Zeroth-Order Methods for Nonconvex-Strongly Concave Stochastic Minimax Problems with Decision-Dependent Distributions

Zili Luo and Yongchao Liu*

March 19, 2026

Abstract

Stochastic minimax problems with decision-dependent distributions (SMDD) have emerged as a crucial framework for modeling complex systems where data distributions drift in response to decision variables. Most existing methods for SMDD rely on an explicit functional relationship between the decision variables and the probability distribution. In this paper, we propose two sample-based zeroth-order algorithms, namely Single-Directional Zeroth-Order Stochastic Gradient Descent Ascent (SD-ZO-SGDA) and Multi-Directional Zeroth-Order Stochastic Gradient Descent Ascent (MD-ZO-SGDA), to find an ϵ-stationary point of the nonconvex-strongly concave SMDD. Under accuracy-independent constant step-sizes, SD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-2})$ with a batch size of $\mathcal{O}(\epsilon^{-4}\kappa^2)$, where d denotes the total variable dimension and κ is the condition number. Under accuracy-dependent step-sizes, SD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-5})$ with a relaxed batch size of $\mathcal{O}(\epsilon^{-2}\kappa^2)$. On the other hand, the multi-directional variant MD-ZO-SGDA further reduces gradient estimation variance, which improves the complexity dependence on the condition number κ and dimension d under both step-size settings. Numerical experiments on a synthetic problem and a real-world competitive EV charging application demonstrate the effectiveness of our algorithms.

Keywords: Zeroth-order methods, stochastic minimax problems, decision-dependent distributions, nonconvex-strongly concave, sample-based.

1 Introduction

Stochastic minimax optimization is a fundamental framework for nested decision-making, which optimizes opposing variables through simultaneous minimization and maximization. Its flexibility in handling competitive objectives enables broad applications, ranging from adversarial training [7, 12] and distributionally robust optimization [18, 24, 36] to strategic game theory [3, 26] and policy evaluation in reinforcement learning [30, 31]. In many modern machine learning and cyber-physical systems, the data distribution is rarely static. Specifically, the data distributions drift in response to decision variables. For example, consumer demand reacts to price changes [5, 6],

*School of Mathematical Sciences, Dalian University of Technology, Dalian 116024, China, e-mail: zililuo@mail.dlut.edu.cn (Zili Luo), lyc@dlut.edu.cn (Yongchao Liu)

navigation algorithms redirect traffic flows [20, 27], and credit scoring models affect loan terms and repayment behaviors [16, 29]. In scenarios where the probability distribution depends on the decision variables through a distribution map, the corresponding stochastic minimax problems with decision-dependent distributions (SMDD) [33] can be formulated as follows:

$$\min_{x\in\mathbb{R}^{d_1}} \max_{y\in\mathcal{Y}} F(x,y) := \mathbb{E}_{\xi\sim\mathcal{D}(x,y)}[f(x,y,\xi)], \tag{1}$$

where ξ is a random variable supported on $\mathcal{Z}\subset\mathbb{R}^n$, $\mathcal{Y}\subset\mathbb{R}^{d_2}$, $f(\cdot):\mathbb{R}^{d_1}\times\mathbb{R}^{d_2}\times\mathbb{R}^n\to\mathbb{R}$ is a scalar-valued function, $D(\cdot):\mathbb{R}^{d_1}\times\mathbb{R}^{d_2}\to\mathcal{P}(\mathbb{R}^n)$ is a distribution map and $\mathbb{E}[\cdot]$ denotes the expectation with respect to the distribution $D(x,y)$.

The model of SMDD is proposed by Wood and Dall'Anese [33]. The authors introduce the notation of the performative equilibrium points which are the saddle point for the stochastic minimax problem whose probability distribution is induced by itself, and provide sufficient conditions for existence and uniqueness of the performative equilibrium points. They develop a stochastic primal-dual algorithm, and prove linear and $\mathcal{O}(1/\sqrt{t})$ convergence rates under constant and dynamic step-sizes, respectively. Narang et al. [25] study a decision-dependent stochastic game model and propose the concept of performatively stable equilibrium. The authors introduce repeated retraining and the repeated (stochastic) gradient methods for finding performatively stable equilibrium. When the decision-dependent stochastic game is strongly monotone, they propose an adaptive stochastic gradient method for finding Nash equilibrium, and demonstrate $\mathcal{O}(1/\sqrt{t})$ convergence when the distribution follows a location-scale model. On the other hand, for the saddle point of the SMDD, Wood and Dall'Anese [33] propose a zeroth-order algorithm for finding the saddle point in the strongly convex–strongly concave regime. Notably, their analysis necessitates specific structural assumptions on the distribution map $D(x,y)$, namely opposing mixture dominance. Wood et al. [34] explore a two-stage method for seeking Nash equilibrium which estimates the distribution map first and then solves the game by the stochastic gradient descent method. More recently, Gao and Liu [9] propose the adaptive stochastic gradient descent ascent algorithm and its alternating variant to find the stationary points of SMDD, which learns the unknown distribution map dynamically and optimizes the minimax problem simultaneously via (alternating) stochastic gradient descent ascent at each iteration. When the distribution map follows a location-scale model, the authors analyze the non-asymptotic complexity under nonconvex–(strongly) concave and nonconvex–PŁ settings of the objective function, respectively. Beyond the gradient-based framework, Gao et al. [10] propose a stochastic trust region algorithm for SMDD, which employs a local regression model to locally learn the distribution map from samples. By solving trust-region subproblems within the learned environment, the algorithm attains almost sure convergence to a stationary point under the nonconvex-strongly concave setting.

While these pioneering works have significantly advanced the study of SMDD, most of them mainly rely on the explicit parametric models of the distribution map. In this work, we consider a more general and challenging setting that relies only on the accessibility of samples following the decision-dependent distribution, without requiring explicit knowledge of its internal structure. To this end, we propose two sample-based zeroth-order algorithms to solve the nonconvex-strongly concave SMDD. As far as we are concerned, the contributions of this paper can be summarized as follows.

- We propose two sample-based zeroth-order algorithms, namely Single-Directional Zeroth-

Order Stochastic Gradient Descent Ascent (SD-ZO-SGDA) and Multi-Directional Zeroth-Order Stochastic Gradient Descent Ascent (MD-ZO-SGDA), to solve nonconvex-strongly concave SMDD. Specifically, SD-ZO-SGDA generalizes the standard stochastic gradient descent ascent (SGDA) method by replacing the unavailable stochastic gradients with single-directional two-point gradient estimators. On the other hand, compared with SD-ZO-SGDA, MD-ZO-SGDA leverages multiple random directions per iteration to construct more accurate gradient estimators, which effectively reduces estimation variance and improves optimization stability. SD-ZO-SGDA and MD-ZO-SGDA are sample-based, requiring only the accessibility of sample realizations from the decision-dependent distribution without any prior knowledge of its explicit functional relationship or internal structure.

- Under accuracy-independent constant step-sizes, SD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-2})$ with a batch size of $\mathcal{O}(\epsilon^{-4}\kappa^2)$, where d denotes the total variable dimension and κ is the condition number. Under accuracy-dependent step-sizes, SD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-5})$ with a relaxed batch size of $\mathcal{O}(\epsilon^{-2}\kappa^2)$. Our analysis quantifies the impact of distributional shifts on zeroth-order estimation, and identifies that the coupling between decisions and distributions introduces an irreducible variance term of order $\mathcal{O}(\sigma^2/(\mu^2 b))$, where σ^2 is a uniform upper bound of the variance of the stochastic function f, μ is the smoothing parameter, and b is the batch size. Compared to existing literature [9, 10, 33, 34] that rely on the Wasserstein continuity of the distribution map, our theoretical framework demonstrates that convergence can be established without this assumption. This implies that our algorithms remain robust under a broader class of decision-dependent distributions.

- Under accuracy-independent constant step-sizes, MD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-2})$ with a batch size of $\mathcal{O}(\epsilon^{-4}\kappa^2 d^2)$. Under accuracy-dependent step-sizes, MD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-5})$ with a relaxed batch size of $\mathcal{O}(\epsilon^{-2}\kappa^2 d^2)$. These results demonstrate that leveraging multi-directional sampling effectively reduces estimation variance. This leads to two key theoretical improvements. MD-ZO-SGDA improves the dependence on the condition number from $\mathcal{O}(\kappa^3)$ to $\mathcal{O}(\kappa^2)$ compared to the single-directional method. When $m \propto d^2$, MD-ZO-SGDA achieves a dimension-independent iteration complexity of $\mathcal{O}(\kappa^2\epsilon^2)$. To verify the effectiveness of the algorithms, we conduct numerical experiments on a synthetic minimax problem and a real-world competitive EV charging application.

