|F(z; \xi) - F(z)\|^2 \leq \sigma^2$ for all $z = (x, y)$ and random index ξ .*

Assumption 2.2 (smoothness) *We suppose the function $f(x, y)$ is L -smooth, that is there exists some $L > 0$ such that*

$$\|F(z) - F(z')\| \leq L\|z - z'\|$$

for all $z = (x, y)$ and $z' = (x', y')$.

Assumption 2.3 (convex-concave) *We suppose the function $f(x, y)$ is convex-concave (CC), that is the function $f(\cdot, y)$ is convex for any given y and the function $f(x, \cdot)$ is concave for any given x .*

Assumption 2.4 (strongly-convex-strongly-concave) *We suppose the function $f(x, y)$ is λ -strongly-convex- λ -concave for some $\lambda > 0$, that is the function $f(x, y) + \frac{\lambda}{2}\|x\|^2 - \frac{\lambda}{2}\|y\|^2$ is convex-concave.*

For smooth and convex-concave objective function, the corresponding gradient operator F has the monotonicity properties Rockafellar (1970) as follows.

Lemma 2.1 (monotonicity) *Under Assumption 2.2 and 2.3, it holds that*

$$(F(z) - F(z'))^\top (z - z') \geq 0$$

for all $z = (x, y)$ and $z' = (x', y')$.

Lemma 2.2 (strong monotonicity) *Under Assumption 2.2 and 2.4, it holds that*

$$(F(z) - F(z'))^\top (z - z') \geq \lambda\|z - z'\|^2$$

for all $z = (x, y)$ and $z' = (x', y')$.

We are interested in finding an ϵ -stationary point of the differentiable function $f(x, y)$, that is the point where the norm of its gradient (gradient operator) is small.

Definition 2.1 (nearly-stationary point) *We say $\hat{z} = (\hat{x}, \hat{y})$ is an ϵ -stationary point of the differentiable function $f(x, y)$ if it satisfies $\|\nabla f(\hat{z})\| \leq \epsilon$, or equivalently $\|F(\hat{z})\| \leq \epsilon$.*

Throughout this paper, we always assume there exists a stationary point for $f(x, y)$ and the initial point $z_0 = (x_0, y_0)$ of the considered algorithm is in a bounded set.

Assumption 2.5 *We suppose there exists some $z^* = (x^*, y^*)$ such that $\nabla f(z^*) = 0$, or equivalently, $F(z^*) = 0$.*

Assumption 2.6 *We suppose the initial point $z_0 = (x_0, y_0)$ satisfies $\|z_0 - z^*\| \leq D$ for some $D > 0$, where $z^* = (x^*, y^*)$ is a stationary point of $f(x, y)$.*

3. The Recursive Anchored Iteration

In this section, we focus on stochastic minimax optimization in the convex-concave setting. We propose the Recursive Anchored Iteration (RAIN) method in Algorithm 1. The RAIN calls the subroutine Epoch-SEG (Algorithm 2) to find the point $z_{s+1} = (x_{s+1}, y_{s+1})$ that is an approximate solution of the two-sided regularized minimax problem

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f^{(s)}(x, y), \quad (2)$$

where $f^{(s)}(x, y)$ is defined as follows

$$f^{(s)}(x, y) \triangleq \begin{cases} f(x, y), & s = 0, \\ f^{(s-1)}(x, y) + \frac{\lambda_s}{2} \|x - x_s\|^2 - \frac{\lambda_s}{2} \|y - y_s\|^2, & s \geq 1. \end{cases} \quad (3)$$

We call the sequence $\{(x_s, y_s)\}_{s=1}^S$ as the anchors of RAIN, which plays an important role in our convergence analysis.

For strongly-convex-strongly-concave objective function $f(x, y)$, the corresponding gradient operator $F(x, y)$ is strongly-monotone and the output of RAIN has the following property.

Lemma 3.1 (recursively anchoring lemma) *Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave. Let z_s^* be the unique solution to sub-problem (2), then the output of RAIN (Algorithm 1) with $\gamma = 1$ holds that*

$$\|F(z_S)\| \leq 16\lambda\gamma \sum_{s=1}^S (1 + \gamma)^{s-1} \|z_{s-1}^* - z_s\|.$$

Remark 3.1 *We directly set $\gamma = 1$ for the simplification of the proof. In fact, setting γ to be any positive constant would lead to the upper bound of the same order.*

Algorithm 1 RAIN ($f, z_0, \lambda, L, \{N_s\}_{s=0}^{S-1}, \{K_s\}_{s=0}^{S-1}, \gamma$)

- 1: $f_0 \triangleq f, \lambda_0 = \lambda\gamma, S = \lfloor \log_{(1+\gamma)}(L/\lambda) \rfloor$
 - 2: **for** $s = 0, 1, \dots, S-1$
 - 3: $z_{s+1} \leftarrow \text{Epoch-SEG}(f_s, z_s, \lambda_s, 2L, N_s, K_s)$
 - 4: $\lambda_{s+1} \leftarrow (1 + \gamma)\lambda_s$
 - 5: $f_{s+1}(z) \triangleq f_s(z) + \frac{\lambda_{s+1}}{2}\|x - x_{s+1}\|^2 - \frac{\lambda_{s+1}}{2}\|y - y_{s+1}\|^2$
 - 6: **return** z_S
-

Lemma 3.1 indicates we can make $\|F(z_S)\|$ small if the subroutine provide z_s that is sufficiently close to z_{s-1}^* at each round. The recursive definition of $f_s(x, y)$ means the condition numbers of the sub-problems is decreasing with the iteration of RAIN. Hence, achieving an accurate solution to sub-problem (2) will not be too expensive for large s . We will show that the total SFO complexity of RAIN could nearly match the lower bound by designing the sub-problem solver carefully.

For general convex-concave objective function $f(x, y)$, we use the initial point $z_0 = (x_0, y_0)$ as the additional anchor and then we apply RAIN to find an ϵ -stationary point of the following strongly-convex-strongly-concave function

$$g(x, y) \triangleq f(x, y) + \frac{\lambda}{2}\|x - x_0\|^2 - \frac{\lambda}{2}\|y - y_0\|^2. \quad (4)$$

We denote the gradient operator of $g(x, y)$ as $G(x, y)$. Then the following lemma provides the connection between the norms of $F(x, y)$ and $G(x, y)$.

Lemma 3.2 (anchoring lemma) *Suppose the function $f(x, y)$ is smooth and convex-concave. We define $g(x, y) \triangleq f(x, y) + \frac{\lambda}{2}\|x - x_0\|^2 - \frac{\lambda}{2}\|y - y_0\|^2$ for some (x_0, y_0) and denote its gradient operator as*

$$G(x, y) \triangleq \begin{bmatrix} \nabla_x g(x, y) \\ -\nabla_y g(x, y) \end{bmatrix},$$

then it holds that

$$\|F(\tilde{z})\| \leq 2\|G(\tilde{z})\| + \lambda\|z_0 - z^*\|,$$

for any $\tilde{z} = (\tilde{x}, \tilde{y})$, where $z^* = (x^*, y^*)$ is the stationary point of $f(x, y)$ and $z_0 = (x_0, y_0)$.

According to Assumption 2.6 and Lemma 3.2, setting $\lambda = \Theta(\epsilon/D)$ leads to finding an ϵ -stationary point of $f(x, y)$ can be reduced into finding an $\mathcal{O}(\epsilon)$ -stationary point of $g(x, y)$, which can be done by applying RAIN (Algorithm 1) on the strongly-convex-strongly-concave function $g(x, y)$.

Connection to Related Work We provide some discussion on comparing RAIN with related work.

- In convex optimization, Allen-Zhu (2018) proposed the recursive regularization technique for finding the nearly stationary point stochastically, which can be regarded as

a special case of our anchored framework. However, Allen-Zhu’s analysis depends on the convexity of the objective function, which is not suitable for minimax problem. In contrast, the analysis of RAIN is mainly based on the monotonicity of the gradient operator, which is more general than convex optimization.

- In minimax optimization, Yoon and Ryu (2021); Lee and Kim (2021) considered variants of (stochastic) extragradient method by using initial point z_0 as the fixed anchor. In contrast, the proposed algorithm RAIN adjusts the anchoring point z_s with iterations, which leads to the sequence of anchoring points $\{z_s\}$ converges to z^* . As a result, the RAIN achieves near-optimal SFO complexity in the stochastic setting (see Table 1).
- Several existing methods (Lin et al., 2020b; Zhao, 2022; Kovalev and Gasnikov, 2022; Luo et al., 2021; Yang et al., 2020b) introduce the proximal point iteration

$$(x_{s+1}, y_{s+1}) \approx \arg \min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f(x, y) + \frac{\beta}{2} \|x - x_s\|^2,$$

which is useful to establish the near-optimal algorithms for unbalanced minimax optimization in the offline scenario. However, it is questionable whether the one-sided regularization is helpful in finding near-stationary points in stochastic minimax problem (1).

4. Complexity Analysis for RAIN

In this section, we analyze the sub-problem solver Epoch Stochastic ExtraGradient (Epoch-SEG) and show our RAIN has near-optimal SFO upper bound for finding ϵ -stationary point of stochastic convex-concave minimax problem.

The procedure of Epoch-SEG (Algorithm 2) depends on the ordinary stochastic extragradient method (SEG, Algorithm 3), which has the following property.

Lemma 4.1 (SEG) *Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave, and the SFO $F(x, y; \xi)$ is unbiased and has variance bounded by σ^2 . Then SEG (Algorithm 3) holds that*

$$\lambda \mathbb{E} \|z_{t+1/2} - z^*\|^2 \leq \frac{1}{\eta} \mathbb{E} [\|z_{t+1} - z^*\|^2 - \|z_t - z^*\|^2] + 16\eta\sigma^2 \quad (5)$$

for any $0 < \eta < 1/(4L)$.

Taking the average on (5) over $t = 0, \dots, T - 1$ and applying Lemma 2.2, we know the output of SEG satisfies

$$\mathbb{E} \|\bar{z} - z^*\|^2 \leq \frac{1}{\lambda\eta T} \mathbb{E} \|z_0 - z^*\|^2 + \frac{16\eta\sigma^2}{\lambda},$$

which means SEG is able to decrease the distance from the output $\bar{z}$ to the optimal solution z^* with iterations. However, it only converges to a neighborhood of z^* by using the fixed stepsize.

Algorithm 2 Epoch-SEG (f, z_0, λ, L, N, K)

```

1: for  $k = 0, 1, \dots, N - 1$  do
2:    $z_{k+1} \leftarrow \text{SEG}(f, z_k, \frac{1}{4L}, \frac{8L}{\lambda})$ 
3: for  $k = N, N + 1, \dots, N + K - 1$  do
4:    $z_{k+1} \leftarrow \text{SEG}(f, z_k, \frac{1}{2^{k-N+3}L}, \frac{2^{k-N+5}L}{\lambda})$ 
5: return  $z_{N+K}$ .

```

Algorithm 3 SEG (f, z_0, η, T)

```

1: for  $t = 0, 1, \dots, T - 1$  do
2:    $\xi_i \leftarrow$  a random index
3:    $z_{t+1/2} \leftarrow z_t - \eta F(z_t; \xi_i)$ 
4:    $\xi_j \leftarrow$  a random index
5:    $z_{t+1} \leftarrow z_t - \eta F(z_{t+1/2}; \xi_j)$ 
6: return  $\bar{z}$  by uniformly sampling from  $\{z_{t+1/2}\}_{t=0}^{T-1}$ 

```

Then we consider the epoch stochastic extragradient (Epoch-SEG, Algorithm 2), which consists of two phases and each of them calling SEG as subroutine by different parameters. The Epoch-SEG targets to reduce both the optimization error and statistical error in the iterations:

- In the first phase, we call SEG by fixed stepsize and fixed iteration numbers to decrease the optimization error, which is related to the number $\kappa \triangleq L/\lambda$ and the distance $\|z_0 - z^*\|$.
- In the second phase, the statistical error aroused from the variance of stochastic oracle has accumulated. Hence, we call SEG by decreasing stepsizes and increasing iteration numbers to reduce the statistical error.

The formal theoretical guarantee of Epoch-SEG is shown in Lemma 4.2.

Lemma 4.2 (Epoch-SEG) *Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave, and the stochastic first-order oracle $F(x, y; \xi)$ is an unbiased estimator of $F(x, y)$ and has variance bounded by σ^2 . Then Epoch-SEG (Algorithm 2) holds that*

$$\mathbb{E}\|z_{N+K} - z^*\|^2 \leq \underbrace{\frac{1}{2^{N+2K}} \mathbb{E}\|z_0 - z^*\|^2}_{\text{optimization error}} + \underbrace{\frac{8\sigma^2}{2^K \lambda L}}_{\text{statistical error}}.$$

Additionally, the total number of SFO calls is no more than $16\kappa N + 2^{K+6}\kappa$, where $\kappa = L/\lambda$.

Combining the above results, we obtain the SFO upper bound complexity of RAIN (Algorithm 1) for the strongly-convex-strongly-concave case.

Theorem 4.1 (RAIN, SCSC) *Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave, and the stochastic first-order oracle $F(x, y; \xi)$ is an unbiased estimator of $F(x, y)$ and has variance bounded by σ^2 . If we run RAIN (Algorithm 1) with*

$$N_0 = \begin{cases} \left\lceil \log_2 \left(\frac{512\lambda^2 S^2 D^2}{\epsilon^2} \right) \right\rceil, & s = 0, \\ 3, & s \geq 1, \end{cases} \quad \text{and} \quad K_s = \left\lceil \log_2 \left(\frac{2048\lambda_s S^2 \sigma^2}{L\epsilon^2} \right) \right\rceil, \quad (6)$$

then the output z_S satisfies $\mathbb{E}\|F(z_S)\| \leq \epsilon$ and the total number of SFO calls is no more than

$$\mathcal{O} \left(\frac{L}{\lambda} + \frac{L}{\lambda} \log \left(\frac{\lambda D}{\epsilon} \log \left(\frac{L}{\lambda} \right) \right) + \frac{\sigma^2}{\epsilon^2} \log^3 \left(\frac{L}{\lambda} \right) \right).$$

We can find nearly stationary points for the general convex-concave case by introducing the regularized function $g(x, y)$ defined in (4). Applying Theorem 4.1 and Theorem 3.2, we achieve the SFO upper bound complexity as follows.

Theorem 4.2 (RAIN, CC) *Suppose the function $f(x, y)$ is L -smooth and convex-concave, and the stochastic first-order oracle $F(x, y; \xi)$ is an unbiased estimator of $F(x, y)$ and has variance bounded by σ^2 . Running RAIN (Algorithm 1) on function*

$$g(x, y) \triangleq f(x, y) + \frac{\lambda}{2} \|x - x_0\|^2 - \frac{\lambda}{2} \|y - y_0\|^2$$

with $\lambda = \min \{\epsilon/D, L\}$ outputs a 3ϵ -stationary point in expectation, and the total number of SFO calls is no more than

$$\mathcal{O} \left(\frac{LD}{\epsilon} + \frac{LD}{\epsilon} \log \log \left(\frac{LD}{\epsilon} \right) + \frac{\sigma^2}{\epsilon^2} \log^3 \left(\frac{LD}{\epsilon} \right) \right).$$

The comparison in Table 1 shows both the results of Theorem 4.1 and Theorem 4.2 are better than existing algorithms and nearly match the lower bound. Additionally, the theoretical results in this section are not limited to convex-concave minimax optimization. In fact, our analysis is mainly based on the Lipschitz continuity and the monotonicity of the gradient operator, which means the results are also applicable to more general problems of variational inequality with Lipschitz continuous and monotone operator (Rockafellar, 1970).