Next, we review the literature on zeroth-order minimax optimization and zeroth-order methods in decision-dependent settings that is most relevant to our work.

Zeroth-Order Minimax Optimization. Zeroth-order methods are a class of powerful optimization tools to solve many complex problems, whose explicit gradients are infeasible to access. For nonconvex-strongly concave minimax optimization problems, Liu et al. [23] pioneer this area by proposing ZO-Min-Max, an alternating projected stochastic gradient descent ascent method equipped with a zeroth-order gradient estimator. The authors establish that ZO-Min-Max finds an ϵ-stationary point with a sample complexity of $\mathcal{O}((d_x + d_y)\epsilon^{-4})$. Subsequently, Wang et al. [32] expand this framework by analyzing both Zeroth-Order Gradient Descent Ascent (ZO-GDA) and Zeroth-Order Gradient Descent Multi-Step Ascent (ZO-GDMSA). For deterministic

nonconvex-strongly concave problems, they prove that these algorithms require $\mathcal{O}(\kappa^5(d_x+d_y)\epsilon^{-2})$ and $\mathcal{O}(\kappa(d_x+\kappa d_y\log(\epsilon^{-1}))\epsilon^{-2})$ oracle calls, respectively. In the stochastic setting, the corresponding sample complexities increase to $\mathcal{O}(\kappa^5(d_x+d_y)\epsilon^{-4})$ and $\mathcal{O}(\kappa(d_x+\kappa d_y\log(\epsilon^{-1}))\epsilon^{-4})$. To further improve sample efficiency, researchers have integrated variance reduction and momentum techniques. For instance, Xu et al. [35] introduce the ZO-VRGDA algorithm, leveraging variance reduction to achieve an improved sample complexity of $\mathcal{O}(\kappa^3(d_x+d_y)\epsilon^{-3})$ in stochastic environments. Similarly, Huang et al. [15] propose an accelerated zeroth-order momentum descent ascent (Acc-ZOMDA) method, which attains an ϵ-stationary point with a sample complexity of $\mathcal{O}(\kappa^3(d_x+d_y)^{3/2}\epsilon^{-3})$. These methods primarily focus on scenarios with a static data distribution. In contrast, our work addresses environments characterized by decision-dependent distributions. This dependency presents a fundamental technical challenge: the distributional shift prevents the effective elimination of stochastic noise through common random numbers, a standard technique in static settings. Consequently, this leads to an irreducible variance term of order $\mathcal{O}(\sigma^2/(\mu^2 b))$. A critical implication of this result is that the smoothing parameter μ cannot be chosen arbitrarily small to reduce smoothing bias, as it would lead to a significant expansion of the variance. Our framework requires a careful balance between the smoothing parameter μ and the batch size b to ensure robust convergence in the presence of dynamic distributional responses.

Zeroth-Order Methods in Decision-Dependent Settings. Research on applying zeroth-order methods to decision-dependent environments is still in its nascent stage. The most relevant work is by Wood and Dall'Anese [33], who briefly discuss a zeroth-order algorithm for SMDD. Their theoretical results are established in the strongly convex-strongly concave setting. For minimization problems with decision-dependent distributions, some zeroth-order methods have been proposed. For instance, Ray et al. [28] and Liu et al. [22] propose zeroth-order methods employing one-point gradient estimators to address optimization problems where the distribution evolves dynamically over time. While these works handle distribution shifts, one-point estimators are subject to higher variance than multi-point schemes, which often leads to slower convergence. More recently, Chen et al. [4] explore a two-point gradient estimator for decision-dependent problems. However, their theoretical guarantees rely on reducing the problem to convex optimization, which is only valid under restrictive combinations of the objective function and the distribution map. Hikima and Takeda [13] advance this field by introducing a variance-reduced one-point estimator and a two-point gradient estimator for decision-dependent minimization. Crucially, their two-point method is designed to operate without stringent structural assumptions on the distribution map or the objective function. Furthermore, Hikima and Takeda [14] conduct a unified sample complexity analysis of zeroth-order methods across gradient estimators with different search directions. Despite these advancements in zeroth-order gradient estimation, existing theoretical guarantees in decision-dependent settings remain confined to either standard minimization tasks or strongly convex-strongly concave setting minimax setting. Our work bridges this gap by integrating these foundational estimation techniques, ranging from two-point to multi-directional schemes, into the more challenging nonconvex-strongly concave minimax framework.

The remainder of this paper is organized as follows. Section 3 introduces the necessary notations and presents the standard assumptions. Section 4 introduces SD-ZO-SGDA and provides its theoretical convergence analysis under accuracy-independent constant as well as accuracy-dependent step-sizes. Section **??** describes MD-ZO-SGDA and establishes its improved theoretical convergence

guarantees under both step-size schemes. Finally, Section 5 conducts numerical experiments on a synthetic problem and a competitive EV charging application to demonstrate the effectiveness of the proposed algorithms.

2 Preliminaries

Throughout this paper, $\langle x, y\rangle$ denotes the inner product of two vectors x and y. $\mathbb{R}^d$ denotes the d-dimensional Euclidean space endowed with norm $\|x\| = \sqrt{\langle x, x\rangle}$. I_d denotes a d-dimensional identity matrix. Given a convex closed set $\mathcal{X}$, we define a projection operation to $\mathcal{X}$ as $P_{\mathcal{X}}(\cdot)$. We denote $a = \mathcal{O}(b)$ if $|a| \leq C|b|$ for some constant $C > 0$, and $a = \Theta(b)$ if there exist constants $C_1, C_2 > 0$ such that $C_1|b| \leq |a| \leq C_2|b|$. The minimax problem (1) can be reformulated as following minimization of the primal function $\Phi(x)$:

$$\min_{x\in\mathbb{R}^{d_1}} \Phi(x) := \max_{y\in\mathcal{Y}} F(x, y) = F(x, y^*(x)), \tag{2}$$

where $y^*(x) = \arg\max_{y\in\mathcal{Y}} F(x, y)$. Next, we introduce some standard assumptions.

Assumption 1. The function $\Phi(x)$ is bounded from below in $\mathbb{R}^{d_1}$, i.e., $\Phi^* = \inf_{x\in\mathbb{R}^{d_1}} \Phi(x)$.

Assumption 2. For any $x \in \mathbb{R}^{d_1}$, $y \in \mathbb{R}^{d_2}$, there exists a constant $\sigma \geq 0$ such that

$$\mathbb{E}_{\xi\sim\mathcal{D}(x,y)}[(F(x, y) - f(x, y, \xi))^2] \leq \sigma^2. \tag{3}$$

Assumption 3. The function $F(x, y)$ is L_f-smooth, i.e.,

$$||\nabla F(x, y) - \nabla F(x', y')|| \leq L_f||(x, y) - (x', y')||, \forall x, x' \in \mathbb{R}^{d_1}, y, y' \in \mathcal{Y}.$$

Assumption 4. The objective function $F(x, y)$ is τ-strongly concave in variable y, i.e.,

$$\|\nabla_y F(x, y) - \nabla_y F(x, y')\| \geq \tau\|y - y'\|, \forall x \in \mathbb{R}^{d_1}, y, y' \in \mathcal{Y}, \tag{4}$$

and the constraint set $\mathcal{Y}$ is a compact convex set with a diameter $D > 0$ such that $\|y - y'\| \leq D$ for all $y, y' \in \mathcal{Y}$.

Assumption 1 ensures that the optimization problem is well-posed by bounding the objective from below, preventing the iterative process from diverging. Assumption 2 is essential for estimating the objective value via finite samples and controlling the variance in zeroth-order gradient estimation. Assumption 3 is a standard smoothness requirement that ensures the reliability of gradient-based updates and facilitates the convergence analysis via Taylor expansion. Finally, Assumption 4 guarantees the existence and uniqueness of the optimal response $y^*(x)$, which ensures that the primal function $\Phi(x)$ is well-defined and differentiable.

2.1 Single-Directional Two-Point Gradient Estimators

As previously discussed, zeroth-order methods have been extensively adopted for solving saddle-point and decision-dependent problems. To address the SMDD (1), we apply the Uniform smoothing Gradient Estimator (UniGE) [8, 17] to generate two-point zeroth-order gradient estimators,

$$G_{\mu_1}(x, y, u_1) = \frac{F(x + \mu_1 u_1, y) - F(x, y)}{\mu_1/d_1} u_1, \tag{5}$$

$$H_{\mu_2}(x,y,u_2) = \frac{F(x,y+\mu_2 u_2) - F(x,y)}{\mu_2/d_2} u_2, \tag{6}$$

where $\mu_1, \mu_2 > 0$ are the smoothing parameters, and the random vectors $u_1 \in \mathbb{R}^{d_1}$ and $u_2 \in \mathbb{R}^{d_2}$ are drawn from the uniform distributions over the unit spheres, denoted by $\mathcal{U}_{Sp_1}$ and $\mathcal{U}_{Sp_2}$, respectively.

Note that $F(x,y)$ is an expectation function, its exact value is typically inaccessible in practical scenarios. Assuming that sample realizations from the decision-dependent distribution $D(x,y)$ are available for any $(x,y) \in \mathbb{R}^{d_1} \times \mathcal{Y}$, the sample-based approximations of the zeroth-order gradient estimators in (5) and (6) are defined as follows:

$$G_{\mu_1}(x,y,u_1,\mathcal{B},\mathcal{B}_{\mu_1}) = \frac{1}{b}\sum_{i=1}^{b} G_{\mu_1}(x,y,u_1,\tilde{\xi}_i,\xi_{1,i}), \tag{7}$$

$$H_{\mu_2}(x,y,u_2,\mathcal{B},\mathcal{B}_{\mu_2}) = \frac{1}{b}\sum_{i=1}^{b} H_{\mu_1}(x,y,u_2,\tilde{\xi}_i,\xi_{2,i}), \tag{8}$$

where b is the batch size, and the mini-batches are drawn as follows: $\mathcal{B} = \{\tilde{\xi}_i\}_{i=1}^b$ is drawn i.i.d. from $D(x,y)$, $\mathcal{B}_{\mu_1} = \{\xi_{1,i}\}_{i=1}^b$ is drawn i.i.d. from $D(x+\mu_1 u_1, y)$, and $\mathcal{B}_{\mu_2} = \{\xi_{2,i}\}_{i=1}^b$ is drawn i.i.d. from $D(x, y+\mu_2 u_2)$.

Denote $F_{\mu_1}(x,y) = \mathbb{E}_{u_1\sim\mathcal{U}_{b_1}}[F(x+\mu_1 u_1, y)]$ and $F_{\mu_2}(x,y) = \mathbb{E}_{u_2\sim\mathcal{U}_{b_2}}[F(x, y+\mu_2 u_2)]$ as the smoothing functions, where $\mathcal{U}_{b_1}$ and $\mathcal{U}_{b_2}$ represent the uniform distributions over the unit Euclidean balls in $\mathbb{R}^{d_1}$ and $\mathbb{R}^{d_2}$, respectively. Following [8, Lemma 4.1], we have

$$\begin{aligned} \mathbb{E}_{u_1\sim\mathcal{U}_{Sp_1}}\mathbb{E}_{\mathcal{B},\mathcal{B}_{\mu_1}}[G_{\mu_1}(x,y,u_1,\mathcal{B},\mathcal{B}_{\mu_1})] &= \nabla_x F_{\mu_1}(x,y), \\ \mathbb{E}_{u_2\sim\mathcal{U}_{Sp_2}}\mathbb{E}_{\mathcal{B},\mathcal{B}_{\mu_2}}[H_{\mu_2}(x,y,u_2,\mathcal{B},\mathcal{B}_{\mu_2})] &= \nabla_y F_{\mu_2}(x,y). \end{aligned} \tag{9}$$

2.2 Multi-Directional Two-Point Gradient Estimators

In stochastic optimization with static distributions, the randomness of the data ξ is independent of the perturbation direction u. However, in decision-dependent settings, the distribution shifts with every perturbation $x+\mu u$, creating a tight coupling between the exploration direction and the data-generating mechanism. Under a finite query budget, a critical design choice arises between allocating m samples to a single exploration direction to suppress sampling noise, or querying m distinct directions with a single sample each to reduce estimation variance.

We adopt the latter multi-directional strategy. By exploring m independent random directions per iteration, we can effectively suppress the variance between the gradient estimator and the true gradient. This approach simultaneously reduces the directional bias inherent in zeroth-order estimation and the sampling noise from the shifting distributions. Accordingly, we consider the following multi-directional zeroth-order gradient estimators:

$$G^M_{\mu_1}(x,y,u_1) = \frac{d_1}{\mu_1 m}\sum_{i=1}^{m}\left(F(x+\mu_1 u_{1,i}, y) - F(x,y)\right) u_{1,i}, \tag{10}$$

$$H^M_{\mu_2}(x,y,u_2) = \frac{d_2}{\mu_2 m}\sum_{i=1}^{m}\left(F(x, y+\mu_2 u_{2,i}) - F(x,y)\right) u_{2,i}, \tag{11}$$

where μ_1 and μ_2 are the smoothing parameters, $u_{1,i} \in \mathbb{R}^{d_1}$ and $u_{2,i} \in \mathbb{R}^{d_2}$ are generated from the uniform distribution over the unit sphere $\mathcal{U}_{Sp_1}$ and $\mathcal{U}_{Sp_2}$, respectively.

Similarly, we define the sample-based multi-directional zeroth-order gradient estimators as follows:

$$G_{\mu_1}(x, y, u_{M_1}, \tilde{\xi}_M, \xi_{1,M}) = \frac{d_1}{\mu_1 m} \sum_{i=1}^{m} \Big(f(x + \mu_1 u_{1,i}, y, \xi_{1.i}) - f(x, y, \tilde{\xi}_i) \Big) \cdot u_{1,i}, \tag{12}$$

$$H_{\mu_2}(x, y, u_{M_2}, \tilde{\xi}_M, \xi_{2,M}) = \frac{d_2}{\mu_2 m} \sum_{i=1}^{m} \Big(f(x, y + \mu_2 u_{2,i}, \xi_{2,i}) - f(x, y, \tilde{\xi}_i) \Big) \cdot u_{2,i}, \tag{13}$$

where m denotes the number of random perturbation directions, $u_{M_1} = \{u_{1,i}\}_{i=1}^m$ and $u_{M_2} = \{u_{2,i}\}_{i=1}^m$ represent the collections of independent random directions sampled from the unit spheres $\mathcal{U}_{Sp_1}$ and $\mathcal{U}_{Sp_2}$, respectively. Furthermore, $\tilde{\xi}_M = \{\tilde{\xi}_i\}_{i=1}^m$, $\xi_{1,M} = \{\xi_{1,i}\}_{i=1}^m$, and $\xi_{2,M} = \{\xi_{2,i}\}_{i=1}^m$ denote the sets of sample realizations drawn independently from the decision-dependent distributions $D(x, y)$, $D(x + \mu_1 u_{1,i}, y)$, and $D(x, y + \mu_2 u_{2,i})$, respectively.

By [8, Lemma 4.1], we have

$$\mathbb{E}_{u_{M_1} \sim \mathcal{U}_{Sp_1}} \mathbb{E}_{\tilde{\xi}_M, \xi_{1,M}} [G_{\mu_1}(x, y, u_{M_1}, \tilde{\xi}_M, \xi_{1,M})] = \nabla_x F_{\mu_1}(x, y),$$
$$\mathbb{E}_{u_{M_2} \sim \mathcal{U}_{Sp_2}} \mathbb{E}_{\tilde{\xi}_M, \xi_{2,M}} [H_{\mu_2}(x, y, u_{M_2}, \tilde{\xi}_M, \xi_{2,M})] = \nabla_y F_{\mu_2}(x, y).$$

3 SD-ZO-SGDA: Algorithm and Convergence Analysis

This section is devoted to the study of SD-ZO-SGDA. We first present the detailed steps of the algorithm in Section 3.1, followed by a rigorous establishment of its convergence theory in Section 3.2.

3.1 SD-ZO-SGDA algorithm

The proposed algorithm is as follows.