5. Extension to Nonconvex-Nonconcave Settings

In this section, we extend RAIN to solve nonconvex-nonconcave minimax problems. We focus on two settings of comonotone and intersection dominant conditions. The comonotonicity is defined as follows.

Definition 5.1 (comonotonicity) *We say the operator $F(\cdot)$ is ρ -comonotone if there exists some $\rho \in \mathbb{R}$ such that*

$$(F(z) - F(z'))^\top (z - z') \geq \rho \|F(z) - F(z')\|^2$$

for all z and z' . We also say $F(\cdot)$ is a negative comonotonic operator when $\rho < 0$.

Remark 5.1 In the case of $\rho = 0$, the ρ -comonotonicity reduces to monotonicity shown in Lemma 2.1. In the case of $\rho < 0$, the ρ -comonotone gradient operator allow the objective function be nonconvex-nonconcave. Typically, we additionally require $\rho \geq -\mathcal{O}(1/L)$ in the convergence analysis for optimization algorithms in the negative comonotone setting (Dikonikolas et al., 2021; Lee and Kim, 2021).

The intersection dominant condition (Grimmer et al., 2023) also allows the objective function be nonconvex-nonconcave, which is defined as follows.

Definition 5.2 (intersection dominant condition) For a twice differentiable function $f(x, y)$, we say it satisfies the (τ, α) -intersection-dominant condition if there exist some $\tau > 0$ and $\alpha > 0$ such that

$$\nabla_{xx}^2 f(x, y) + \nabla_{xy}^2 f(x, y)(\tau I - \nabla_{yy}^2 f(x, y))^{-1} \nabla_{yx}^2 f(x, y) \succeq \alpha I$$

and

$$-\nabla_{yy}^2 f(x, y) + \nabla_{yx}^2 f(x, y)(\tau I + \nabla_{xx}^2 f(x, y))^{-1} \nabla_{xy}^2 f(x, y) \succeq \alpha I$$

for all (x, y) .

Remark 5.2 Typically, we also require $\tau \geq \Omega(L)$ in conditions of intersection dominant for the analysis for minimax optimization (Grimmer et al., 2023; Lee and Kim, 2021).

Remark 5.3 For smooth minimax optimization, the intersection dominant condition is stronger than negative comonotonicity. Concretely, the intersection dominant condition needs the function $f(x, y)$ to be twice differentiable, which is unnecessary to negative comonotonicity. Furthermore, if the function $f(x, y)$ is L -smooth and satisfies the (τ, α) -intersection-dominant condition for some $\alpha > 0$ and $\tau > L$, then its gradient operator $F(\cdot)$ must satisfy the $-1/\tau$ -comonotonic condition (see Example 1 of Lee and Kim (2021)).

It turns out that both of these two conditions are related to the saddle envelope, which is a natural generalization of Moreau envelope for minimax problems.

Definition 5.3 (saddle envelope) Given some $\tau > 0$, we define the saddle envelope of function $f(x, y)$ as

$$f_\tau(x, y) \triangleq \min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x', y') + \frac{\tau}{2} \|x' - x\|^2 - \frac{\tau}{2} \|y' - y\|^2.$$

The saddle envelope has the following properties.

Proposition 5.1 Suppose function $f(x, y)$ is L -smooth and satisfies the (τ, α) -intersection dominant condition for $\tau \geq 2L$, then $f_{2L}(x, y)$ is λ -strongly-convex- λ -strongly-concave, where

$$\lambda = \left(\frac{1}{2L} + \frac{1}{\alpha} \right)^{-1} = \Theta(\alpha). \quad (7)$$

Proposition 5.2 *Suppose the function $f(x, y)$ is L -smooth and convex-concave, and its gradient operator $F(x, y)$ is $-\rho$ -comonotonic with $0 < \rho \leq 1/(2L)$, then its saddle envelope $f_{2L}(x, y)$ is convex-concave.*

Based on the observation that the stationary points of $f_{2L}(x, y)$ are exactly the same as those of $f(x, y)$ (Grimmer et al., 2023, Corollary 2.2), it is possible to apply RAIN on the saddle envelope $f_{2L}(x, y)$ for nonconvex-nonconcave minimax optimization. However, accessing the (stochastic) gradient operator of $f_{2L}(x, y)$ is non-trivial since it cannot be obtained by the gradient of $f(x, y)$ directly. Hence, we maintain the stochastic estimator for the gradient operator of $f_{2L}(x, y)$ as follows:

1. We first denote the gradient operator of $f_{2L}(x, y)$ by

$$F_{2L}(x, y) = \begin{bmatrix} \nabla_x f_{2L}(x, y) \\ -\nabla_y f_{2L}(x, y) \end{bmatrix},$$

which also can be written as (Grimmer et al., 2023):

$$F_{2L}(z) = F(z^+) = 2L(z - z^+),$$

where $z = (x, y)$ and

$$z^+ = (x^+, y^+) = \arg \min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x', y') + L\|x' - x\|^2 - L\|y' - y\|^2.$$

2. Then we estimate z^+ by $\hat{z}^+$, which is obtained by solving the minimax problem:

$$\hat{z}^+ = (\hat{x}^+, \hat{y}^+) \approx \arg \min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} g(x', y'; x, y), \quad (8)$$

where $g(x', y'; x, y) \triangleq f(x', y') + L\|x' - x\|^2 - L\|y' - y\|^2$.

3. Finally, we construct $\hat{F}_{2L}(z) = 2L(z - \hat{z}^+)$ as the estimator of $F_{2L}(z)$.

We would like to regard $\hat{F}_{2L}(\cdot)$ as the stochastic estimator of $F_{2L}(\cdot)$, then it looks reasonable to run RAIN on the saddle envelope $f_{2L}(\cdot)$ by using its approximate gradient operator $\hat{F}_{2L}(\cdot)$. Since the function $g(x', y'; x, y)$ is strongly-convex in x' and strongly-concave in y' , it is desired to obtain $\hat{z}^+$ by Epoch-SEG efficiently. However, directly using Epoch-SEG on g only leads to $\mathbb{E}\|\hat{z}^+ - z^+\|^2$ be small, while the output $\hat{z}^+$ may be a biased estimator of z^+ . Consequently, the constructed $\hat{F}_{2L}(z)$ would also be biased, violating the assumption of unbiased stochastic oracle in RAIN. We address this issue by the following strategies:

1. We propose Epoch-SEG⁺ (Algorithm 4) by integrating the step of multilevel Monte-Carlo (MLMC) (Asi et al., 2021) with Epoch-SEG, which makes the bias of $\hat{z}^+$ decrease exponentially. The formal theoretical result is described in Theorem 5.1.
2. We show that RAIN also works for stochastic oracles with low bias. Since the basic component of RAIN is SEG, it is sufficient to analyze the complexity of SEG with low biased stochastic oracle. The formal theoretical result is described in Theorem 5.2.

Algorithm 4 Epoch-SEG⁺ ($f, z_0, \lambda, L, N, K, M$)

-
- 1: **for** $m = 0, 1, \dots, M-1$ **do**
 - 2: draw $J \sim \text{Geom}(1/2)$
 - 3: get $z_{m,N}, z_{m,N+J-1}, z_{m,N+J}$ by Epoch-SEG(f, z_0, λ, L, N, J)
 - 4: $\hat{z}_m \leftarrow z_{m,N} + 2^J(z_{m,N+J} - z_{m,N+J-1})\mathbb{I}[J \leq K]$
 - 5: **return** $\hat{z} \leftarrow \frac{1}{M} \sum_{m=0}^{M-1} \hat{z}_m$.
-

Theorem 5.1 Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave, and the stochastic first-order oracle $F(x, y; \xi)$ is unbiased estimator of $F(x, y)$ and has variance bounded with σ^2 , then Epoch-SEG⁺ (Algorithm 4) holds that

$$\|\mathbb{E}[\hat{z}^*] - z^*\|^2 \leq \frac{1}{2^{N+2K}} \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{2^K \lambda L}$$

and

$$\mathbb{E}\|\hat{z}^* - \mathbb{E}[\hat{z}^*]\|^2 \leq \frac{22}{2^N M} \mathbb{E}\|z_0 - z^*\|^2 + \frac{112K\sigma^2}{M\lambda L}.$$

Additionally, the total number of SFO calls is no more than $16\kappa MN + 64MKL/\lambda$ in expectation.

Theorem 5.2 (SEG with biased oracle) Suppose the function $f(x, y)$ is L -smooth and λ -strongly-convex- λ -strongly-concave. We run SEG (Algorithm 3) on $f(x, y)$ and denote

$$\begin{aligned} b_t &\triangleq \|\mathbb{E}F(z_t; \xi_i) - F(z_t)\|, & b_{t+1/2} &\triangleq \|\mathbb{E}F(z_{t+1/2}; \xi_j) - F(z_{t+1/2})\|, \\ \sigma_t^2 &\triangleq \mathbb{E}\|F(z_t; \xi_i) - F(z_t)\|^2, & \sigma_{t+1/2}^2 &\triangleq \mathbb{E}\|F(z_{t+1/2}; \xi_j) - F(z_{t+1/2})\|^2. \end{aligned}$$

Then it holds that

$$\begin{aligned} \mathbb{E}[\lambda\|z_{t+1/2} - z^*\|^2] &\leq \frac{1}{\eta} \mathbb{E}[\|z_t - z^*\|^2 - \|z_{t+1} - z^*\|^2] - \frac{1}{2\eta} \mathbb{E}\|z_{t+1} - z_{t+1/2}\|^2 \\ &\quad + \underbrace{6\eta(e_t^2 + e_{t+1/2}^2)}_{\text{mean square error}} + \underbrace{\frac{2b_{t+1/2}^2}{\lambda}}_{\text{bias}} \end{aligned}$$

for any $0 < \eta < 1/(4L)$, where $e_t^2 \triangleq b_t^2 + \sigma_t^2$ and $e_{t+1/2}^2 \triangleq b_{t+1/2}^2 + \sigma_{t+1/2}^2$.

The above theorem suggests that if it satisfies

$$\sigma_t^2 \leq \frac{\sigma^2}{2}, \quad \sigma_{t+1/2}^2 \leq \frac{\sigma^2}{2}, \quad b_t^2 \leq \frac{\sigma^2}{2}, \quad \text{and} \quad b_{t+1/2}^2 \leq 2\lambda\eta\sigma^2, \quad (9)$$

we are able to obtain the result of (5) in Lemma 4.1. Then we can follow the complexity analysis of RAIN to establish the theoretical guarantee of applying RAIN on the saddle envelope.¹ As a result, we can find an ϵ -stationary point of $f_{2L}(x, y)$:

1. We only need to replace every L with $6L$ in Theorem 4.1 and Theorem 4.2 since $f_{2L}(x, y)$ is $6L$ -smooth.

Algorithm 5 RAIN⁺⁺ ($f, z_0, \lambda, L, \{N_s\}_{s=0}^{S-1}, \{K_s\}_{s=0}^{S-1}$)

```

1:  $\lambda_0 = \lambda\gamma, \quad S = \lfloor \log_{(1+\gamma)}(6L/\lambda) \rfloor, \quad w_0 \leftarrow z_0$ 
2: for  $s = 0, 1, \dots, S-1$ 
3:    $z_{s,0} \leftarrow z_s, \quad w_{s,0} \leftarrow w_s$ 
4:   for  $k = 0, 1, \dots, N_s - 1$  do
5:      $(z_{s,k+1}, w_{s,k+1}) \leftarrow \text{SEG}^{++}(f, z_{s,k}, w_{s,k}, \frac{1}{48L}, \frac{96L}{\lambda_s}, s, \{z_i\}_{i=0}^s, \gamma, \lambda)$ 
6:   for  $k = N_s, N_s + 1, \dots, N_s + K_s + 1$  do
7:      $(z_{s,k+1}, w_{s,k+1}) \leftarrow \text{SEG}^{++}(f, z_{s,k}, w_{s,k}, \frac{1}{2^{k-N_s+3} \times 12L}, \frac{2^{k-N_s+5} \times 12L}{\lambda_s}, s, \{z_i\}_{i=0}^s, \gamma, \lambda)$ 
8:   end for
9:    $(z_{s+1}, w_{s+1}) \leftarrow (z_{N_s+K_s}, w_{N_s+K_s})$ 
10:   $\lambda_{s+1} \leftarrow (1 + \gamma)\lambda_s$ 
11: end for
12: return  $z_S$ 
    
```

- For (τ, α) -intersection-dominant condition with $\tau \geq 2L$, it requires $\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + L/\alpha)$ times evaluations of $\hat{F}_{2L}(\cdot)$.
- For $-\rho$ -comonotone condition with $\rho \leq 1/(2L)$, it requires $\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + L/\epsilon)$ times evaluations of $\hat{F}_{2L}(\cdot)$.

Based on the above ideas, we proposed RAIN⁺⁺ in Algorithm 5 as an extension of RAIN. The iterations of RAIN⁺⁺ consider the regularized function

$$f_{2L}^{(s)}(x, y) \triangleq \begin{cases} f_{2L}(x, y), & s = 0, \\ f_{2L}^{(s-1)}(x, y) + \frac{\lambda_s}{2}\|x - x_s\|^2 - \frac{\lambda_s}{2}\|y - y_s\|^2, & s \geq 1, \end{cases} \quad (10)$$

Compared with the iterations of RAIN applying SEG to solve the sub-problem

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

the iterations of RAIN⁺⁺ apply SEG⁺⁺ to solve the sub-problem

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f_{2L}^{(s)}(x, y).$$

For SEG⁺⁺, we apply Epoch-SEG⁺ (Epoch-SEG with the debias step of MLMC) to achieve the low-biased gradient estimation of the saddle envelope $f_{2L}^{(s)}(x, y)$. Note that the sub-problems solved by Epoch-SEG⁺ (Line 5 and 9 of Algorithm 6) are well-conditioned because they are $3L$ -smooth and L -strongly-convex- L -strongly-concave. Therefore the SFO complexity for Epoch-SEG⁺ only require the complexity of $\tilde{\mathcal{O}}(1)$ to obtain a sufficiently accurate solution by carefully choosing the initialization and parameters. As a result, the total SFO complexity of RAIN⁺⁺ would be $\tilde{\mathcal{O}}(1)$ times the one of RAIN.

Below, we present the theoretical result under the intersection-dominant condition.

Theorem 5.3 (RAIN⁺⁺, ID) *Suppose the function $f(x, y)$ is L -smooth and satisfies the (τ, α) -intersection-dominant condition with $\tau \geq 2L$, and the stochastic first-order oracle $F(x, y; \xi)$ is unbiased and has variance bounded by σ^2 , then running RAIN⁺⁺ with $\gamma = 1$, λ as defined in (7) and*

$$N_s = \begin{cases} \left\lceil \log_2 \left(\frac{512\lambda^2 S^2 D^2}{\epsilon^2} \right) \right\rceil, & s = 0, \\ 3, & s \geq 1, \end{cases} \quad \text{and} \quad K_s = \left\lceil \log_2 \left(\frac{1024\lambda_s S^2 \sigma^2}{3L\epsilon^2} \right) \right\rceil$$

holds that $\mathbb{E}\|F_{2L}(z_S)\| \leq \epsilon$, where λ is defined in (7). Additionally, the total number of SFO calls is no more than $\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + L/\alpha)$ in expectation.