SD-ZO-SGDA operates through a sequential "sample-estimate-update" process at each iteration t. Specifically, Step 4 executes the sampling step, where three independent mini-batches, $\mathcal{B}_t, \mathcal{B}_{t,\mu_1}$, and $\mathcal{B}_{t,\mu_2}$ are drawn from the distributions induced by the current iterates $D(x_t, y_t)$ and its perturbed counterparts. These samples are utilized in Step 5 to construct the zeroth-order gradient estimators G_{μ_1} and H_{μ_2}, which approximate the first-order information of the objective function through randomized perturbations. Subsequently, Steps 6–7 perform the core SGDA updates: a stochastic gradient descent step is applied to the minimization variable x, and a projected stochastic gradient ascent step for the maximization variable variable y using the projection operator $\mathcal{P}_{\mathcal{Y}}(\cdot)$. By re-sampling from the shifting distribution before each gradient estimation, SD-ZO-SGDA ensures that the optimization process remains tightly synchronized with the evolving environment, which allows the iterates to converge toward an ϵ-stationary point without requiring explicit knowledge of the distribution map.

Algorithm 1 Single-Directional Zeroth-Order Stochastic Gradient Descent Ascent (SD-ZO-SGDA)

1: **Input**: iteration limit $T > 0$, initial input $x_0 \in \mathbb{R}^{d_1}, y_0 \in \mathcal{Y}$, step sizes (η_1, η_2), smoothing parameters $\mu_1, \mu_2 > 0$ and batch size b.

2: **Initialization**: Draw sample sets $\mathcal{B}^0 = \{\tilde{\xi}_i^0\}_{i=1}^b \overset{\text{i.i.d.}}{\sim} D(x_0, y_0)$, $\mathcal{B}_{\mu_1}^0 = \{\xi_{1,i}^0\}_{i=1}^b \overset{\text{i.i.d.}}{\sim} D(x_0 + \mu_1 u_1^0, y_0)$, and $\mathcal{B}_{\mu_2}^0 = \{\xi_{2,i}^0\}_{i=1}^b \overset{\text{i.i.d.}}{\sim} D(x_0, y_0 + \mu_2 u_2^0)$. Additionally, sample $u_1^0 \sim \mathcal{U}_{Sp_1}$ and $u_2^0 \sim \mathcal{U}_{Sp_2}$. Then, compute $G_{\mu_1}(x_0, y_0, u_1^0; \mathcal{B}^0, \mathcal{B}_{\mu_1}^0)$ and $H_{\mu_2}(x_0, y_0, u_2^0; \mathcal{B}^0, \mathcal{B}_{\mu_2}^0)$ using the estimators defined in (7) and (8).

3: **For** $t = 1, 2, \cdots, T$ **do**

4: Draw a mini-batch samples $\mathcal{B}^t = \{\tilde{\xi}_i^t\}_{i=1}^b$ from $D(x_t, y_t)$, $\mathcal{B}_{\mu_1}^t = \{\xi_{1,i}^t\}_{i=1}^b$ from $D(x_t + \mu_1 u_1^t, y_t)$ and $\mathcal{B}_{\mu_2}^t = \{\xi_{2,i}^t\}_{i=1}^b$ from $D(x_t, y_t + \mu_2 u_2^t)$;

5: Calculate $G_{\mu_1}(x_t, y_t, u_1^t; \mathcal{B}^t, \mathcal{B}_{\mu_1}^t)$, $H_{\mu_2}(x_t, y_t, u_2^t; \mathcal{B}^t, \mathcal{B}_{\mu_2}^t)$;

6: Update $x_{t+1} = x_t - \eta_1 G_{\mu_1}(x_t, y_t, u_1^t; \mathcal{B}^t, \mathcal{B}_{\mu_1}^t)$ with $u_1^t \sim \mathcal{U}_{Sp_1}$;

7: Update $y_{t+1} = \mathcal{P}_{\mathcal{Y}}(y_t + \eta_y H_{\mu_2}(x_t, y_t, u_2^t; \mathcal{B}^t, \mathcal{B}_{\mu_2}^t)$ with $u_2^t \sim \mathcal{U}_{Sp_2}$;

8: **end for**

3.2 Convergence analysis of SD-ZO-SGDA

In this subsection, we establish the theoretical convergence guarantees for the proposed SD-ZO-SGDA in the nonconvex-strongly concave setting. We first provide several technical lemmas to characterize the properties of our stochastic gradient estimators and then derive the convergence rate toward an ϵ-stationary point.

Lemma 1. Suppose that Assumption 2 holds. Then, for any $x \in \mathbb{R}^{d_1}, y \in \mathcal{Y}$ and any $\mu_1, \mu_2 > 0$,

$$\mathbb{E}_{u_1 \sim \mathcal{U}_{Sp_1}} \left[\left\| \frac{d_1}{\mu_1} (F(x + \mu_1 u_1, y) - F(x, y)) u_1 \right\|^2 \right] \leq \frac{\mu_1^2}{2} L_f^2 d_1^2 + 2 d_1 \|\nabla_x F(x, y)\|^2, \tag{14}$$

$$\mathbb{E}_{u_2 \sim \mathcal{U}_{Sp_2}} \left[\left\| \frac{d_2}{\mu_2} (F(x, y + \mu_2 u_2) - F(x, y)) u_2 \right\|^2 \right] \leq \frac{\mu_2^2}{2} L_f^2 d_2^2 + 2 d_2 \|\nabla_y F(x, y)\|^2. \tag{15}$$

Proof. Noting that F is L_f-smooth, we have

$$F(x + \mu_1 u_1, y) - F(x, y) \leq \mu_1 \langle \nabla_x F(x, y), u_1 \rangle + \frac{\mu_1^2}{2} L_f \|u_1\|^2,$$

and

$$F(x + \mu_1 u_1, y) - F(x, y) \geq \mu_1 \langle \nabla_x F(x, y), u_1 \rangle - \frac{\mu_1^2}{2} L_f \|u_1\|^2.$$

Combining the above two inequalities, we obtain

$$(F(x + \mu_1 u_1, y) - F(x, y))^2 \leq 2\mu_1^2 (\nabla_x F(x, y)^\top u_1)^2 + \frac{\mu_1^4}{2} L_f^2 \|u_1\|^4.$$

Taking the expectation with respect to u_1 on both sides of the above inequality, we arrive at

$$\begin{aligned}
&\mathbb{E}_{u_1}\left[\left\|\frac{d_1}{\mu_1}(F(x+\mu_1u_1,y)-F(x,y))u_1\right\|^2\right]\\
&=\frac{d_1^2}{\mu_1^2}\mathbb{E}_{u_1}\left[(F(x+\mu_1u_1,y)-F(x,y))^2\|u_1\|^2\right]\\
&\le\frac{d_1^2}{\mu_1^2}\left(\mathbb{E}_{u_1}\left[\frac{\mu_1^4}{2}L_f^2\|u_1\|^6\right]+\mathbb{E}_{u_1}\left[2\mu_1^2\langle\nabla_xF(x,y),u_1\rangle^2\|u_1\|^2\right]\right)\\
&\le\frac{\mu_1^2}{2}L_f^2d_1^2+2d_1\|\nabla_xF(x,y)\|^2,
\end{aligned}$$

where the last inequality follows from [1, (2.39)]. Then, we complete the proof of (14).

Similarly, we can prove (15) through an analogous derivation. This proof is complete. □

Lemma 2. Suppose that Assumptions 2 and 3 hold. Then, for any $x\in\mathbb{R}^{d_1}$, $y\in\mathcal{Y}$,

$$\mathbb{E}_{u_1\sim\mathcal{U}_{Sp_1}}\left[\mathbb{E}_{\mathcal{B},\mathcal{B}_{\mu_1}}\left[\|G_{\mu_1}(x,y,u_1,\mathcal{B},\mathcal{B}_{\mu_1})\|^2\right]\right]\le\frac{3}{2}\mu_1^2d_1^2L_f^2+6d_1\|\nabla_xF(x,y)\|^2+\frac{6\sigma^2d_1^2}{\mu_1^2b},\tag{16}$$

$$\mathbb{E}_{u_2\sim\mathcal{U}_{Sp_2}}\left[\mathbb{E}_{\mathcal{B},\mathcal{B}_{\mu_2}}\left[\|H_{\mu_2}(x,y,u_2,\mathcal{B},\mathcal{B}_{\mu_2})\|^2\right]\right]\le\frac{3}{2}\mu_2^2d_2^2L_f^2+6d_2\|\nabla_yF(x,y)\|^2+\frac{6\sigma^2d_2^2}{\mu_2^2b},\tag{17}$$

where $\mathcal{B}=\{\tilde{\xi}_i\}_{i=1}^b\overset{\text{i.i.d.}}{\sim}D(x,y)$, $\mathcal{B}_{\mu_1}=\{\xi_{1,i}\}_{i=1}^b\overset{\text{i.i.d.}}{\sim}D(x+\mu_1u_1,y)$, and $\mathcal{B}_{\mu_2}=\{\xi_{2,i}\}_{i=1}^b\overset{\text{i.i.d.}}{\sim}D(x,y+\mu_2u_2)$.