We can also use the regularization trick to obtain the theoretical guarantee in the negative comonotone setting. The idea is applying RAIN⁺⁺ on

$$g_{2L}(x, y) \triangleq f_{2L}(x, y) + \frac{\lambda}{2}\|x - x_0\|^2 - \frac{\lambda}{2}\|y - y_0\|^2$$

for some $\lambda = \Theta(\epsilon/D)$. Note that the function $g_{2L}(x, y)$ is λ -strongly-convex- λ -strongly-concave and the algorithm solves the minimax sub-problems with objective functions

$$f_{2L}^{(s)}(x, y) \triangleq \begin{cases} f_{2L}(x, y) + \frac{\lambda}{2}\|x - x_0\|^2 - \frac{\lambda}{2}\|y - y_0\|^2, & s = 0, \\ f_{2L}^{(s-1)}(x, y) + \frac{\lambda_s}{2}\|x - x_s\|^2 - \frac{\lambda_s}{2}\|y - y_s\|^2, & s \geq 1, \end{cases} \quad (11)$$

which is similar to what we have done in the general convex-concave case. We formally present the result in the following theorem.

Theorem 5.4 (RAIN⁺⁺, NC) *Suppose the function $f(x, y)$ is L -smooth and its gradient operator is $-\rho$ -comonotone with $\rho \leq 1/(2L)$, and the stochastic first-order oracle $F(x, y; \xi)$ is unbiased and has variance bounded by σ^2 . If we run RAIN⁺⁺ (Algorithm 5) with $\gamma = 1$ and $\lambda = \min\{\epsilon/D, 6L\}$ then it holds that $\mathbb{E}\|F_{2L}(z_S)\| \leq 3\epsilon$. Additionally, the total number of SFO calls is no more than $\tilde{\mathcal{O}}(\sigma^2\epsilon^{-2} + LD\epsilon^{-1})$.*

Corollary 5.3 and Corollary 5.4 indicate RAIN⁺⁺ can find an ϵ -stationary point of the envelope $f_{2L}(x, y)$ in corresponding settings, which easily leads to a nearly-stationary point of $f(x, y)$ by the following proposition.

Proposition 5.3 *Suppose the function $f(x, y)$ is L -smooth and the point $\hat{z} = (\hat{x}, \hat{y})$ is an ϵ -stationary point of the function $f_{2L}(x, y)$, then we can find a 2ϵ -stationary point of $f(x, y)$ within $\mathcal{O}(\log(\epsilon^{-1}) + \sigma^2\epsilon^{-2})$ SFO complexity in expectation.*

6. The Lower Complexity Bounds

We first provide the lower complexity bounds for finding near-stationary points of the stochastic convex-concave minimax problem. Combining the ideas of Luo et al. (2021); Foster et al. (2019) and Yoon and Ryu (2021), we obtain Theorem 6.2 and Theorem 6.1 for general convex-concave case and strongly-convex-strongly-concave case respectively, which nearly match the upper bounds shown in Theorem 4.1 and Theorem 4.2. Hence, the proposed RAIN (Algorithm 1) is near-optimal.

Algorithm 6 $\text{SEG}^{++}(f, z_0, w_0, \eta, T, s, \{\tilde{z}_i\}_{i=0}^s, \gamma, \lambda)$

```

1: let  $L$  be the smoothness coefficient of  $f(x, y)$ 
2: set  $N, K, M$  according to equation (24) and  $\hat{z}_{-1/2}^+ \leftarrow w_0$ 
3: for  $t = 0, 1, \dots, T - 1$  do
4:    $g_t(x, y) \triangleq f(x, y) + L\|x - x_t\|^2 - L\|y - y_t\|^2$ 
5:    $\hat{z}_t^+ \leftarrow \text{Epoch-SEG}^+(g_t, \hat{z}_{t-1/2}^+, L, 3L, N, K, M)$ 
6:    $\hat{F}_{2L}^{(s)}(z_t) \leftarrow \begin{cases} 2L(z_t - \hat{z}_t^+) + \lambda\gamma \sum_{i=1}^s (1 + \gamma)^i (z_t - \tilde{z}_i), & \text{ID case} \\ 2L(z_t - \hat{z}_t^+) + \lambda\gamma \sum_{i=0}^s (1 + \gamma)^i (z_t - \tilde{z}_i), & \text{NC case} \end{cases}$ 
7:    $z_{t+1/2} \leftarrow z_t - \eta \hat{F}_{2L}^{(s)}(z_t)$ 
8:    $g_{t+1/2}(x, y) \triangleq f(x, y) + L\|x - x_{t+1/2}\|^2 - L\|y - y_{t+1/2}\|^2$ 
9:    $\hat{z}_{t+1/2}^+ \leftarrow \text{Epoch-SEG}^+(g_{t+1/2}, \hat{z}_t^+, L, 3L, N, K, M)$ 
10:   $\hat{F}_{2L}^{(s)}(z_{t+1/2}) \leftarrow \begin{cases} 2L(z_{t+1/2} - \hat{z}_{t+1/2}^+) + \lambda\gamma \sum_{i=1}^s (1 + \gamma)^i (z_{t+1/2} - \tilde{z}_i), & \text{ID case} \\ 2L(z_{t+1/2} - \hat{z}_{t+1/2}^+) + \lambda\gamma \sum_{i=0}^s (1 + \gamma)^i (z_{t+1/2} - \tilde{z}_i), & \text{NC case} \end{cases}$ 
11:   $z_{t+1} \leftarrow z_t - \eta \hat{F}_{2L}^{(s)}(z_{t+1/2})$ 
12: end for
13: Draw  $J \sim \text{Unif}([T])$ 
14: return  $(z_{J+1/2}, \hat{z}_{J+1/2}^+)$ 

```

Theorem 6.1 *For any stochastic algorithm $\mathcal{A}$ based on stochastic first-order oracle (SFO) under Assumption 2.6 and parameters $L \geq 2, \lambda \leq 1$ and $\epsilon \leq 0.01\lambda$, there exist an L -smooth and λ -strongly-convex- λ -strongly-concave function $f(x, y)$ that $\mathcal{A}$ needs at least*

$$\Omega(\sigma^2 \epsilon^{-2} \log(L\epsilon^{-1}) + \kappa \log(\lambda \epsilon^{-1}))$$

SFO calls to find an ϵ -stationary point of $f(x, y)$.

Theorem 6.2 *For any stochastic algorithm $\mathcal{A}$ based on stochastic first-order oracle (SFO) under Assumption 2.6 and parameters of $L \geq 2$ and $\epsilon \leq 0.01$, there exists an L -smooth and convex-concave function $f(x, y)$ such that $\mathcal{A}$ needs at least*

$$\Omega(\sigma^2 \epsilon^{-2} \log(L\epsilon^{-1}) + LD\epsilon^{-1})$$

SFO calls to find an ϵ -stationary point of $f(x, y)$.

Since the gradient operator of any convex-concave function is negative comonotone, the lower bound shown in Theorem 6.1 is also valid for the problem under the negative comonotone condition.

Corollary 6.1 *For any stochastic algorithm $\mathcal{A}$ based on stochastic first-order oracle (SFO) under Assumption 2.6 and parameters of $L \geq 2$ and $\epsilon \leq 0.01$, there exist an L -smooth function $f(x, y)$ whose gradient operator is negative comonotone such that $\mathcal{A}$ needs at least*

$$\Omega(\sigma^2 \epsilon^{-2} \log(L\epsilon^{-1}) + LD\epsilon^{-1})$$

SFO calls to find an ϵ -stationary point of $f(x, y)$.





Figure 1: The results of the number of SFO calls against gradient norm on problem (12).

Figure 2: The results of the number of SFO calls against gradient norm on problem (13). SEAG diverges in (c), which does not contradict its convergence guarantee as the condition $\sigma_k^2 \leq \epsilon/(k+1)$ in Theorem 6.1 (Lee and Kim, 2021) is unsatisfied.

Similarly, Theorem 6.2 leads to the lower bound for intersection-dominant condition.

Corollary 6.2 *For any stochastic algorithm $\mathcal{A}$ based on stochastic first-order oracle (SFO) under Assumption 2.6 and parameters of $L \geq 2$, $\alpha \leq 1$ and $\epsilon \leq 0.01\alpha$, there exist an L -smooth function $f(x, y)$ satisfying the (τ, α) -intersection-dominant condition with $\tau \geq 2L$ such that $\mathcal{A}$ needs at least*

$$\Omega(\sigma^2 \epsilon^{-2} \log(L\epsilon^{-1}) + L\alpha^{-1} \log(\alpha\epsilon^{-1}))$$

SFO calls to find an ϵ -stationary point of $f(x, y)$.

Corollary 6.1 and Corollary 6.2 suggest that the proposed RAIN⁺⁺ (Algorithm 5) is also near-optimal in negative comonotone and intersection-dominant conditions respectively.

7. Numerical Experiments

In this section, we compare our algorithms with baselines on both convex-concave and nonconvex-nonconcave minimax problems.

7.1 The Convex-Concave Case

We conduct numerical experiments by comparing our RAIN with the following baselines:

- SEG: the ordinary stochastic extragradient;
- R-SEG: the stochastic extragradient with regularization trick (Nesterov, 2012);
- SEAG: the stochastic extra anchored gradient (Lee and Kim, 2021, Algorithm 3);
- PDHG: the primal-dual hybrid gradient (Zhao, 2022, Algorithm 1).

The experiments consider two minimax problems as follows.

- The first one is the bilinear minimax problem:

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^d} f_b(x, y) \triangleq x^\top y, \quad (12)$$

which reveals some important issues in minimax optimization. For example, the duality gap is not well defined except at its unique saddle point $(0, 0)$; the classical method stochastic gradient descent ascent diverges and other first-order algorithms also converge slowly due to the cycling behaviors (Pethick et al., 2022).

- The second one is the hard case of convex-concave minimax problem

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^d} f_{\delta, \nu}(x, y) \triangleq (1 - \delta)g_\nu(x) + \delta x^\top y - (1 - \delta)g_\nu(y), \quad (13)$$

where

$$g_\nu(u_i) = \begin{cases} \nu|u_i| - \frac{1}{2}\nu^2, & |u_i| \geq \nu, \\ \frac{1}{2}u_i^2, & |u_i| < \nu, \end{cases}$$

and we set $\nu = 5 \times 10^{-5}$ and $\delta = 10^{-2}$ by following Yoon and Ryu (2021)’s setting.

We use the stochastic first-order oracles

$$F_b(x, y; \xi) = F_b(x, y) + \xi \quad \text{and} \quad F_{\delta, \nu}(x, y; \xi) = F_{\delta, \nu}(x, y) + \xi$$

for problems (12) and (13) respectively, where $F_b(x, y)$ is the gradient operator of $f_b(x, y)$, $F_{\delta, \nu}(x, y)$ is the gradient operator of $f_{\delta, \nu}(x, y)$ and $\xi \sim \mathcal{N}(0, \sigma^2 I_{2d})$. We set $d = 1000$ for problem (12) and $d = 100$ for problem (13). We provide the detailed implementation for the algorithms in Appendix E and the source code is available². We present the experimental results under different levels of noise in Figure 1 and Figure 2, which shows the proposed RAIN obviously performs better than baseline methods.

2. <https://github.com/TrueNobility303/RAIN>




Figure 3: The results of the number of SFO calls against gradient norm on problem (14) with $L = 1$ and $\rho = -1/(8\sqrt{2})$.

7.2 The Nonconvex-Nonconcave Case

We also consider the nonconvex-nonconcave problem

$$\min_{x \in \mathbb{R}} \max_{y \in \mathbb{R}} f_{\rho, L}(x, y) \triangleq \frac{\rho L^2}{2} x^2 + L \sqrt{1 - \rho^2 L^2} xy - \frac{\rho L^2}{2} y^2, \quad (14)$$

where $L, \rho > 0$ such that $\rho L < 1$. According to the verification of Lee and Kim (2021), the gradient operator of function $f_{\rho, L}(x, y)$ is L -Lipschitz continuous and ρ -comonotone. We use the stochastic first-order oracles

$$F_{\rho, L}(x, y; \xi) = f_{\rho, L}(x, y) + \xi$$

in our experiments, where $\xi \sim \mathcal{N}(0, \sigma^2 I_2)$.

We compare our proposed RAIN⁺⁺ with the following baselines:

- SEG⁺: the extension of SEG under the negative comonotonicity assumption (Diakonikolas et al., 2021, Equation (EG_p⁺));
- SFEG: the stochastic fast extra gradient method (Lee and Kim, 2021, Algorithm 1), where we replace the exact gradient operator with the stochastic gradient operator.

We test the algorithms on Problem (14) with $L = 1$ and $\rho = -1/8\sqrt{2}$ or $-1/3$, and the experimental results are shown in Figure 3 and 4, respectively. The implementation details of the algorithm are deferred to Appendix E. We can observe that our proposed RAIN⁺⁺ performs better than baselines. We remark that the convergence analysis of SFEG by Lee and Kim (2021) requires either $\rho = 0$ or $\sigma = 0$, while it also works on our stochastic nonconvex-nonconcave problem with $\rho > 0$ and $\sigma > 0$ in practice. Note that Figure 4 shows that SEG⁺ diverges on Problem (14) with $\rho = -1/3$, which is also observed by the empirical results of Lee and Kim (2021). The reason is the convergence guarantee of SEG⁺ (Diakonikolas et al., 2021, Theorem 4.5)) requires the condition $\rho \in [-1/(4\sqrt{2}L), 0)$, which is not satisfied in the setting of $\rho = -1/3$.




Figure 4: The results of the number of SFO calls against gradient norm on problem (14) with $L = 1$ and $\rho = -1/3$.

8. Conclusion

In this work, we propose the Recursive Anchor IterationN (RAIN) algorithm for stochastic minimax optimization. The theoretical analysis has shown that the framework of RAIN with appropriate sub-problem solvers could achieve the near-optimal SFO complexity for finding nearly stationary points in convex-concave and strongly-convex-strongly-concave minimax optimization problems. We also extend the idea of RAIN to solve two specific nonconvex-nonconcave minimax problems and the proposed method RAIN^{++} also achieves the near-optimal SFO complexity in these settings.

Acknowledgments

This work was done while Lesi Chen was at Fudan University. This work is supported by National Natural Science Foundation of China (No. 62206058), Shanghai Sailing Program (22YF1402900), Shanghai Basic Research Program (23JC1401000), and the Major Key Project of Pengcheng Laboratory (No. PCL2024A06)

Appendix A. The Proofs in Section 3

We first provide the non-expansiveness lemma, then give detailed proofs for results in Section 3.