Proof. By the definition of the estimator (5), we have

$$\begin{aligned}
&\mathbb{E}_{u_1\sim\mathcal{U}_{Sp_1}}\left[\mathbb{E}_{\mathcal{B},\mathcal{B}_{\mu_1}}\left[\|G_{\mu_1}(x,y,u_1,\mathcal{B},\mathcal{B}_{\mu_1})\|^2\right]\right]\\
&=\mathbb{E}_{u_1,\mathcal{B},\mathcal{B}_{\mu_1}}\left[\left\|\frac{\frac{1}{b}\sum_{i=1}^bf(x+\mu_1u_1,y;\xi_{1,i})-\frac{1}{b}\sum_{i=1}^bf(x,y;\tilde{\xi}_i)}{\mu_1/d_1}u_1\right\|^2\right]\\
&=\mathbb{E}_{u_1,\mathcal{B},\mathcal{B}_{\mu_1}}\left[\left\|\frac{F(x+\mu_1u_1,y)-F(x,y)}{\mu_1/d_1}u_1+\frac{1}{b}\sum_{i=1}^b\frac{f(x+\mu_1u_1,y;\xi_{1,i})-F(x+\mu_1u_1,y)}{\mu_1/d_1}u_1\right.\right.\\
&\quad\left.\left.-\frac{1}{b}\sum_{i=1}^b\frac{f(x,y;\tilde{\xi}_i)-F(x,y)}{\mu_1/d_1}u_1\right\|^2\right]\\
&\le3\mathbb{E}_{u_1}\left[\left\|\frac{F(x+\mu_1u_1,y)-F(x,y)}{\mu_1/d_1}u_1\right\|^2\right]\\
&\quad+\frac{3d_1^2}{\mu_1^2}\mathbb{E}_{u_1,\mathcal{B}_{\mu_1}}\left[\left\|\frac{1}{b}\sum_{j=1}^b(f(x+\mu_1u_1,y;\xi_{1,j})-F(x+\mu_1u_1,y))\right\|^2\right]\\
&\quad+\frac{3d_1^2}{\mu_1^2}\mathbb{E}_{u_1,\mathcal{B}}\left[\left\|\frac{1}{b}\sum_{j=1}^b(f(x,y;\tilde{\xi}_j)-F(x,y))\right\|^2\right]
\end{aligned}$$

$$\leq \left(\frac{3}{2}\mu_1^2 L_f^2 d_1^2 + 6d_1\|\nabla_x F(x,y)\|^2\right) + \frac{3\sigma^2 d_1^2}{\mu_1^2 b} + \frac{3\sigma^2 d_1^2}{\mu_1^2 b}$$
$$= \frac{3}{2}\mu_1^2 L_f^2 d_1^2 + 6d_1\|\nabla_x F(x,y)\|^2 + \frac{6\sigma^2 d_1^2}{\mu_1^2 b},$$

where the first inequality follows from the fact that $\|a+b+c\|^2 \leq 3\|a\|^2 + 3\|b\|^2 + 3\|c\|^2$, and the second inequality follows from Lemma 1 and Assumption 2. Then, we complete the proof of (16).

Similarly, we can prove (17) through an analogous derivation. This proof is complete. □

At each iteration t, we denote the combined search directions and sample realizations as:

$$\boldsymbol{u}^t = (u_1^t, u_2^t), \quad \boldsymbol{\xi}^t = (\tilde{\xi}^t, \xi_1^t, \xi_2^t), \quad \text{and} \quad \zeta_t = \{\boldsymbol{u}^t, \boldsymbol{\xi}^t\}. \tag{18}$$

Furthermore, we define $\zeta_{[0,t]} = \{\zeta_0, \zeta_1, \ldots, \zeta_t\}$ as the collection of all random variables up to iteration t.

Lemma 3. Let $\{(x_t, y_t)\}_{t=0}^T$ be the sequence generated by Algorithm 1, $\delta_t := \mathbb{E}[\|y^*(x_t) - y_t\|^2]$. Suppose that Assumptions 2 and 3 hold, and $\eta_y = \min(\frac{1}{6L_f d_2}, \frac{4}{3\tau})$. Then, for any $t \geq 1$:

$$\delta_t \leq \gamma_1^t \delta_0 + \frac{48d_1}{\eta_y \tau}\kappa^2 \eta_x^2 \sum_{j=0}^{t-1} \gamma_1^{t-1-j} \mathbb{E}\left[\|\nabla\Phi(x_j)\|^2\right] + C_1 \sum_{j=0}^{t-1} \gamma_1^{t-1-j}, \tag{19}$$

where

$$\gamma_1 := 1 - \frac{\eta_y \tau}{4} + \frac{48d_1}{\eta_y \tau}\kappa^2 \eta_x^2 L_f^2,$$
$$C_1 := \frac{4\eta_y \mu_2^2 L_f^2}{\tau} + 3\mu_2^2 d_2^2 L_f^2 \eta_y^2 + \frac{12\sigma^2 d_2^2 \eta_y^2}{\mu_2^2 b} + \frac{6\kappa^2 \eta_x^2 \mu_1^2 d_1^2 L_f^2}{\eta_y \tau} + \frac{24\kappa^2 \eta_x^2 \sigma^2 d_1^2}{\eta_y \tau \mu_1^2 b}.$$

Proof. By the definition of δ_t, we have

$$\begin{aligned}
\delta_t &= \mathbb{E}\left[\|y^*(x_t) - y^*(x_{t-1}) + y^*(x_{t-1}) - y_t\|^2\right] \\
&\leq (1+\lambda)\mathbb{E}\left[\|y^*(x_t) - y^*(x_{t-1})\|^2\right] + \left(1 + \frac{1}{\lambda}\right)\mathbb{E}\left[\|y_t - y^*(x_{t-1})\|^2\right] \\
&\leq (1+\lambda)\kappa^2 \eta_x^2 \mathbb{E}\left[\|G_{t-1}\|^2\right] + \left(1 + \frac{1}{\lambda}\right)\mathbb{E}\left[\|y_t - y^*(x_{t-1})\|^2\right],
\end{aligned} \tag{20}$$

where $G_{t-1} := G_{\mu_1}(x_{t-1}, y_{t-1}, u_1^{t-1}; \mathcal{B}_{t-1}, \mathcal{B}_{t-1,\mu_1})$, the first inequality holds for any $\lambda > 0$ due to Young's inequality, the second inequality follows from the κ-Lipschitz continuity of $y^*(\cdot)$ (see, e.g., [22, Lemma 16]) and the update rule.