Lemma A.1 (non-expansiveness) *Suppose the operator $F(\cdot)$ is monotone. We define $G(z) \triangleq F(z) + \lambda(z - z_0)$ for some $\lambda > 0$ and $z_0 \in \mathbb{R}^d$, then it holds that*

$$\|w^* - z_0\| \leq \|z^* - z_0\| \quad \text{and} \quad \|w^* - z^*\| \leq \|z^* - z_0\|,$$

where z^* is the solution to $F(z) = 0$ and w^* is the solution to $G(w) = 0$.

Proof The monotonicity of F means G is λ -strongly monotone. Therefore, the operator G must has an unique solution w^* . Then we have

$$\begin{aligned} & \lambda \|w^* - z^*\|^2 \\ & \leq G(z^*)^\top (z^* - w^*) \\ & = (F(z^*) + \lambda(z^* - z_0))^\top (z^* - w^*) \\ & = \lambda(z^* - z_0)^\top (z^* - w^*) \\ & = \frac{\lambda}{2} \|w^* - z^*\|^2 + \frac{\lambda}{2} \|z^* - z_0\|^2 - \frac{\lambda}{2} \|w^* - z_0\|^2, \end{aligned}$$

which implies the result of this lemma. ■

A.1 The Proof of Lemma 3.1

Proof Let $F^{(s)}(\cdot)$ be the gradient operator of $f^{(s)}(x, y)$. It is clear that $F^{(s)}(\cdot)$ is equivalent to the following definition:

$$F^{(s)}(z) \triangleq F(z) + \lambda \sum_{i=1}^s 2^i (z - z_i), \quad s = 0, \dots, S-1.$$

Based on the setting of $S = \lfloor \log_2(L/\lambda) \rfloor$, we have

$$\lambda 2^S \leq L \leq \lambda 2^{S+1},$$

then it holds that

$$\lambda + \lambda \sum_{i=1}^s 2^i > 2^s \lambda$$

and

$$L + \lambda \sum_{i=1}^{S-1} 2^i \leq 2L.$$

Hence, we know that each $F^{(s)}$ is at least $2^s\lambda$ -strongly monotone and $2L$ -smooth. Further, we have

$$\begin{aligned}
 \|F(z_S)\| &\leq \|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_S - z_i\| \\
 &\leq \|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_S - z_{S-1}^*\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{S-1}^* - z_i\| \\
 &\leq 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{S-1}^* - z_i\| \\
 &\leq 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{S-1}^* - z_{i-1}^*\| \\
 &\leq 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| + \lambda \sum_{i=1}^{S-1} 2^i \sum_{j=i}^{S-1} \|z_j^* - z_{j-1}^*\| \\
 &= 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| + \lambda \sum_{j=1}^{S-1} \|z_j^* - z_{j-1}^*\| \sum_{i=1}^j 2^i \\
 &\leq 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| + \lambda \sum_{j=1}^{S-1} 2^{j+1} \|z_j^* - z_{j-1}^*\| \\
 &= 2\|F^{(S-1)}(z_S)\| + \lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| + \lambda \sum_{i=1}^{S-1} 2^{i+1} \|z_i^* - z_{i-1}^*\| \\
 &\leq 4L\|z_S - z_{S-1}^*\| + 3\lambda \sum_{i=1}^{S-1} 2^i \|z_{i-1}^* - z_i\| \\
 &\leq 16\lambda \sum_{i=1}^S 2^{i-1} \|z_{i-1}^* - z_i\|.
 \end{aligned}$$

Above, the third line follows from $F^{(S-1)}$ is at least $(\lambda \sum_{i=1}^{S-1} 2^i)$ -strongly monotone; the second last line relies on the non-expansiveness after anchoring shown in Lemma A.1 such that $\|z_i^* - z_{i-1}^*\| \leq \|z_{i-1}^* - z_i\|$; in the last line we use $F^{(S-1)}$ is at most $2L$ -Lipschitz and $L \leq \lambda 2^{S+1}$; the other steps only requires triangle inequality and simple calculation. $\blacksquare$

A.2 The Proof of Lemma 3.2

Proof Let $z_g^* = (x_g^*, y_g^*)$ be the unique stationary point of $g(x, y)$ such that $G(z_g^*) = 0$. Then we have

$$\begin{aligned}
 \|F(\tilde{z})\| &\leq \|G(\tilde{z})\| + \lambda \|\tilde{z} - z_0\| \\
 &\leq \|G(\tilde{z})\| + \lambda \|\tilde{z} - z_g^*\| + \lambda \|z_g^* - z_0\| \\
 &\leq 2\|G(\tilde{z})\| + \lambda \|z_g^* - z_0\|,
 \end{aligned}$$

where the last step we use G is λ -strongly-monotone and $\|z_g^* - z_0\| \leq \|z^* - z_0\|$ from Lemma A.1. $\blacksquare$

Appendix B. The Proofs in Section 4

We provide detailed proofs for results in Section 4 except Lemma 4.1. Note that Lemma 4.1 is a special case of Theorem 5.2 with $b_t = b_{t+1/2} = 0$ and $\sigma_t = \sigma_{t+1/2} = \sigma$. Hence, it can be proved by directly using the analysis for Theorem 5.2 in Appendix C.4.

B.1 The Proof of Lemma 4.2

Proof Denote η_k and T_k be the step size and epoch length in the k_{th} epoch. From Lemma 4.1 and strong monotone we know that

$$\mathbb{E}\|z_{k+1} - z^*\|^2 \leq \frac{1}{\lambda\eta_k T_k} \mathbb{E}\|z_k - z^*\|^2 + \frac{16\eta_k \sigma^2}{\lambda}. \quad (15)$$

Telescoping from $k = 0, 1, \dots, N-1$, we obtain

$$\mathbb{E}\|z_N - z^*\|^2 \leq \frac{1}{2^N} \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{\lambda L}.$$

Then we can use induction from $N, N+1, \dots, N+K-1$ to show inequality in this theorem: suppose it is true for the case $N+k$, then for the case $N+k+1$, we have

$$\begin{aligned} & \mathbb{E}\|z_{N+k+1} - z^*\|^2 \\ & \leq \frac{1}{4} \mathbb{E}\|z_{N+k} - z^*\|^2 + \frac{2\sigma^2}{2^k \lambda L} \\ & \leq \frac{1}{4} \left(\frac{1}{2^{N+2k}} \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{2^k \lambda L} \right) + \frac{2\sigma^2}{2^k \lambda L} \\ & \leq \frac{1}{2^{N+2(k+1)}} \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{2^{k+1} \lambda L}, \end{aligned}$$

where the first inequality is by plugging η_k, T_k into (15) and the second one comes from the induction hypothesis.

In addition, as a side effect, it is clear that in this procedure, for any $k = 0, 1, \dots, N+K-1$, we can always maintain the bound of

$$\mathbb{E}\|z_k - z^*\|^2 \leq \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{\lambda L}. \quad (16)$$

For each iteration of SEG, we need to quest the stochastic operator twice. Then summing up all the iterations in every epoch yields the complexity as claimed. $\blacksquare$

B.2 The Proof of Theorem 4.1

Proof By Lemma 3.1 and taking expectation, we know that to find z_S such that $\mathbb{E}\|F(z_S)\| \leq \epsilon$ it is sufficient to guarantee

$$256\lambda_s^2 S^2 \mathbb{E}\|z_{s+1} - z_s^*\|^2 \leq \epsilon^2, \quad (17)$$

for $s = 0, 1, \dots, S-1$, where z_s^* denotes the solution to

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f_s(x, y).$$

Then we prove (17) holds for $s = 0, 1, \dots, S-1$ by induction. By the definition of λ_s , we know that $\lambda_{s+1} = 2\lambda_s$ and each sub-problem $f_s(x, y)$ is at least λ_s -strongly-convex- λ_s -strongly-concave. Now we specify the parameters in Epoch-SEG by

$$z_{s+1} \leftarrow \text{Epoch-SEG}(f_s, z_s, \lambda_s, 2L, N_s, K_s).$$

By results of (15), we know that

$$\begin{aligned} & \mathbb{E}\|z_{s+1} - z_s^*\|^2 \\ & \leq \frac{1}{2^{N_s+2K_s}} \mathbb{E}\|z_s - z_s^*\|^2 + \frac{4\sigma^2}{2^{K_s}\lambda_s L} \\ & \leq \frac{1}{2^{N_s+2K_s}} \mathbb{E}\|z_s - z_{s-1}^*\|^2 + \frac{4\sigma^2}{2^{K_s}\lambda_s L}, \end{aligned}$$

where we use Lemma 3.2 to obtain $\|z_s - z_s^*\| \leq \|z_s - z_{s-1}^*\|$.

For the case of $s = 0$, we have

$$\mathbb{E}\|z_1 - z_0^*\|^2 \leq \frac{1}{2^{N_0}} \mathbb{E}\|z_0 - z^*\|^2 + \frac{4\sigma^2}{2^{K_0}\lambda L},$$

where we use $w_0^* = z^*$, $\lambda_0 = \lambda$ and $K_0 \geq 1$. Then we can let

$$N_0 \geq \log_2 \left(\frac{512\lambda^2 S^2 D^2}{\epsilon^2} \right) \quad \text{and} \quad 2^{K_0} \geq \frac{2048\lambda S^2 \sigma^2}{L\epsilon^2}$$

to guarantee $\mathbb{E}\|z_1 - z_0^*\|^2$ meet our requirement.

Suppose we already have $256\lambda_{s-1}^2 S^2 \mathbb{E}\|z_s - z_{s-1}^*\|^2 \leq \epsilon^2$. By observing that

$$\begin{aligned} \mathbb{E}[\|z_{s+1} - z_s^*\|^2] & \leq \frac{1}{2^{N_s+2K_s}} \mathbb{E}[\|z_s - z_s^*\|^2] + \frac{4\sigma^2}{2^{K_s}\lambda_s L} \\ & \leq \frac{1}{2^{N_s+2K_s}} \mathbb{E}[\|z_s - z_{s-1}^*\|^2] + \frac{4\sigma^2}{2^{K_s}\lambda_s L} \\ & \leq \frac{1}{2^{K_s}} \times \frac{\epsilon^2}{256\lambda_{s-1}^2 S^2} + \frac{4\sigma^2}{2^{N_s}\lambda_s L}, \end{aligned}$$

then let

$$N_s \geq 3 \quad \text{and} \quad 2^{K_s} \geq \frac{2048\lambda_s S^2 \sigma^2}{L\epsilon^2}$$

for all $s \geq 1$ leads to

$$256\lambda_s^2 S^2 \mathbb{E} \|z_{s+1} - z_s^*\|^2 \leq \epsilon^2.$$

The total SFO complexity is

$$\begin{aligned} & \sum_{s=0}^{S-1} \frac{2L}{\lambda_s} \times (16N_s + 64 \times 2^{K_s}) \\ &= \frac{2LN_0}{\lambda} + \sum_{s=1}^{S-1} \frac{96L}{\lambda_s} + \sum_{s=0}^{S-1} \frac{128L}{\lambda_s} \times 2^{K_s} \\ &\leq \frac{2LN_0}{\lambda} + \sum_{s=1}^{S-1} \frac{96L}{\lambda_s} + \sum_{s=0}^{S-1} \frac{512L}{\lambda_s} \times \frac{2048\lambda_s S^2 \sigma^2}{L\epsilon^2} \\ &\leq \frac{2LN_0}{\lambda} + \frac{96L}{\lambda} + \frac{1048576S^3 \sigma^2}{\epsilon^2}. \end{aligned}$$

Plugging the value of N_0 and S into above equation completes the proof. $\blacksquare$

B.3 The Proof of Theorem 4.2

Proof Let w be the output of applying RAIN on the function $g(x, y)$ and we define the operator $G_0(z) \triangleq F(z) + \lambda(z - z_0)$. By the anchoring lemma with any $0 < \lambda \leq \epsilon/D$ and taking expectation, we know that if we can make sure $\mathbb{E}\|G_0(w)\| \leq \epsilon$, then we can obtain $\mathbb{E}\|F(w)\| \leq 3\epsilon$. By Theorem 4.1, we can ensure $\mathbb{E}\|G(w)\| \leq \epsilon$ within a SFO complexity of

$$\mathcal{O}\left(\frac{L+\lambda}{\lambda} + \frac{L+\lambda}{\lambda} \log\left(\frac{\lambda D}{\epsilon} \log\left(\frac{L+\lambda}{\lambda}\right)\right) + \frac{\sigma^2}{\epsilon^2} \log^3\left(\frac{L+\lambda}{\lambda}\right)\right).$$

Then we complete the proof by plugging the value of λ into above equation. $\blacksquare$

Appendix C. The Proofs in Section 5

We provide the detailed proofs for results in Section 5.

C.1 The Proof of Proposition 5.1

Proof Since $\tau \geq 2L$, then we know

$$\nabla_{xx}^2 f(x, y) + \nabla_{xy}^2 f(x, y)(2LI - \nabla_{yy}^2 f(x, y))^{-1} \nabla_{yx}^2 f(x, y) \succeq \alpha I$$

and

$$-\nabla_{yy}^2 f(x, y) + \nabla_{yx}^2 f(x, y)(2LI + \nabla_{xx}^2 f(x, y))^{-1} \nabla_{xy}^2 f(x, y) \succeq \alpha I,$$

which implies the desired result by Proposition 2.6 of Grimmer et al. (2023). $\blacksquare$

C.2 The Proof of Proposition 5.2

Proof According to the proof of Example 1 from Lee and Kim (2021), we know that

$$(F(z) - F(z'))^\top (z - z') \geq -\frac{1}{2L} \|F(z) - F(z')\|^2$$

for any $z = (x, y)$ and $z' = (x', y')$, which implies the gradient operator of $f_{2L}(x, y)$ is monotone. It further indicates the saddle envelope $f_{2L}(x, y)$ is convex-concave (Liu et al., 2021, Lemma 3). $\blacksquare$

C.3 The Proof of Theorem 5.1

Proof We know that

$$\mathbb{P}(J = j) = 2^{-j}, \quad j = 1, \dots, K.$$

Since $\hat{z}_m$ is i.i.d for all m , we have

$$\mathbb{E}[\hat{z}] = \mathbb{E}[\hat{z}_0] = \mathbb{E}[z_N] + \sum_{j=1}^K \mathbb{P}(J = j) 2^j \mathbb{E}[z_{N+j} - z_{N+j-1}] = \mathbb{E}[z_{N+K}].$$

Using Lemma 4.2 and Jensen's inequality, we obtain the upper bound for bias:

$$\|\mathbb{E}\hat{z} - z^*\|^2 \leq \mathbb{E}\|z_{N+K} - z^*\|^2 \leq \frac{1}{2^{N+2K}} \times \mathbb{E}\|z_0 - z^*\|^2 + \frac{8\sigma^2}{2^K \lambda L};$$

as well as the upper bound for variance:

$$\begin{aligned} & \mathbb{E}\|\hat{z} - \mathbb{E}\hat{z}\|^2 \\ & \leq \frac{1}{M} \mathbb{E}\|\hat{z}_0 - \mathbb{E}\hat{z}_0\|^2 \\ & \leq \frac{1}{M} \mathbb{E}\|\hat{z}_0 - z^*\|^2 \\ & \leq \frac{2}{M} \mathbb{E}\|z_N - z^*\|^2 + \frac{2}{M} \mathbb{E}\|2^J (z_{N+J} - z_{N+J-1})\|^2 \\ & = \frac{2}{M} \mathbb{E}\|z_N - z^*\|^2 + \frac{2}{M} \sum_{j=1}^K \mathbb{P}(J = j) 2^j \|z_{N+j} - z_{N+j-1}\|^2 \\ & \leq \frac{2}{M} \mathbb{E}\|z_N - z^*\|^2 + \frac{4}{M} \sum_{j=1}^K 2^j (\|z_{N+j} - z^*\|^2 + \|z_{N+j-1} - z^*\|^2) \\ & \leq \frac{2}{M} \mathbb{E}\|z_N - z^*\|^2 + \frac{1}{M} \sum_{j=1}^K \left\{ \frac{1}{2^{N+j}} \times 20 \mathbb{E}\|z_0 - z^*\|^2 + \frac{96\sigma^2}{\lambda L} \right\} \\ & \leq \frac{1}{M} \left\{ \frac{2}{2^N} \mathbb{E}\|z_0 - z^*\|^2 + \frac{16\sigma^2}{\lambda L} + \frac{20}{2^N} \mathbb{E}\|z_0 - z^*\|^2 + \frac{96K\sigma^2}{\lambda L} \right\} \\ & \leq \frac{22}{2^N M} \mathbb{E}[\|z_0 - z^*\|^2] + \frac{112K\sigma^2}{M\lambda L}, \end{aligned}$$

where the second inequality follows from the fact that variance is always bounded by mean square error, i.e. it holds that

$$\begin{aligned}
& \mathbb{E}\|\hat{z}_0 - \mathbb{E}\hat{z}_0\|^2 \\
&= \mathbb{E}\|\hat{z}_0\|^2 - \|\mathbb{E}\hat{z}_0\|^2 \\
&\leq \mathbb{E}\|\hat{z}_0\|^2 - 2\langle \mathbb{E}\hat{z}_0, z^* \rangle + \|z^*\|^2 \\
&= \mathbb{E}\|\hat{z}_0 - z^*\|^2;
\end{aligned}$$

where the third inequality is from the definition of $\hat{z}_m$; the second last, third last and fourth last ones are all by Lemma 4.2; others are dependent on the Young's inequality or simple algebra. The complexity in expectation can be derived by Lemma 4.2, which is

$$\begin{aligned}
& M \times \mathbb{E}[16\kappa N + 64\kappa 2^J] \\
&= 16\kappa MN + 64\kappa M \sum_{j=1}^K \mathbb{P}(J = j) 2^j \\
&= 16\kappa MN + 64\kappa MK,
\end{aligned}$$

where $\kappa = L/\lambda$. ■

C.4 The Proof of Theorem 5.2

Proof We begin from the strong monotonicity and an identity:

$$\begin{aligned}
& 2\mathbb{E}[F(z_{t+1/2})^\top (z_{t+1/2} - z^*)] \\
&= 2\mathbb{E}[(F(z_{t+1/2}) - F(z_{t+1/2}; \xi_j))^\top (z_{t+1/2} - z^*) + F(z_{t+1/2}; \xi_j)^\top (z_{t+1/2} - z^*)] \\
&+ 2\mathbb{E}[(F(z_{t+1/2}; \xi_j) - F(z_t; \xi_i))^\top (z_{t+1/2} - z_{t+1}) + F(z_t; \xi_i)^\top (z_{t+1/2} - z_{t+1})].
\end{aligned} \tag{18}$$

Note that we have

$$\begin{aligned}
& 2\mathbb{E}[(F(z_{t+1/2}) - F(z_{t+1/2}; \xi_j))^\top (z_{t+1/2} - z^*)] \\
&= 2\mathbb{E}[F(z_{t+1/2}) - F(z_{t+1/2}; \xi_j)]^\top \mathbb{E}[z_{t+1/2} - z^*] \\
&\leq \frac{\lambda}{2} \mathbb{E}[\|z_{t+1/2} - z^*\|^2] + \frac{2b_{t+1/2}^2}{\lambda},
\end{aligned} \tag{19}$$

where we use the independence and Young's inequality. Also, we have

$$\begin{aligned}
& 2\mathbb{E}[F(z_{t+1/2}; \xi_j)^\top (z_{t+1} - z^*)] \\
&= \frac{2}{\eta} \mathbb{E}[(z_t - z_{t+1})^\top (z_{t+1} - z^*)] \\
&= \frac{1}{\eta} \mathbb{E}[\|z_t - z^*\|^2 - \|z_t - z_{t+1}\|^2 - \|z_{t+1} - z^*\|^2].
\end{aligned} \tag{20}$$

and

$$\begin{aligned}
 & 2\mathbb{E}[F(z_t; \xi_i)^\top (z_{t+1/2} - z_{t+1})] \\
 &= \frac{2}{\eta} \mathbb{E}[(z_t - z_{t+1/2})^\top (z_{t+1/2} - z_{t+1})] \\
 &= \frac{1}{\eta} \mathbb{E}[\|z_t - z_{t+1}\|^2 - \|z_t - z_{t+1/2}\|^2 - \|z_{t+1/2} - z_{t+1}\|^2]
 \end{aligned} \tag{21}$$

Using $\eta \leq 1/(4L)$ and Young's inequality we obtain

$$\begin{aligned}
 & 2\mathbb{E}[(F(z_{t+1/2}; \xi_j) - F(z_t; \xi_i))^\top (z_{t+1/2} - z_{t+1})] \\
 &\leq \mathbb{E} \left[2\eta \|F(z_{t+1/2}; \xi_j) - F(z_t; \xi_i)\|^2 + \frac{1}{2\eta} \|z_{t+1/2} - z_{t+1}\|^2 \right] \\
 &\leq \mathbb{E} [6\eta \|F(z_{t+1/2}) - F(z_{t+1/2}; \xi_j)\|^2 + 6\eta \|F(z_t) - F(z_t; \xi_i)\|^2] \\
 &\quad + \mathbb{E} \left[6\eta \|F(z_{t+1/2}) - F(z_t)\|^2 + \frac{1}{2\eta} \|z_{t+1/2} - z_{t+1}\|^2 \right] \\
 &\leq 6\eta(e_t^2 + e_{t+1/2}^2) + 6\eta L^2 \mathbb{E}[\|z_t - z_{t+1/2}\|^2] + \frac{1}{2\eta} \mathbb{E}[\|z_{t+1/2} - z_{t+1}\|^2] \\
 &\leq 6\eta(e_t^2 + e_{t+1/2}^2) + \frac{1}{2\eta} \mathbb{E}[\|z_t - z_{t+1/2}\|^2 + \|z_{t+1/2} - z_{t+1}\|^2],
 \end{aligned} \tag{22}$$

where the second last inequality follows from L -Lipschitz property.

The final step is to plug (19) (20) (21) (22) into (18) and then apply the strongly monotonicity which implying the fact that

$$F(z_{t+1/2})^\top (z_{t+1/2} - z^*) \geq \lambda \|z_{t+1/2} - z^*\|^2.$$

■

C.5 The Proof of Theorem 5.3 and Theorem 5.4

First of all, we present some useful lemmas.

Lemma C.1 (non-expansiveness of the resolvent) *Denote*

- $z^+ = (x^+, y^+)$ as the solution to $\min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x', y') + L\|x' - x\|^2 - L\|y' - y\|^2$;
- $w^+ = (u^+, v^+)$ as the solution to $\min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x', y') + L\|x' - u\|^2 - L\|y' - v\|^2$.

Then it holds that $\|z^+ - w^+\| \leq 2\|z - w\|$, where $z = (x, y)$ and $w = (u, v)$.

Proof It follows from

$$\|z^+ - w^+\| \leq \|z - w\| + \frac{1}{2L} \|F(z^+) - F(w^+)\| \leq \|z - w\| + \frac{1}{2} \|z^+ - w^+\|,$$

where the relationship $F(z^+) = 2L(z - z^+)$ and $F(w^+) = 2L(w - w^+)$ and the triangle inequality is used. ■

Lemma C.2 *If F is a L -Lipschitz continuous operator, then F_{2L} is a $6L$ -Lipschitz continuous operator.*

Proof Note that we have the following relationship of

$$F_{2L}(z) = F(z^+) = 2L(z - z^+).$$

Hence we know that the $F_{2L}(\cdot)$ is $6L$ -Lipschitz continuous by noting

$$\|F(z^+) - F(w^+)\| \leq 2L\|z - w\| + 2L\|z^+ - w^+\| \leq 6L\|z - w\|$$

where we use Lemma C.1 to obtain that $\|z^+ - w^+\| \leq 2\|z - w\|$. ■

As we have mentioned, for given k, s , the key is to control the bias as well as the variance term of the estimated gradient operator on envelope in SEG^{++} (Algorithm 6) within $\tilde{\mathcal{O}}(1)$ SFO complexity. We present the details in the following theorem.

Theorem C.1 *Under the setting of both Theorem 5.3 and Theorem 5.4, if the input of SEG^{++} (Algorithm 6) holds that*

$$\lambda\eta \geq \delta_{\min}, \quad T \leq T_{\max}, \quad \|z_s^* - z_0\| \leq D_{\max}, \quad \|w_0 - z_0^+\|^2 \leq \frac{\sigma^2}{4L^2} + D^2 \quad (23)$$

for some $D_{\max} \geq D$, $\delta_{\min} > 0$, and, $T_{\max} > 0$ and we set

$$K = \left\lceil \log_2 \left(\frac{12}{\delta_{\min}} \right) \right\rceil, \quad N = \left\lceil \log_2 \left(\max \left\{ \frac{2T_{\max}}{\delta_{\min}}, \frac{8D_{\max}^2 L^2}{\delta_{\min} \sigma^2} \right\} \right) \right\rceil, \quad M = 1792K, \quad (24)$$

then it guarantees

$$\lambda \mathbb{E}[\|z_{t+1/2} - z_s^*\|^2] \leq \frac{1}{\eta} \mathbb{E}[\|z_t - z_s^*\|^2 - \|z_{t+1} - z_s^*\|^2] - \frac{1}{2\eta} \mathbb{E}[\|z_{t+1} - z_{t+1/2}\|^2] + 16\eta\sigma^2 \quad (25)$$

for any $s = 0, 1, \dots, S-1$, where $z_s^* = (x_s^*, y_s^*)$ is the solution to

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} f_{2L}^{(s)}(x, y)$$

and $f_{2L}^{(s)}(x, y)$ is defined in (10) for ID case or (11) for NC case, respectively. Furthermore, the output holds that

$$\mathbb{E}\|\hat{z}_{J+1/2} - z_{J+1/2}^+\|^2 \leq \frac{\sigma^2}{4L^2}, \quad (26)$$

where $z^+ = (x^+, y^+)$ is the solution to $\min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x, y) + L\|x' - x\|^2 - L\|y' - y\|^2$; and each call of Epoch-SEG⁺ in SEG^{++} (line 6 and line 10 of Algorithm 6) can be finished within the SFO complexity of

$$\mathcal{O} \left(\max \left\{ \log \left(\frac{T_{\max}}{\delta_{\min}} \right), \log \left(\frac{D_{\max}^2 L^2}{\delta_{\min} \sigma^2} \right) \right\} \log \left(\frac{1}{\delta_{\min}} \right) \right). \quad (27)$$

Proof We introduce the following notations for our proof:

- We denote $F_{2L}^{(s)}$ as the gradient operator of $f_{2L}^{(s)}(x, y)$.
- We denote $z^+ = (x^+, y^+)$ as the solution to the minimax problem

$$\min_{x' \in \mathbb{R}^{d_x}} \max_{y' \in \mathbb{R}^{d_y}} f(x', y') + L\|x' - x\|^2 - L\|y' - y\|^2.$$

- In the procedure of Algorithm 5, we denote

$$\hat{F}_{2L}^{(s)}(z) = \begin{cases} 2L(z - \hat{z}^+) + \sum_{i=1}^s \lambda_i(z - z_i), & \text{ID case;} \\ 2L(z - \hat{z}^+) + \sum_{i=0}^s \lambda_i(z - z_i), & \text{NC case;} \end{cases}$$

for given $z = (x, y)$ and s . Note that $\hat{F}_{2L}^{(s)}$ is used to approximate $F_{2L}^{(s)}$ in our algorithm and analysis.

- We denote

$$\begin{aligned} b_t^{(s)} &\triangleq \|\mathbb{E}\hat{F}_{2L}^{(s)}(z_t) - F_{2L}^{(s)}(z_t)\|, & (\sigma_t^{(s)})^2 &\triangleq \mathbb{E}[\|\hat{F}_{2L}^{(s)}(z_t) - \mathbb{E}\hat{F}_{2L}^{(s)}(z_t)\|^2], \\ b_{t+1/2}^{(s)} &\triangleq \|\mathbb{E}\hat{F}_{2L}^{(s)}(z_{t+1/2}) - F_{2L}^{(s)}(z_{t+1/2})\|, & (\sigma_{t+1/2}^{(s)})^2 &\triangleq \mathbb{E}[\|\hat{F}_{2L}^{(s)}(z_{t+1/2}) - \mathbb{E}\hat{F}_{2L}^{(s)}(z_{t+1/2})\|^2]. \end{aligned}$$

for the bias and variance of the approximate gradient operator $F_{2L}^{(s)}$.

If the the bias and variance of $F_{2L}^{(s)}(\cdot)$ satisfy

$$b_t^{(s)} \leq 2\lambda\eta\sigma^2, \quad b_{t+1/2}^{(s)} \leq 2\lambda\eta\sigma^2, \quad (\sigma_t^{(s)})^2 \leq \frac{\sigma^2}{2} \quad \text{and} \quad (\sigma_{t+1/2}^{(s)})^2 \leq \frac{\sigma^2}{2}, \quad (28)$$

then applying Theorem 5.2 on $f_{2L}^{(s)}$ leads to the result of (25). The Lipschitz continuity of $\hat{F}_{2L}$ indicates the conditions (28) holds if we can prove the following claims:

$$\begin{aligned} \text{Claim I: } & 4L^2\|\mathbb{E}\hat{z}_t^+ - z_t^+\|^2 \leq 2\lambda\eta\sigma^2, & 4L^2\mathbb{E}[\|\hat{z}_t^+ - \mathbb{E}\hat{z}_t^+\|^2] &\leq \sigma^2/2; \\ \text{Claim II: } & 4L^2\|\mathbb{E}\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|^2 \leq 2\lambda\eta\sigma^2, & 4L^2\mathbb{E}[\|\hat{z}_{t+1/2}^+ - \mathbb{E}\hat{z}_{t+1/2}^+\|^2] &\leq \sigma^2/2; \end{aligned} \quad (29)$$

Now start to prove (29) holds for all t by induction, which implies the result of (25) in the theorem.

Induction Base: For Claim I with $t = 0$, applying Theorem 5.1 and using the fact $\hat{z}_{-1/2}^+ = w_0$, we have

$$\|\mathbb{E}\hat{z}_0^+ - z_0^+\|^2 \leq \frac{1}{2^{N+2K}} \times \|w_0 - z_0^+\|^2 + \frac{8\sigma^2}{2^K \times 3L^2}$$

and

$$\mathbb{E}\|\hat{z}_0^+ - \mathbb{E}\hat{z}_0^+\|^2 \leq \frac{1}{2^{NM}} \times 22\|w_0 - z_0^+\|^2 + \frac{112K\sigma^2}{M \times 3L^2}.$$

Plugging the setting of N, M and K as (24) into above inequalities and using the bound $\|w_0 - z_0^+\| \leq \sigma^2/(4L^2) + D^2$ followed from the assumption, we obtain Claim I for $t = 0$. The next step should be showing Claim II holds for $t = 0$. We can prove this result by the same way as proving Claim I $\Rightarrow$ Claim II, which will be detailed presented in upcoming paragraph.