For the second term on the right-hand side of (20), we have

$$\begin{aligned}
&\mathbb{E}\left[\|y_t - y^*(x_{t-1})\|^2\right] \\
&\leq \mathbb{E}\left[\|y_{t-1} + \eta_y H_{t-1} - y^*(x_{t-1})\|^2\right] \\
&= \mathbb{E}\left[\|y_{t-1} - y^*(x_{t-1})\|^2 + 2\eta_y (y_{t-1} - y^*(x_{t-1}))^\top H_{t-1} + \eta_y^2 \|H_{t-1}\|^2\right] \\
&= \delta_{t-1} + 2\eta_y \mathbb{E}_{\zeta_{[0,t-2]}}\left[(y_{t-1} - y^*(x_{t-1}))^\top \mathbb{E}_{\zeta_{t-1}}\left[H_{t-1} \mid \zeta_{[0,t-2]}\right]\right] + \eta_y^2 \mathbb{E}[\|H_{t-1}\|^2]
\end{aligned}$$

$$
\begin{aligned}
&= \delta_{t-1} + 2\eta_y \mathbb{E}\left[(y_{t-1} - y^*(x_{t-1}))^\top (\nabla_y F(x_{t-1}, y_{t-1}) + B_{t-1})\right] + \eta_y^2 \mathbb{E}[\|H_{t-1}\|^2] \\
&\leq \delta_{t-1} - 2\eta_y \mathbb{E}\left[(y^*(x_{t-1}) - y_{t-1})^\top \nabla_y F(x_{t-1}, y_{t-1})\right] \\
&\quad + 2\eta_y \mathbb{E}\left[\|y_{t-1} - y^*(x_{t-1})\| \cdot \|B_{t-1}\|\right] + \eta_y^2 \left(\frac{3}{2}\mu_2^2 d_2^2 L_f^2 + 6d_2 \|\nabla_y F(x_{t-1}, y_{t-1})\|^2 + \frac{6\sigma^2 d_2^2}{\mu_2^2 b}\right) \\
&\leq \delta_{t-1} - 2\eta_y \mathbb{E}\left[\frac{\tau}{2}\|y^*(x_{t-1}) - y_{t-1}\|^2 + \frac{1}{2L_f}\|\nabla_y F(x_{t-1}, y_{t-1})\|^2\right] \\
&\quad + \frac{\eta_y \tau}{2}\delta_{t-1} + \frac{2\eta_y \mu_2^2 L_f^2}{\tau} + 6d_2 \eta_y^2 \mathbb{E}\left[\|\nabla_y F(x_{t-1}, y_{t-1})\|^2\right] + \frac{3}{2}\mu_2^2 d_2^2 L_f^2 \eta_y^2 + \frac{6\sigma^2 d_2^2 \eta_y^2}{\mu_2^2 b} \\
&\leq \left(1 - \frac{\eta_y \tau}{2}\right)\delta_{t-1} + \frac{2\eta_y \mu_2^2 L_f^2}{\tau} + \frac{3}{2}\mu_2^2 d_2^2 L_f^2 \eta_y^2 + \frac{6\sigma^2 d_2^2 \eta_y^2}{\mu_2^2 b},
\end{aligned}
$$

where $H_{t-1} := H_{\mu_2}(x_{t-1}, y_{t-1}, u_2^{t-1}; \mathcal{B}_{t-1}, \mathcal{B}_{t-1,\mu_2})$, $B_{t-1} := \nabla_y F_{\mu_2}(x_{t-1}, y_{t-1}) - \nabla_y F(x_{t-1}, y_{t-1})$ represents the smoothing bias bounded as $\|B_{t-1}\| \leq \mu_2 L_f$ by [1, (2.35)], and the third inequality follows from the co-coercivity of strongly concave functions.

Then,

$$
\begin{aligned}
\delta_t &\leq \left(1 + \frac{1}{\lambda}\right)\left[\left(1 - \frac{\eta_y \tau}{2}\right)\delta_{t-1} + \frac{2\eta_y \mu_2^2 L_f^2}{\tau} + \frac{3}{2}\mu_2^2 d_2^2 L_f^2 \eta_y^2 + \frac{6\sigma^2 d_2^2 \eta_y^2}{\mu_2^2 b}\right] \\
&\quad + (1+\lambda)\kappa^2 \eta_x^2 \mathbb{E}\left[\frac{3}{2}\mu_1^2 d_1^2 L_f^2 + 12 d_1 \|\nabla \Phi(x_{t-1})\|^2 + 12 d_1 L_f^2 \|y^*(x_{t-1}) - y_{t-1}\|^2 + \frac{6\sigma^2 d_1^2}{\mu_1^2 b}\right] \\
&\leq \left(1 - \frac{\eta_y \tau}{4}\right)\delta_{t-1} + \frac{4\eta_y \mu_2^2 L_f^2}{\tau} + 3\mu_2^2 d_2^2 L_f^2 \eta_y^2 + \frac{12\sigma^2 d_2^2 \eta_y^2}{\mu_2^2 b} + \frac{6\kappa^2 \eta_x^2 \mu_1^2 d_1^2 L_f^2}{\eta_y \tau} \\
&\quad + \frac{48 d_1 \kappa^2 \eta_x^2}{\eta_y \tau}\mathbb{E}\left[\|\nabla \Phi(x_{t-1})\|^2\right] + \frac{48 d_1 \kappa^2 \eta_x^2 L_f^2}{\eta_y \tau}\delta_{t-1} + \frac{24\kappa^2 \eta_x^2 \sigma^2 d_1^2}{\eta_y \tau \mu_1^2 b} \\
&= \gamma_1 \delta_{t-1} + \frac{48 d_1}{\eta_y \tau}\kappa^2 \eta_x^2 \mathbb{E}\left[\|\nabla \Phi(x_{t-1})\|^2\right] + C_1 \\
&\leq \gamma_1^t D^2 + \frac{48 d_1}{\eta_y \tau}\kappa^2 \eta_x^2 \sum_{j=0}^{t-1} \gamma_1^{t-1-j} \mathbb{E}\left[\|\nabla \Phi(x_j)\|^2\right] + C_1 \sum_{j=0}^{t-1} \gamma_1^{t-1-j},
\end{aligned}
$$

where the second inequality follows by choosing $\lambda = \frac{4 - 2\eta_y \tau}{\eta_y \tau}$ and the condition $\eta_y \leq \frac{4}{3\tau}$. The proof is complete. □

Theorem 1 (Convergence Results)**.** Let $\{(x_t, y_t)\}_{t=0}^T$ be the sequence generated by Algorithm 1, and $d := d_1 + d_2$. Suppose that Assumptions 1-4 hold. Then for any $\epsilon > 0$, the following statements hold.

(i) **Accuracy-Independent Constant Step-size:** Suppose that the parameters are chosen such that

   (a) the smoothing parameters $\mu_1 = \Theta(\epsilon)$ and $\mu_2 = \Theta(\epsilon \kappa^{-1})$;

   (b) the step-sizes $\eta_x = \Theta(d^{-3}\kappa^{-3})$ and $\eta_y = \Theta(d^{-2}\kappa^{-1})$ such that $\eta_x \leq \min\left(\frac{\eta_y \tau}{8\sqrt{6 d_1 \kappa L_f}}, \frac{1}{72 d_1 L_g}\right)$ and $\eta_y \leq \min\left(\frac{1}{6 L_f d_2}, \frac{4}{3\tau}\right)$;

(c) the batch size $b = \Theta(\epsilon^{-4}\kappa^2)$.

Then, SD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-2})$ and a sample complexity of $\mathcal{O}(d^3\kappa^5\epsilon^{-6})$.

(ii) **Accuracy-Dependent Step-size:** Suppose that the parameters are chosen such that

(a) the smoothing parameters $\mu_1 = \Theta(\epsilon)$ and $\mu_2 = \Theta(\epsilon\kappa^{-1})$;

(b) the step-sizes $\eta_x = \Theta(\epsilon^3 d^{-3}\kappa^{-3})$ and $\eta_y = \Theta(\epsilon^2 d^{-2}\kappa^{-1})$ such that $\eta_x \leq \min\left(\frac{\eta_y \tau}{8\sqrt{6d_1}\kappa L_f}, \frac{1}{72 d_1 L_g}\right)$ and $\eta_y \leq \min\left(\frac{1}{6L_f d_2}, \frac{4}{3\tau}\right)$;

(c) the batch size $b = \Theta(\epsilon^{-2}\kappa^2)$.

Then, SD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-5})$ and a sample complexity of $\mathcal{O}(d^3\kappa^5\epsilon^{-7})$.

Proof. By [21, Lemma 16], Φ is L-smooth with $L := 2\kappa L_f$. Then,

$$\begin{aligned}
\Phi(x_{t+1}) &\leq \Phi(x_t) + \nabla\Phi(x_t)^\top (x_{t+1} - x_t) + \frac{L}{2}\|x_{t+1} - x_t\|^2 \\
&= \Phi(x_t) - \eta_x \left\langle \nabla\Phi(x_t), G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\right\rangle + \frac{L}{2}\eta_x^2 \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2 \\
&= \Phi(x_t) - \eta_x \nabla\Phi(x_t)^\top \nabla_x F_{\mu_1}(x_t, y_t) - \eta_x \nabla\Phi(x_t)^\top \Delta_t + \frac{L}{2}\eta_x^2 \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2 \\
&= \Phi(x_t) - \eta_x \nabla\Phi(x_t)^\top (\nabla_x F_{\mu_1}(x_t, y_t) - \nabla_x F(x_t, y_t)) - \eta_x \nabla\Phi(x_t)(\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)) \\
&\quad - \eta_x \|\nabla\Phi(x_t)\|^2 - \eta_x \nabla\Phi(x_t)^\top \Delta_t + \frac{L}{2}\eta_x^2 \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2,
\end{aligned}$$

where $\Delta_t := G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t) - \nabla_x F_{\mu_1}(x_t, y_t)$.

Rearranging the terms in the above inequality yields,

$$\begin{aligned}
\eta_x \|\nabla\Phi(x_t)\|^2 &\leq \Phi(x_t) - \Phi(x_{t+1}) - \eta_x \nabla\Phi(x_t)^\top (\nabla_x F_{\mu_1}(x_t, y_t) - \nabla_x F(x_t, y_t)) \\
&\quad - \eta_x \nabla\Phi(x_t)^\top (\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)) - \eta_x \nabla\Phi(x_t)^\top \Delta_t \\
&\quad + \frac{L}{2}\eta_x^2 \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2 \\
&\leq \Phi(x_t) - \Phi(x_{t+1}) + \frac{\eta_x}{4}\|\nabla\Phi(x_t)\|^2 + \eta_x \|\nabla_x F_{\mu_1}(x_t, y_t) - \nabla_x F(x_t, y_t)\|^2 \\
&\quad + \frac{\eta_x}{4}\|\nabla\Phi(x_t)\|^2 + \eta_x \|\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)\|^2 - \eta_x \nabla\Phi(x_t)^\top \Delta_t \\
&\quad + \frac{L}{2}\eta_x^2 \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2,
\end{aligned}$$

where the second inequality follows from Young's inequality. Summing up the above inequalities for $0 \leq t \leq T$, we obtain

$$\begin{aligned}
\frac{\eta_x}{2}\sum_{t=0}^{T} \|\nabla\Phi(x_t)\|^2 &\leq \Phi(x_0) - \Phi(x^*) + \eta_x \sum_{t=0}^{T} \|\nabla_x F_{\mu_1}(x_t, y_t) - \nabla_x F(x_t, y_t)\|^2 \\
&\quad + \eta_x \sum_{t=0}^{T} \|\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)\|^2 - \eta_x \sum_{t=0}^{T} \nabla\Phi(x_t)^\top \Delta_t
\end{aligned}$$

$$
\begin{aligned}
&+ \frac{L}{2}\eta_x^2 \sum_{t=0}^{T} \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2 \\
&\leq \Phi(x_0) - \Phi(x^*) + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^{T} \|y^*(x_t) - y_t\|^2 \\
&- \eta_x \sum_{t=0}^{T} \nabla\Phi(x_t)^\top \Delta_t + \frac{L}{2}\eta_x^2 \sum_{t=0}^{T} \|G_{\mu_1}(x_t, y_t, u_1^t, \mathcal{B}^t, \mathcal{B}_1^t)\|^2.
\end{aligned}
$$

Taking the expectation with respect to $u_1, \tilde{\xi}, \xi_1$ on both sides of the above inequality, we have

$$
\begin{aligned}
&\frac{\eta_x}{2} \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|\nabla\Phi(x_t)\|^2\right] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|y^*(x_t) - y_t\|^2\right] \\
&\quad - \eta_x \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,t-1]} \left[\nabla\Phi(x_t)^\top \mathbb{E}_{\zeta_t}[\Delta_t \mid \zeta_{[0,t-1]}]\right] \\
&\quad + \frac{L}{2}\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\frac{3}{2}\mu_1^2 L_f^2 d_1^2 + 6d_1 \|\nabla_x F(x_t, y_t)\|^2 + \frac{6\sigma^2 d_1^2}{\mu_1^2 b}\right] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|y^*(x_t) - y_t\|^2\right] \\
&\quad + 3Ld_1\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|\nabla_x F(x_t, y_t) - \nabla\Phi(x_t) + \nabla\Phi(x_t)\|^2\right] \\
&\quad + \frac{3L\eta_x^2 \mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2 \sigma^2 d_1^2}{\mu_1^2 b}(T+1) \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + (\eta_x L_f^2 + 6Ld_1\eta_x^2 L_f^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|y^*(x_t) - y_t\|^2\right] \\
&\quad + 6Ld_1\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|\nabla\Phi(x_t)\|^2\right] + \frac{3L\eta_x^2 \mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2 \sigma^2 d_1^2}{\mu_1^2 b}(T+1).
\end{aligned}
$$

Rearranging the terms in the above inequality yields,

$$
\begin{aligned}
&\left(\frac{\eta_x}{2} - 6Ld_1\eta_x^2\right) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|\nabla\Phi(x_t)\|^2\right] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 d_1^2 (T+1) + (\eta_x L_f^2 + 6Ld_1\eta_x^2 L_f^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]} \left[\|y^*(x_t) - y_t\|^2\right] \\
&\quad + \frac{3L\eta_x^2 \mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2 \sigma^2 d_1^2}{\mu_1^2 b}(T+1) \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1)
\end{aligned}
$$

$$
\begin{aligned}
&+ (\eta_x L_f^2 + 6Ld_1\eta_x^2 L_f^2)\left(D^2\sum_{t=0}^{T}\gamma_1^t + \frac{48d_1}{\eta_y\tau}\kappa^2\eta_x^2\sum_{t=0}^{T}\sum_{j=0}^{t-1}\gamma_1^{t-1-j}\mathbb{E}\left[\|\nabla\Phi(x_j)\|^2\right] + C_1\sum_{t=0}^{T}\sum_{j=0}^{t-1}\gamma_1^{t-1-j}\right)\\
&+ \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 b}(T+1)\\
&\le \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x\mu_1^2 L_f^2(T+1)\\
&+ (\eta_x L_f^2 + 6Ld_1\eta_x^2 L_f^2)\left(\frac{8D^2}{\eta_y\tau} + \frac{384d_1}{\eta_y^2\tau^2}\kappa^2\eta_x^2\sum_{t=0}^{T}\mathbb{E}\left[\|\nabla\Phi(x_t)\|^2\right] + \frac{8C_1}{\eta_y\tau}(T+1)\right)\\
&+ \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 b}(T+1),
\end{aligned}
$$

where the third inequality follows from $\sum_{t=0}^{\infty}\gamma_1^t \le \frac{8}{\eta_y\tau}$.

By further rearranging the terms, we obtain

$$
\begin{aligned}
&\left(\frac{\eta_x}{2} - 6Ld_1\eta_x^2 - \frac{384d_1}{\eta_y^2\tau^2}\kappa^2\eta_x^3 L_f^2 - \frac{2304}{\eta_y^2\tau^2}\kappa^2\eta_x^4 L d_1^2 L_f^2\right)\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}\left[\|\nabla\Phi(x_t)\|^2\right]\\
&\le \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x\mu_1^2 L_f^2(T+1)\\
&\quad + (\eta_x L_f^2 + 6Ld_1\eta_x^2 L_f^2)\left(\frac{8D^2}{\eta_y\tau} + \frac{8C_1}{\eta_y\tau}(T+1)\right)\\
&\quad + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 b}(T+1).
\end{aligned}
$$

Dividing both sides of the above inequality by

$$
(\frac{\eta_x}{2} - 6Ld_1\eta_x^2 - \frac{384d_1}{\eta_y^2\tau^2}\kappa^2\eta_x^3 L_f^2 - \frac{2304}{\eta_y^2\tau^2}\kappa^2\eta_x^4 L d_1^2 L_f^2)(T+1),
$$

and applying the parameter constraint $\eta_x \le \min\left(\frac{\eta_y\tau}{8\sqrt{6d_1}\kappa L_f}, \frac{1}{72d_1 L}\right)$, we obtain

$$
\begin{aligned}
&\frac{1}{T+1}\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}[\|\nabla\Phi(x_t)\|^2]\\
&\le \frac{4}{(T+1)\eta_x}\mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \frac{(L_f^2 + 6Ld_1\eta_x^2 L_f^2)32D^2}{(T+1)\eta_y\tau} + 4\mu_1^2 L_f^2\\
&\quad + (L_f^2 + 6Ld_1\eta_x L_f^2)\frac{32C_1}{\eta_y\tau} + 3L\eta_x\mu_1^2 L_f^2 d_1^2 + \frac{12L\eta_x\sigma^2 d_1^2}{\mu_1^2 b},
\end{aligned}
$$