Induction Step: Suppose the results of (29) hold for all $t' \leq t$, then we target to show **Claim I** and **Claim II** both hold for $t + 1$.

We first consider **Claim I**. Since we already have (29) for all $t' \leq t$, then it implies that (25) also holds for all $t' \leq t$ by using Theorem 5.2. Telescoping the result of (25) for $t' = 0, \dots, t$, we obtain

$$\mathbb{E}\|z_{t+1/2} - z_{t+1}\|^2 \leq 2\mathbb{E}\|z_0 - z_s^*\|^2 + 32\eta^2\sigma^2T_{\max} \quad (30)$$

and

$$\mathbb{E}\|z_{t+1} - z_s^*\|^2 \leq \mathbb{E}\|z_0 - z_s^*\|^2 + 16\eta^2\sigma^2T_{\max}. \quad (31)$$

Using the bias-variance decomposition, i.e.

$$\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|^2 = \|z_{t+1/2}^+ - \mathbb{E}\hat{z}_{t+1/2}^+\|^2 + \mathbb{E}\|\hat{z}_{t+1/2}^+ - \mathbb{E}\hat{z}_{t+1/2}^+\|^2$$

and the induction hypothesis, we know that

$$4L^2\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|^2 \leq \sigma^2.$$

Then we have

$$\begin{aligned} & \mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1}^+\|^2 \\ & \leq 2\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|^2 + 2\mathbb{E}\|z_{t+1/2}^+ - z_{t+1}^+\|^2 \\ & \leq 2\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|^2 + 8\mathbb{E}\|z_{t+1/2} - z_{t+1}\|^2 \\ & \leq \frac{\sigma^2}{2L^2} + 16\mathbb{E}\|z_0 - z_s^*\|^2 + 256\eta^2\sigma^2T_{\max} \\ & \leq \frac{\sigma^2T_{\max}}{L^2} + 16\mathbb{E}\|z_0 - z_s^*\|^2, \end{aligned} \quad (32)$$

where we use the bound for $\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1/2}^+\|$ by induction hypothesis and the bound of $\mathbb{E}\|z_{t+1/2} - z_{t+1}\|^2$ shown in (30).

Note that the algorithm apply the update rule

$$\hat{z}_{t+1}^+ \leftarrow \text{Epoch-SEG}^+(g_{t+1}, \hat{z}_{t+1/2}^+, L, 3L, N, K, M).$$

Applying Theorem 5.1, we have

$$\begin{aligned} & \|z_{t+1}^+ - \mathbb{E}\hat{z}_{t+1}^+\|^2 \\ & \leq \frac{1}{2^{N+2K}} \times \mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1}^+\|^2 + \frac{8\sigma^2}{2^K \times 3L^2} \\ & \leq \underbrace{\frac{1}{2^{N+2K}} \times \left(\frac{\sigma^2T_{\max}}{L^2} + 16\mathbb{E}\|z_0 - z_s^*\|^2 \right)}_{A_1} + \underbrace{\frac{8\sigma^2}{2^K \times 3L^2}}_{A_2}. \end{aligned}$$

and

$$\begin{aligned}
 & \mathbb{E}\|\hat{z}_{t+1}^+ - \mathbb{E}\hat{z}_{t+1}^+\|^2 \\
 & \leq \frac{1}{2^N M} \times 22\mathbb{E}\|\hat{z}_{t+1/2}^+ - z_{t+1}^+\|^2 + \frac{112K\sigma^2}{M \times 3L^2} \\
 & \leq \underbrace{\frac{1}{2^N M} \times \left(\frac{22\sigma^2 T_{\max}}{L^2} + 352\mathbb{E}\|z_0 - z_s^*\|^2 \right)}_{B_1} + \underbrace{\frac{112K\sigma^2}{M \times 3L^2}}_{B_2}.
 \end{aligned}$$

The parameters setting as (24) guarantees A_1, A_2, B_1 and B_2 are sufficient small, which leads to

$$4L^2\|\mathbb{E}\hat{z}_{t+1}^+ - z_{t+1}^+\|^2 \leq 2\lambda\eta\sigma^2 \quad \text{and} \quad 4L^2\mathbb{E}[\|\hat{z}_{t+1}^+ - \mathbb{E}\hat{z}_{t+1}^+\|^2] \leq \frac{\sigma^2}{2}. \quad (33)$$

Then we finish the proof of **Claim I** for $t + 1$.

Now we show that **Claim I** $\Rightarrow$ **Claim II** to finish our induction. We begin from

$$\begin{aligned}
 & \mathbb{E}\|z_{t+3/2} - z_{t+1}\|^2 \\
 & = \eta^2 \mathbb{E}\|\hat{F}_{2L}^{(s)}(z_{t+1})\|^2 \\
 & \leq 2\eta^2 \mathbb{E}\|F_{2L}^{(s)}(z_{t+1})\|^2 + 2\eta^2 \mathbb{E}\|F_{2L}^{(s)}(z_{t+1}) - \hat{F}_{2L}^{(s)}(z_{t+1})\|^2 \\
 & \leq 1152\eta^2 L^2 \mathbb{E}\|z_{t+1} - z_s^*\|^2 + 2\eta^2 \sigma^2 \\
 & \leq 1152\eta^2 L^2 (\mathbb{E}\|z_0 - z_s^*\|^2 + 16\eta^2 \sigma^2 T_{\max}) + 2\eta^2 \sigma^2,
 \end{aligned} \quad (34)$$

where the equality is based on the update $z_{t+3/2} \leftarrow z_{t+1} - \eta\hat{F}_{2L}^{(s)}(z_{t+1})$; the inequalities use the fact that $F_{2L}^{(s)}$ is $24L$ -Lipschitz continuous and the upper bound of $\mathbb{E}\|z_{t+1} - z_s^*\|^2$ which we have shown in (31).³

Furthermore, the results of (33) imply

$$4L^2\mathbb{E}\|z_{t+1}^+ - \hat{z}_{t+1}^+\|^2 \leq \sigma^2.$$

Combining Lemma C.1 with the bound of $\mathbb{E}\|z_{t+3/2} - z_{t+1}\|^2$ as shown in (34), we obtain

$$\begin{aligned}
 & \mathbb{E}\|z_{t+3/2}^+ - \hat{z}_{t+1}^+\|^2 \\
 & \leq 2\mathbb{E}\|z_{t+1}^+ - \hat{z}_{t+1}^+\|^2 + 2\mathbb{E}\|z_{t+3/2}^+ - z_{t+1}^+\|^2 \\
 & \leq 2\mathbb{E}\|z_{t+1}^+ - \hat{z}_{t+1}^+\|^2 + 8\mathbb{E}\|z_{t+3/2} - z_{t+1}\|^2 \\
 & \leq \frac{\sigma^2 T_{\max}}{L^2} + 16\mathbb{E}\|z_0 - z_s^*\|^2,
 \end{aligned}$$

which is derived by the similar way to the proof of (32). Note that the algorithm use the update

$$\hat{z}_{t+3/2}^+ \leftarrow \text{Epoch-SEG}^+(g_{t+1}, \hat{z}_{t+1}^+, L, 3L, N, K, M).$$

3. Note that in the case of $t = -1$, bound of $\mathbb{E}[\|z_{t+1} - z_s^*\|^2]$ is independent on the induction hypothesis.

Applying Theorem 5.1, we have

$$\begin{aligned}
& \|\mathbb{E}\hat{z}_{t+3/2}^+ - z_{t+3/2}^+\|^2 \\
& \leq \frac{1}{2^{N+2K}} \times \mathbb{E}\|\hat{z}_{t+1}^+ - z_{t+3/2}^+\|^2 + \frac{8\sigma^2}{2^K \times 3L^2} \\
& \leq \underbrace{\frac{1}{2^{N+2K}} \times \left(\frac{\sigma^2 T_{\max}}{L^2} + 16\mathbb{E}\|z_0 - z_s^*\|^2 \right)}_{A_1} + \underbrace{\frac{8\sigma^2}{2^K \times 3L^2}}_{A_2}.
\end{aligned}$$

and

$$\begin{aligned}
& \mathbb{E}\|\hat{z}_{t+3/2}^+ - \mathbb{E}\hat{z}_{t+3/2}^+\|^2 \\
& \leq \frac{1}{2^N M} \times 22\mathbb{E}\|\hat{z}_{t+1}^+ - \hat{z}_{t+3/2}^+\|^2 + \frac{112K\sigma^2}{M \times 3L^2} \\
& \leq \underbrace{\frac{1}{2^N M} \times \left(\frac{22\sigma^2 T_{\max}}{L^2} + 352\mathbb{E}\|z_0 - z_s^*\|^2 \right)}_{B_1} + \underbrace{\frac{112K\sigma^2}{M \times 3L^2}}_{B_2}.
\end{aligned}$$

The parameters setting of (24) guarantees

$$4L^2\|\mathbb{E}\hat{z}_{t+3/2}^+ - z_{t+3/2}^+\|^2 \leq 2\lambda\eta\sigma^2 \quad \text{and} \quad 4L^2\mathbb{E}\|\hat{z}_{t+3/2}^+ - \mathbb{E}\hat{z}_{t+3/2}^+\|^2 \leq \frac{\sigma^2}{2}, \quad (35)$$

which means **Claim II** holds. Hence, we have complete the induction.

We can verify the condition number of sub-problem is no more than 3: since it is L -strongly-convex- L -strongly-concave and $3L$ -smooth. Then based on the setting of (24) and Theorem 5.1, it guarantees that the SFO complexity of each call of Epoch-SEG⁺ (line 6 and line 10 of Algorithm 6) is no more than

$$48MN + 192MK = \mathcal{O}\left(\max\left\{\log\left(\frac{T_{\max}}{\delta_{\min}}\right), \log\left(\frac{D_{\max}^2 L^2}{\delta_{\min} \sigma^2}\right)\right\} \log\left(\frac{1}{\delta_{\min}}\right)\right).$$

■

C.5.1 THE PROOF OF THEOREM 5.3 AND THEOREM 5.4

Note that the parameters N_s and K_s in Theorem 5.3 follow (6) by replacing L with $6L$. Our proof of Theorem 5.3 and Theorem 5.4 will be described by the notation of Algorithm 5.

Proof The main steps of RAIN⁺⁺ (Algorithm 5) are based on calling SEG⁺⁺ with $s = 0, 1, \dots, S-1$ and $k = 0, 1, \dots, N_s + K_s - 1$. By Theorem C.1, once the initial conditions of (23) hold for the call of SEG⁺⁺ in RAIN⁺⁺ (Algorithm 5), we achieve the bound (25) on $f_{2L}^{(s)}(x, y)$ (which is defined in (10) for ID case and (11) for NC case) for any s (just like running SEG on f_s). Then the SFO complexity of each SEG⁺⁺ call in RAIN⁺⁺ is

$$l = \mathcal{O}\left(\max\left\{\log\left(\frac{T_{\max}}{\delta_{\min}}\right), \log\left(\frac{D_{\max}^2 L^2}{\delta_{\min} \sigma^2}\right)\right\} \log\left(\frac{1}{\delta_{\min}}\right)\right). \quad (36)$$

Note that RAIN^{++} can be regarded as the modification of RAIN by replacing the sub-routines SEG with SEG^{++} . Compared with SEG , SEG^{++} requires additional stochastic oracle calls to estimate $F(\cdot)$. Therefore, the total SFO complexity of RAIN^{++} is the total SFO complexity of RAIN (shown in Theorem 4.1) multiplying l , which implies the SFO complexity of $\tilde{O}(\sigma^2\epsilon^{-2} + L/\lambda)$ we desired. Hence, the remains in the proof only need to verify there exist some $T_{\max}$, $\delta_{\min}$ and $D_{\max}$ to guarantee the conditions of (23) hold for each call of SEG^{++} in RAIN^{++} (line 5 and line 7 in Algorithm 5).

The upper bound of $\|w_{s,k} - z_{s,k}^+\|^2$: Consider the first time when SEG^{++} (Algorithm 6) is called in Algorithm 5 (i.e. $s = 0$ and $k = 0$). Note that we have $w_{s,k} = z_{s,k} = z_0$, then Lemma C.1 means $\|z_0^+ - z^*\| \leq 2\|z_0 - z^*\|$ and we achieve

$$\|w_{s,k} - z_{s,k}^+\| = \|z_0 - z_0^+\| = \frac{1}{2L}\|F(z_0^+)\| = \frac{1}{2L}\|F(z_0^+) - F(z^*)\| \leq \frac{1}{2}\|z_0^+ - z^*\| \leq \|z_0 - z^*\| \leq D.$$

Now we consider the other cases (i.e. $s > 0$ or $k > 0$). In these rounds, SEG^{++} use $(z_{s,k}, w_{s,k})$ as the initial point, which is the output of the previous call of SEG^{++} . Then it holds that

$$\mathbb{E}\|z_{s,k}^+ - w_{s,k}\|^2 \leq \frac{\sigma^2}{4L^2}$$

by the result of (26) in Theorem C.1. Therefore, combining above the two cases, it holds that

$$\mathbb{E}\|z_{s,k}^+ - w_{s,k}\|^2 \leq \frac{\sigma^2}{4L^2} + D^2$$

for all s and k .

The settings of $T_{\max}$ and $1/\delta_{\min}$: We present a simple bound for $T_{\max}$ and $1/\delta_{\min}$ from (6) and (17). For each stage of s in Algorithm 5, we run SEG^{++} for two phases, that is $k = 0, \dots, N_{s-1}$ and $k = N_s, \dots, N_s + K_s + 1$.

- For $k = 0, \dots, N_{s-1}$, we run SEG^{++} by the stepsize of $1/(48L)$ and the iteration number of $96L/\lambda_s$. Since $\lambda_s \geq \lambda$, setting $\delta_{\min}$ and $T_{\max}$ with

$$\frac{1}{\delta_{\min}} \leq \frac{48L}{\lambda} \quad \text{and} \quad T_{\max} \leq \frac{96L}{\lambda} \quad (37)$$

satisfies the conditions in (23).

- For $k = N_s, \dots, N_s + K_s + 1$, the stepsize of SEG^{++} is decreasing and the iteration numbers of SEG^{++} is decreasing. So we only needs to consider the last time we call SEG^{++} , whose stepsize and iteration numbers are $1/(96L \times 2^{K_s - N_s})$ and $384 \times 2^{K_s - N_s}$ respectively. Since we have $\lambda_s \leq 6L$, setting

$$T_{\max} \leq 384 \times 2^{K_{s-1} - N_{s-1}} \leq 786432S^2\sigma^2\epsilon^{-2}, \quad (38)$$

and

$$\frac{1}{\delta_{\min}} \leq \frac{1}{\lambda_{s-1}/(96L \times 2^{K_{s-1} - N_{s-1}})} \leq 8192S^2\sigma^2\epsilon^{-2}. \quad (39)$$

satisfies the conditions in (23).

Combining the bounds of (37), (38) and (39), we obtain

$$\begin{aligned} \frac{1}{\delta_{\min}} &\leq \max \{48L/\lambda, 8192S^2\sigma^2\epsilon^{-2}\} \quad \text{and} \\ T_{\max} &\leq \max \{96L/\lambda, 786432S^2\sigma^2\epsilon^{-2}\}. \end{aligned} \quad (40)$$

The setting of $D_{\max}$: This paragraph shows for all s and k , the term $\mathbb{E}[\|z_{s,k} - z_s^*\|^2]$ can be bounded by some positive constant $D_{\max}^2$. Since Line 4-8 in Algorithm 5 can be regarded as running Epoch-SEG on $f_{2L}^{(s)}$, the result (16) in the proof of Lemma 4.2 means

$$\mathbb{E}\|z_{s,k} - z_s^*\|^2 \leq \mathbb{E}\|z_{s,0} - z_s^*\|^2 + \frac{8\sigma^2}{\lambda L} \quad (41)$$

for any $k = 0, 1, \dots, N_s + K_s - 1$. Now we bound $\mathbb{E}\|z_{s,0} - z_s^*\|^2$.

- For $s = 0$, we have $z_{s,0} = z_0$ and $z_s^* = z^*$, which directly implies $\mathbb{E}\|z_{s,0} - z_s^*\|^2 = D^2$.
- For $s \geq 1$, the result of (17) guarantees

$$256\lambda_{s-1}^2 S^2 \mathbb{E}\|z_{s,0} - z_{s-1}^*\|^2 \leq \epsilon^2.$$

Using the non-expansiveness after anchoring (Lemma A.1), we obtain

$$\mathbb{E}\|z_{s,0} - z_s^*\|^2 \leq \mathbb{E}\|z_{s,0} - z_{s-1}^*\|^2 \leq \frac{\epsilon^2}{256\lambda^2 S^2}.$$

for any $s \geq 1$.

Finally, combining above two cases and (41) means the setting

$$D_{\max}^2 \leq \max \left\{ D^2, \frac{\epsilon^2}{256\lambda^2 S^2} \right\} + \frac{8\sigma^2}{\lambda L}. \quad (42)$$

satisfies the condition in (23).

In summary, we can set $\delta_{\min}$, $T_{\max}$ and $D_{\max}$ by following (40) and (42) to guarantee the conditions in (23) hold, then the SFO complexity of each SEG^{++} call in RAIN^{++} is $l = \text{polylog}(L/\lambda, 1/\epsilon, D, \sigma) = \tilde{O}(1)$ and total SFO complexity of RAIN^{++} is $\tilde{O}(\sigma^2\epsilon^{-2} + L/\lambda)$. Since Proposition 5.1 says $\lambda = \Theta(\alpha)$, the SFO complexity is $\tilde{O}(\sigma^2\epsilon^{-2} + L/\alpha)$ for ID case. For the NC case, we takes $\lambda = \Theta(\epsilon)$, which leads to the SFO complexity of $\tilde{O}(\sigma^2\epsilon^{-2} + L/\epsilon)$. ■

C.6 The Proof of Proposition 5.3

Proof We denote F_{2L} as the gradient operator of $f_{2L}(x, y)$ and F as the gradient operator of $f(x, y)$. Let $z^+ = (x^+, y^+)$ be the solution to

$$\min_{x' \in \mathbb{R}_x^d} \max_{y' \in \mathbb{R}_y^d} g(x, y) \triangleq f(x, y) + L\|x' - x\|^2 - L\|y - y'\|^2.$$

For any $w = (u, v)$, we have

$$\|F(w)\| \leq \|F(z^+)\| + \|F(z^+) - F(w)\|.$$

Since it holds that $F(z^+) = F_{2L}(z)$ and we already have $\|F(z^+)\| \leq \epsilon$, the smoothness of F means

$$\mathbb{E}\|w - z^+\| \leq \frac{\epsilon}{L}, \quad (43)$$

which can make sure that $w = (u, v)$ is a 2ϵ -stationary point of $f(x, y)$ in expectation. We can verify the condition number of $g(x, y)$ is $\Theta(1)$, which means finding $w = (u, b)$ that satisfies (43) can be finished within $\mathcal{O}(\log(\epsilon^{-1}) + \sigma^2\epsilon^{-2})$ SFO complexity by running Epoch-SEG (Algorithm 2) and following the result of Lemma 4.2. $\blacksquare$

Appendix D. The Proofs in Section 6

Foster et al. (2019) showed that the lower bound complexity for finding a point with small gradient can be decomposed as the statistical complexity given by stochastic oracle and the optimization complexity with deterministic oracle. We following their ideas to construct our lower bounds for stochastic minimax optimization.

D.1 The Proof of Theorem 6.1

Proof We first consider the statistical complexity. There exists an L -smooth and convex function f_{sample} such that any $\mathcal{A}$ needs at least an SFO complexity of $\Omega(\sigma^2\epsilon^{-2}\log(L\epsilon^{-1}))$ to find its ϵ -stationary point (Foster et al., 2019, Theorem 2). The worst-case convex function for the sample complexity Foster et al. (2019) is given by its first-order derivative

$$f'_{\text{sample}}(x; Z_t) = \begin{cases} -2\epsilon, & x < 0, \\ 2\epsilon, & x \geq D, \\ \frac{x - a_j}{D/N} \sigma Z_{t,j+1} + \left(1 - \frac{x - a_j}{D/N}\right) \sigma Z_{t,j} & x \in [a_j, a_{j+1}) \text{ for some } j < N, \end{cases}$$

where $Z_{t,j} \in \{-1, 1\}$ is drawn from the following distribution

$$\mathbb{P}(Z_{t,j} = 1) = \begin{cases} \frac{1}{2} - p, & j \leq j^*, \\ \frac{1}{2} + p, & j > j^*. \end{cases}$$

In above equations, the values of $(a_1, \dots, a_N)$ and j^* is given by the lower bound for noisy binary search problem Foster et al. (2019). Setting $N = LD/(4\epsilon)$ with $D = \|x^* - x_0\|$, we obtain the function

$$F(x) = \mathbb{E}_{Z_t}[f_{\text{sample}}(x; Z_t)],$$

which is L -smooth and convex; and leads to a lower bound

$$\Omega(\sigma^2\epsilon^{-2}\log(LD\epsilon^{-1})) \quad (44)$$

for finding an ϵ -stationary point of $F(x)$. For minimax optimization, we construct the function

$$H(x, y) = F(x) - F(y),$$

that is L -smooth and convex-concave. Naturally, it provide the lower bound (44) for finding an ϵ -stationary point for $H(x, y)$.

Then we consider the optimization complexity. We consider the strongly-convex-strongly-concave function Luo et al. (2021) as follows

$$f_{\text{SCSC}}(x, y) = \frac{\lambda' r}{2} \|x\|^2 + \lambda' x^\top (By - c) - \frac{\lambda' r}{2} \|y\|^2,$$

where

$$B = \begin{bmatrix} 1 & & & & & \\ -1 & 1 & & & & \\ & & \ddots & \ddots & & \\ & & & -1 & 1 & \\ & & & & -1 & \sqrt{r\omega} \end{bmatrix} \in \mathbb{R}^{d \times d}, \quad c = \begin{bmatrix} \omega \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad \text{and} \quad \omega = \frac{\sqrt{r^2 + 4} - r}{2}.$$

By setting

$$r = \sqrt{\frac{8}{L^2/\lambda'^2 - 2}}, \quad \lambda' = \frac{\lambda}{r} \quad \text{and} \quad d = \left\lfloor \frac{1}{r} \log \left(\frac{1}{2\epsilon} \right) \right\rfloor - 4,$$

we obtain the special case of the lower bound in Theorem 1 of Luo et al. (2021) with $n = 1$. It is shown that Luo et al. (2021) the function $f_{\text{SCSC}}(x, y)$ is L -smooth, λ -strongly-convex- λ -strongly-concave and any first-order algorithm requires at least $\Omega(\kappa \log(1/\epsilon))$ number of gradient calls to obtain a point z such that $\|z - z^*\| \leq \epsilon$. Since it holds that $\|\nabla f(z)\| \geq \lambda\|z - z^*\|$, any $\mathcal{A}$ needs at least an SFO complexity of

$$\Omega(\kappa \log(\lambda \epsilon^{-1})) \tag{45}$$

to find an ϵ -stationary point of $f_{\text{SCSC}}(x, y)$.

Combining the statistical complexity of (44) and the optimization complexity of (45) completes the proof. $\blacksquare$

D.2 The Proof of Theorem 6.2

Proof The statistical complexity can be obtained by following the proof of Theorem 6.1. For the optimization complexity, Theorem 3 of Yoon and Ryu (2021) provide the convex-concave function of the form

$$f_{\text{CC}}(x, y) = (x - x^*)^\top A(y - x^*),$$

where $A \in \mathbb{R}^{d \times d}$ is symmetric and x^* is the component of a stationary point $z^* = (x^*, y^*)$. Yoon and Ryu (2021) showed that there exist some A and x^* such that $0 \preceq A \preceq L$ and finding an ϵ -stationary point of the corresponding $f_{\text{CC}}(x, y)$ requires at least $\Omega(LD\epsilon^{-1})$ gradient calls. We complete the proof by combining the optimization complexity in Theorem 6.1. $\blacksquare$

Algorithm 7 SEAG (f, z_0, η, T)

```

1: for  $t = 0, 1, \dots, T - 1$  do
2:    $\xi_i \leftarrow$  a random index
3:    $z_{t+1/2} \leftarrow z_t - \left(1 - \frac{1}{t+1}\right) \eta F(z_t; \xi_i) + \frac{1}{t+1}(z_0 - z_t)$ 
4:    $\xi_j \leftarrow$  a random index
5:    $z_{t+1} \leftarrow z_t - \eta F(z_{t+1/2}; \xi_j) + \frac{1}{t+1}(z_0 - z_t)$ 
6: end for
7: return  $z_T$ 
    
```

Algorithm 8 PDHG (f, x_0, y_0, η, T)

```

1:  $\bar{x}_0 = x_0, \bar{y}_0 = y_0$ 
2:  $\xi_y^0 \leftarrow$  a random index
3:  $s_0 = \nabla_y f(x_0, y_0; \xi_y^0)$ 
4: for  $t = 0, 1, \dots, T - 1$  do
5:    $y_{t+1} \rightarrow y_t - \eta s_t$ 
6:    $\xi_x^t \leftarrow$  a random index
7:    $x_{t+1} \leftarrow x_t - \eta \nabla_x f(x_t, y_{t+1}; \xi_x^t)$ 
8:    $\xi_y^t \leftarrow$  a random index
9:    $s_{t+1} \leftarrow \frac{t+2}{t+1} \nabla_y f(x_t, y_t; \xi_y^t) - \frac{t}{t+1} \nabla_y f(x_t, y_t; \xi_y^{t-1})$ 
10:   $\bar{x}_{t+1} \leftarrow \frac{t-1}{t+1} \bar{x}_t + \frac{2}{t+1} x_{t+1}$ 
11:   $\bar{y}_{t+1} \leftarrow \frac{t-1}{t+1} \bar{y}_t + \frac{2}{t+1} y_{t+1}$ 
12: end for
13: return  $(\bar{x}_T, \bar{y}_T)$ 
    
```

D.3 The Proof of Corollary 6.1

Proof Since the gradient operator of convex-concave function is monotone, it is also negative comonotone. Hence, the lower bound in Theorem 6.2 is also applicable for negative comonotone setting and the result of this corollary is obtained. $\blacksquare$

D.4 The Proof of Corollary 6.2

Proof For $\tau \geq 2L$, the assumptions of L -smooth and α -strongly-convex- α -strongly-concave on $f(x, y)$ directly means the function is (τ, α) -intersection-dominant. Hence, the lower bound in Theorem 6.1 is also applicable for (τ, α) -intersection-dominant setting and the result of this corollary is obtained. $\blacksquare$

Algorithm 9 A Single-Loop Variant of RAIN ($f, z_0, \eta, T, \lambda, \gamma$)

```

1: for  $s = 0, 1, \dots, T - 1$  do
2:    $\xi_i \leftarrow$  a random index
3:    $z_{t+1/2} \leftarrow z_t - \eta(F(z_t; \xi_i) + \lambda\gamma \sum_{j=0}^{t-1} (1 + \gamma)^j (z_t - z_j))$ 
4:    $\xi_j \leftarrow$  a random index
5:    $z_{t+1} \leftarrow z_t - \eta(F(z_{t+1/2}; \xi_j) + \lambda\gamma \sum_{j=0}^{t-1} (1 + \gamma)^j (z_{t+1/2} - z_j))$ 
6: end for
7: return  $z_T$ 

```

Appendix E. Details for Numerical Experiments

We give the details for the implementation of the algorithms in our experiments.

E.1 The Concave-Concave Case

The implementations of algorithms in convex-concave case:

- For ordinary stochastic extragradient method (Algorithm 3, SEG), we replace the output from uniform sampling in all the history with the point in the last iteration for better performance.
- For regularized SEG (R-SEG), we use the regularization trick of Nesterov (2012) for minimax optimization, that is using SEG to solve the regularized minimax problem

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} g(x, y) \triangleq f(x, y) + \frac{\lambda}{2} \|x - x_0\|^2 - \frac{\lambda}{2} \|y - y_0\|^2$$

for some small λ .

- For stochastic extra anchor gradient (SEAG), we follow the algorithm by Lee and Kim (2021), which is described in Algorithm 7.
- For primal-dual hybrid gradient (PDHG), we follow the Algorithm 1 by Zhao (2022) with $\alpha_t = \tau_t = \eta$, which is described in Algorithm 8.
- For RAIN (Algorithm 1), we directly set $N_s = 1$ and $K_s = 0$. This simplified variant yields a single-loop implementation described in Algorithm 9, and we observe it has good performance in practice.

E.2 The Nonconcave-Nonconcave Case

The implementations of algorithms in nonconvex-nonconcave case:

- For SEG^+ , we follow Equation (EG_p^+) by Diakonikolas et al. (2021) with $\beta = 1/2$ and $\alpha_k = \eta$, by but output the point in the last iteration for better performance.
- For SFEG, we follow Algorithm 1 by (Lee and Kim, 2021) by replacing the exact gradient with the stochastic gradient.

Algorithm 10 RAIN $(z_0, \lambda, L, \{N_s\}_{s=0}^{S-1}, \{K_s\}_{s=0}^{S-1}, \gamma)$

```

1:  $\lambda_0 = \gamma\lambda, S = \lfloor \log_{(1+\gamma)}(L/\lambda) \rfloor$ 
2: for  $s = 0, 1, \dots, S - 1$ 
3:    $z_{s,0} \leftarrow z_s$ 
4:   for  $k = 0, 1, \dots, N_s - 1$ 
5:      $z_{s,k,0} \leftarrow z_{s,k}, \eta \leftarrow \frac{1}{8L}, T \leftarrow \frac{16L}{\lambda_s}$ 
6:     for  $t = 0, 1, \dots, T - 1$ 
7:        $z_{s,k,t+1/2} = z_{s,k,t} - \eta(F(z_{s,k,t}; \xi_{s,k,t}) + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t} - z_j))$ 
8:        $z_{s,k,t+1} = z_{s,k,t} - \eta(F(z_{s,k,t+1/2}; \xi_{s,k,t+1/2}) + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t+1/2} - z_j))$ 
9:     end for
10:     $z_{s,k+1} \leftarrow$  uniformly samples from  $\{z_{s,k,t+1/2}\}_{t=0}^{T-1}$ 
11:  end for
12:  for  $k = N_s, N_s + 1, \dots, N_s + K_s - 1$ 
13:     $z_{s,k,0} \leftarrow z_{s,k}, \eta \leftarrow \frac{1}{2^{k-N_s+4}L}, T \leftarrow \frac{2^{k-N_s+6}}{\lambda_s}$ 
14:    for  $t = 0, 1, \dots, T - 1$ 
15:       $z_{s,k,t+1/2} = z_{s,k,t} - \eta(F(z_{s,k,t}; \xi_{s,k,t}) + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t} - z_j))$ 
16:       $z_{s,k,t+1} = z_{s,k,t} - \eta(F(z_{s,k,t+1/2}; \xi_{s,k,t+1/2}) + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t+1/2} - z_j))$ 
17:    end for
18:     $z_{s,k+1} \leftarrow$  uniformly samples from  $\{z_{s,k,t+1/2}\}_{t=0}^{T-1}$ 
19:  end for
20:   $z_{s+1} \leftarrow z_{s,N_s+K_s}$ 
21:   $\lambda_{s+1} \leftarrow (1 + \gamma)\lambda_s$ 
22: end for
23: return  $z_S$ 

```

- For RAIN⁺⁺, we follow the setting of Algorithm 9 ($N_s = 1$ and $K_s = 0$) for the subroutine RAIN. For the steps of MLMC, we set $N = 1, M = 1$ and $K = 1$. This simplified variant is easy to implement and performs well in practice.

E.3 Hyperparameter Selection

For each algorithm, we tune the parameters η from $\{0.005, 0.01, 0.05, 0.1, 1, 5, 10\}$, λ from $\{0.001, 0.01, 0.1, 1\}$ and γ from $\{0.001, 0.01, 0.1, 1\}$ and reports the best run.

Appendix F. The Merged Algorithms for Easy Reference

In the main text, we present our algorithm by nested functions to facilitate the theoretical analysis. In this section, we write all steps of RAIN in Algorithm 1 without the presentation of any subroutine call, which is easy to follow for readers who are interested in the implementation. We also provided a merged presentation for RAIN⁺⁺ in Algorithm 11.

Algorithm 11 RAIN⁺⁺ ($z_0, \lambda, L, \{N_s\}_{s=0}^{S-1}, \{K_s\}_{s=0}^{S-1}, \gamma, N, K, M$)

```

1:  $\lambda_0 = \gamma\lambda, S = \lfloor \log_{(1+\gamma)}(3L/\lambda) \rfloor$ 
2: for  $s = 0, 1, \dots, S-1$ 
3:    $z_{s,0} \leftarrow z_s, w_{s,0} \leftarrow w_s$ 
4:   for  $k = 0, 1, \dots, N_s - 1$ 
5:      $z_{s,k,0} \leftarrow z_{s,k}, w_{s,k,-1/2} \leftarrow w_{s,k}, \eta \leftarrow \frac{1}{24L}, T \leftarrow \frac{48L}{\lambda_s}$ 
6:     for  $t = 0, 1, \dots, T-1$ 
7:        $(\hat{F}_{s,k,t}, w_{s,k,t}) = \text{EnvGradEst}(z_{s,k,t}, w_{s,k,t-1/2}, L, N, K, M)$ 
8:        $z_{s,k,t+1/2} = z_{s,k,t} - \eta(\hat{F}_{s,k,t} + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t} - z_j))$ 
9:        $(\hat{F}_{s,k,t+1/2}, w_{s,k,t+1/2}) = \text{EnvGradEst}(z_{s,k,t}, w_{s,k,t}, L, N, K, M)$ 
10:       $z_{s,k,t+1} = z_{s,k,t} - \eta(\hat{F}_{s,k,t+1/2} + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t+1/2} - z_j))$ 
11:    end for
12:    Draw  $J \sim \text{Unif}([T])$ 
13:     $z_{s,k+1} \leftarrow z_{s,k,J+1/2}, w_{s,k+1} \leftarrow w_{s,k,J+1/2}$ 
14:  end for
15:  for  $k = N_s, N_s + 1, \dots, N_s + K_s - 1$ 
16:     $z_{s,k,0} \leftarrow z_{s,k}, w_{s,k,-1/2} \leftarrow w_{s,k}, \eta \leftarrow \frac{1}{2^{k-N_s+4} \cdot 3L}, T \leftarrow \frac{2^{k-N_s+6} \cdot 3L}{\lambda_s}$ 
17:    for  $t = 0, 1, \dots, T-1$ 
18:       $(\hat{F}_{s,k,t}, w_{s,k,t}) = \text{EnvGradEst}(z_{s,k,t}, w_{s,k,t-1/2}, L, N, K, M)$ 
19:       $z_{s,k,t+1/2} = z_{s,k,t} - \eta(\hat{F}_{s,k,t} + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t} - z_j))$ 
20:       $(\hat{F}_{s,k,t+1/2}, w_{s,k,t+1/2}) = \text{EnvGradEst}(z_{s,k,t}, w_{s,k,t}, L, N, K, M)$ 
21:       $z_{s,k,t+1} = z_{s,k,t} - \eta(\hat{F}_{s,k,t+1/2} + \sum_{j=0}^{s-1} \lambda_j(z_{s,k,t+1/2} - z_j))$ 
22:    end for
23:    Draw  $J \sim \text{Unif}([T])$ 
24:     $z_{s,k+1} \leftarrow z_{s,k,J+1/2}, w_{s,k+1} \leftarrow w_{s,k,J+1/2}$ 
25:  end for
26:   $z_{s+1} \leftarrow z_{s,N_s+K_s}, w_{s+1} \leftarrow w_{s,N_s+K_s}$ 
27:   $\lambda_{s+1} \leftarrow (1 + \gamma)\lambda_s$ 
28: end for
29: return  $z_S$ 

```

Since the steps of RAIN⁺⁺ is indeed complicated, it still includes one subroutine (Algorithm 12) to present the evaluation for the gradient of the saddle envelope. Note that EnvGradEst (Algorithm 12) achieves a nearly unbiased gradient estimator of the saddle envelope and the initial point in the next subroutine call. In the main text, EnvGradEst is presented by

Algorithm 12 EnvGradEst($\bar{z}, z_0, L, N, K, M$)

```

1: for  $m = 0, 1, \dots, M - 1$  do
2:   draw  $J \sim \text{Geom}(1/2)$ 
3:   for  $k = 0, 1, \dots, N - 1$ 
4:      $z_{m,k,0} \leftarrow z_0, \eta \leftarrow \frac{1}{12L}, T \leftarrow 24$ 
5:     for  $t = 0, 1, \dots, T - 1$ 
6:        $z_{m,k,t+1/2} \leftarrow z_{m,k,t} - \eta(F(z_{m,k,t}; \xi_{m,k,t}) + 2L(z_{m,k,t} - \bar{z}))$ 
7:        $z_{m,k,t+1} \leftarrow z_{m,k,t} - \eta(F(z_{m,k,t+1/2}; \xi_{m,k,t+1/2}) + 2L(z_{m,k,t+1/2} - \bar{z}))$ 
8:     end for
9:      $z_{m,k+1} \leftarrow$  uniformly samples from  $\{z_{m,k,t+1/2}\}_{t=0}^{T-1}$ 
10:   end for
11:   for  $k = N, N + 1, \dots, N + J - 1$ 
12:      $z_{m,k,0} \leftarrow z_{m,k}, \eta \leftarrow \frac{1}{2^{k-N+3} \cdot 3L}, T \leftarrow 2^{k-N+5} \cdot 3L$ 
13:     for  $t = 0, 1, \dots, T - 1$ 
14:        $z_{m,k,t+1/2} \leftarrow z_{m,k,t} - \eta(F(z_{m,k,t}; \xi_{m,k,t}) + 2L(z_{m,k,t} - \bar{z}))$ 
15:        $z_{m,k,t+1} \leftarrow z_{m,k,t} - \eta(F(z_{m,k,t+1/2}; \xi_{m,k,t+1/2}) + 2L(z_{m,k,t+1/2} - \bar{z}))$ 
16:     end for
17:      $z_{m,k+1} \leftarrow$  uniformly samples from  $\{z_{m,k,t+1/2}\}_{t=0}^{T-1}$ 
18:   end for
19:    $\hat{z}_m = z_{m,N} + 2^J(z_{m,N+J} - z_{m,N+J-1})\mathbb{I}[J \leq K]$ 
20: end for
21:  $\hat{z}^+ \leftarrow \frac{1}{M} \sum_{m=0}^{M-1} \hat{z}_m, \hat{F} = 2L(\bar{z} - \hat{z}^+)$ 
22: return  $(\hat{F}, \hat{z}^+)$ 
    
```

applying Epoch-SEG⁺ on the sub-problem

$$\min_{x \in \mathbb{R}^{d_x}} \max_{y \in \mathbb{R}^{d_y}} g_{t,s,k}(x, y) := f(x_{t,s,k}, y_{t,s,k}) + L\|x - x_{t,s,k}\|^2 - L\|y - y_{t,s,k}\|^2$$

to output $\hat{F}_{s,k,t} = 2L(z_{s,k,t} - w_{s,k,t})$ and $w_{s,k,t}$. We can verify they are equivalent.

References

- Ahmet Alacaoglu and Yura Malitsky. Stochastic variance reduction for variational inequality methods. In *COLT*, 2022.
- Zeyuan Allen-Zhu. How to make the gradients small stochastically: Even faster convex and nonconvex SGD. In *NeurIPS*, 2018.
- Hilal Asi, Yair Carmon, Arun Jambulapati, Yujia Jin, and Aaron Sidford. Stochastic bias-reduced gradient methods. In *NeurIPS*, 2021.
- Xufeng Cai, Chaobing Song, Cristóbal Guzmán, and Jelena Diakonikolas. A stochastic halpern iteration with variance reduction for stochastic monotone inclusion problems. In *NeurIPS*, 2022.
- Bo Dai, Albert Shaw, Niao He, Lihong Li, and Le Song. Boosting the actor with dual critic. In *ICLR*, 2018.
- Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS*, 2021.
- Dylan J. Foster, Ayush Sekhari, Ohad Shamir, Nathan Srebro, Karthik Sridharan, and Blake Woodworth. The complexity of making the gradient small in stochastic convex optimization. In *COLT*, 2019.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014a.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014b.
- Benjamin Grimmer, Haihao Lu, Pratik Worah, and Vahab Mirrokni. The landscape of the proximal point method for nonconvex–nonconcave minimax optimization. *Mathematical Programming*, 201(1):373–407, 2023.
- Zhishuai Guo, Zhuoning Yuan, Yan Yan, and Tianbao Yang. Fast objective & duality gap convergence for nonconvex-strongly-concave min-max problems. *arXiv preprint arXiv:2006.06889*, 2020.
- Dmitry Kovalev and Alexander Gasnikov. The first optimal algorithm for smooth and strongly-convex-strongly-concave minimax optimization. *arXiv preprint arXiv:2205.05653*, 2022.
- Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. In *NeurIPS*, 2021.
- Tianyi Lin, Chi Jin, and Michael I. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *ICML*, 2020a.
- Tianyi Lin, Chi Jin, and Michael I. Jordan. Near-optimal algorithms for minimax optimization. In *COLT*, 2020b.

- Mingrui Liu and Francesco Orabona. On the initialization for convex-concave min-max problems. In *International Conference on Algorithmic Learning Theory*, 2022.
- Mingrui Liu, Zhuoning Yuan, Yiming Ying, and Tianbao Yang. Stochastic AUC maximization with deep neural networks. *arXiv preprint arXiv:1908.10831*, 2019.
- Mingrui Liu, Wei Zhang, Youssef Mroueh, Xiaodong Cui, Jarret Ross, Tianbao Yang, and Payel Das. A decentralized parallel algorithm for training generative adversarial nets. In *NeurIPS*, 2020.
- Mingrui Liu, Hassan Rafique, Qihang Lin, and Tianbao Yang. First-order convergence theory for weakly-convex-weakly-concave min-max problems. *Journal of Machine Learning Research*, 2021.
- Luo Luo, Guangzeng Xie, Tong Zhang, and Zhihua Zhang. Near optimal stochastic algorithms for finite-sum unbalanced convex-concave minimax optimization. *arXiv preprint arXiv:2106.01761*, 2021.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Yurii Nesterov. How to make the gradients small. *Optima. Mathematical Optimization Society Newsletter*, 2012.
- Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P. How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *ICML*, 2017.
- Thomas Pethick, Panagiotis Patrinos, Olivier Fercoq, Volkan Cevher, et al. Escaping limit cycles: Global convergence for constrained nonconvex-nonconcave minimax problems. In *ICLR*, 2022.
- R. Tyrrell Rockafellar. Monotone operators associated with saddle-functions and minimax problems. *Nonlinear functional analysis*, 18(part 1):397–407, 1970.
- Hoi-To Wai, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong. Multi-agent reinforcement learning via double averaging primal-dual optimization. In *NeurIPS*, 2018.
- Wenhan Xian, Feihu Huang, Yanfu Zhang, and Heng Huang. A faster decentralized algorithm for nonconvex minimax problems. In *NeurIPS*, 2021.
- Tengyu Xu, Zhe Wang, Yingbin Liang, and H Vincent Poor. Enhanced first and zeroth order variance reduced algorithms for min-max optimization. *arXiv preprint arXiv:2006.09361*, 2020.
- Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In *NeurIPS*, 2020a.
- Junchi Yang, Siqi Zhang, Negar Kiyavash, and Niao He. A catalyst framework for minimax optimization. *NeurIPS*, 2020b.

- Tianbao Yang and Yiming Ying. AUC maximization in the era of big data and ai: A survey. *ACM Computing Surveys*, 2022.
- TaeHo Yoon and Ernest K. Ryu. Accelerated algorithms for smooth convex-concave minimax problems with $\mathcal{O}(1/k^2)$ rate on squared gradient norm. In *ICML*, 2021.
- Zhuoning Yuan, Zhishuai Guo, Yi Xu, Yiming Ying, and Tianbao Yang. Federated deep AUC maximization for heterogeneous data with a constant communication complexity. *arXiv preprint arXiv:2102.04635*, 2021.
- Liang Zhang, Kiran Koshy Thekumparampil, Sewoong Oh, and Niao He. Bring your own algorithm for optimal differentially private stochastic minimax optimization. *arXiv preprint arXiv:2206.00363*, 2022a.
- Siqi Zhang, Junchi Yang, Cristóbal Guzmán, Negar Kiyavash, and Niao He. The complexity of nonconvex-strongly-concave minimax optimization. In *UAI*, 2021.
- Siqi Zhang, Yifan Hu, Liang Zhang, and Niao He. Uniform convergence and generalization for nonconvex stochastic minimax problems. *arXiv preprint arXiv:2205.14278*, 2022b.
- Renbo Zhao. Accelerated stochastic algorithms for convex-concave saddle-point problems. *Mathematics of Operations Research*, 2022.