which implies

$$
\begin{aligned}
&\frac{1}{T+1}\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}\left[\|\nabla\Phi(x_t)\|^2\right]\\
&= \mathcal{O}\left((T+1)^{-1}\eta_x^{-1}\right) + \mathcal{O}\left(\kappa(T+1)^{-1}\eta_y^{-1}\right) + \mathcal{O}\left(d_1\kappa^2(T+1)^{-1}\eta_y^{-1}\eta_x^2\right) + \mathcal{O}\left(\mu_1^2\right)\\
&\quad + \mathcal{O}\left(\mu_2^2\kappa^2\right) + \mathcal{O}\left(\mu_2^2 d_2^2\eta_y\kappa\right) + \mathcal{O}\left(d_2^2\eta_y\kappa\mu_2^{-2}b^{-1}\right) + \mathcal{O}\left(\kappa^4\eta_x^2\mu_1^2 d_1^2\eta_y^{-2}\right) + \mathcal{O}\left(\kappa^4\eta_x^2 d_1^2\eta_y^{-2}\mu_1^{-2}b^{-1}\right)\\
&\quad + \mathcal{O}\left(d_1\eta_x\mu_2^2\kappa^3\right) + \mathcal{O}\left(d_1\eta_x\mu_2^2 d_2^2\eta_y\kappa^2\right) + \mathcal{O}\left(d_1\eta_x d_2^2\eta_y\kappa^2\mu_2^{-2}b^{-1}\right) + \mathcal{O}\left(\kappa^5\eta_x^3\mu_1^2 d_1^3\eta_y^{-2}\right)
\end{aligned}
$$

$$+\mathcal{O}\left(\kappa^5\eta_x^3 d_1^3\eta_y^{-2}\mu_1^{-2}b^{-1}\right)+\mathcal{O}\left(\eta_x\kappa\mu_1^2 d_1^2\right)+\mathcal{O}\left(\eta_x\kappa d_1^2\mu_1^{-2}b^{-1}\right).$$

For the accuracy-independent constant step-sizes setting, $\mu_1=\Theta(\epsilon)$, $\mu_2=\Theta(\epsilon\kappa^{-1})$, $\eta_x=\Theta(d^{-3}\kappa^{-3})$, $\eta_y=\Theta(d^{-2}\kappa^{-1})$, $b=\Theta(\epsilon^{-4}\kappa^2)$, we have

$$\frac{1}{T+1}\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}\left[\|\nabla\Phi(x_t)\|^2\right]\leq\mathcal{O}(\epsilon^2),$$

if $T=\Theta(\epsilon^{-2}d^3\kappa^3)$. This means that, SD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-2})$, and a sample complexity of $\mathcal{O}(d^3\kappa^5\epsilon^{-6})$.

On the other hand, for the accuracy-dependent step-sizes setting, $\mu_1=\Theta(\epsilon)$, $\mu_2=\Theta(\epsilon\kappa^{-1})$, $\eta_x=\Theta(\epsilon^3d^{-3}\kappa^{-3})$, $\eta_y=\Theta(\epsilon^2d^{-2}\kappa^{-1})$, $b=\Theta(\epsilon^{-2}\kappa^2)$, we have

$$\frac{1}{T+1}\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}\left[\|\nabla\Phi(x_t)\|^2\right]\leq\mathcal{O}(\epsilon^2).$$

If $T=\Theta(\epsilon^{-5}d^3\kappa^3)$. This means that, SD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(d^3\kappa^3\epsilon^{-5})$, and a sample complexity of $\mathcal{O}(d^3\kappa^5\epsilon^{-7})$. The proof is complete. □

Remark 1. The convergence results in Theorem 1 demonstrate a fundamental trade-off between different step-size settings. Specifically, under the accuracy-independent constant step-sizes, SD-ZO-SGDA achieves a fast iteration complexity of $\mathcal{O}(\epsilon^{-2})$, though it requires a large batch size $b=\mathcal{O}(\epsilon^{-4})$ to suppress estimation variance. In contrast, the accuracy-dependent step-sizes relaxes the per-iteration requirement to $b=\mathcal{O}(\epsilon^{-2})$, making the algorithm significantly more practical for memory-constrained or high-frequency sampling scenarios, but at the cost of a higher total sample complexity of $\mathcal{O}(\epsilon^{-7})$.

Furthermore, comparing our results with existing literature, the established sample complexity of $\mathcal{O}(d^3\kappa^5\epsilon^{-6})$ involves an additional ϵ^{-2} factor relative to standard zeroth-order minimax algorithms designed for static distributions, such as ZO-SGDA [32], which typically achieves $\mathcal{O}(\epsilon^{-4})$. This complexity gap arises from the irreducible variance induced by the coupling between the decision variables and the data-generating mechanism. Unlike static environments where common random numbers can be employed to effectively cancel out noise during finite-difference estimation, the decision-dependent shift in SMDD prevents such cancellation, which necessitates a more careful balance between the smoothing parameter μ and the batch size b to ensure robust convergence.

4 MD-ZO-SGDA: Algorithm and Convergence Analysis

This section is devoted to the study of MD-ZO-SGDA. We first present the detailed steps of MD-ZO-SGDA in Section 4.1, followed by a rigorous establishment of its convergence theory in Section 4.2.

4.1 MD-ZO-SGDA algorithm

Building upon the foundational framework of SD-ZO-SGDA, we now introduce its multi-directional variant MD-ZO-SGDA, whose full procedure is summarized in Algorithm 2.

The overall structure of MD-ZO-SGDA inherits the alternating "sample-estimate-update" process from Algorithm 1, yet the key of algorithm lies in the structure of the zeroth-order gradient estimators. Specifically, instead of relying on a single random perturbation direction per iteration, MD-ZO-SGDA draws m independent random directions. Accordingly, as outlined in Step 4, the algorithm queries the objective function by drawing sample sets along all m directions from the shifted distributions. These samples are then utilized to construct the multi-directional zeroth-order gradient estimators defined in (12) and (13). Finally, the decision variables x and y are updated using these newly constructed multi-directional estimators, following the standard stochastic gradient descent ascent rules.

Algorithm 2 Multi-directional Zeroth-Order Stochastic Gradient Descent Ascent (MD-ZO-SGDA)

1: **Input**: iteration limit $T > 0$, initial input $x_0 \in \mathbb{R}^{d_1}, y_0 \in \mathcal{Y}$, step sizes (η_1, η_2), smoothing parameters $\mu_1, \mu_2 > 0$ and number of directions m.

2: **Initialization**: Draw m random vectors $\{u_{1,i}^0\}_{i=1}^m$ and $\{u_{2,i}^0\}_{i=1}^m$ from uniform distribution over unit spheres $\mathcal{U}_{Sp_1}$ and $\mathcal{U}_{Sp_2}$. Draw sample sets $\tilde{\xi}_M^0 = \{\tilde{\xi}_i^0\}_{i=1}^m$ from $D(x_0, y_0)$, $\xi_{1,M}^0 = \{\xi_{1,i}^0\}_{i=1}^m$ from $D(x_0 + \mu_1 u_{1,i}^0, y_0)$ and $\xi_{2,M}^0 = \{\xi_{2,i}^0\}_{i=1}^m$ from $D(x_0, y_0 + \mu_2 u_{2,i}^0)$. Compute initial gradients $G_{\mu_1}(x_0, y_0, u_{M_1}^0; \tilde{\xi}_M^0, \xi_{1,M}^0)$ and $H_{\mu_2}(x_0, y_0, u_{M_2}^0; \tilde{\xi}_M^0, \xi_{2,M}^0)$, where the zeroth-order gradients are estimated from (12) and (13).

3: **For** $t = 1, 2, \cdots, T$ **do**

4: Draw m random vectors $\{u_{1,i}^t\}_{i=1}^m \sim \mathcal{U}_{Sp_1}$ and $\{u_{2,i}^t\}_{i=1}^m \sim \mathcal{U}_{Sp_2}$.
Draw sample sets $\tilde{\xi}_M^t = \{\tilde{\xi}_i^t\}_{i=1}^m$ from $D(x_t, y_t)$, $\xi_{1,M}^t = \{\xi_{1,i}^t\}_{i=1}^m$ from $D(x_t + \mu_1 u_{1,i}^t, y_t)$ and $\xi_{2,M}^t = \{\xi_{2,i}^t\}_{i=1}^m$ from $D(x_t, y_t + \mu_2 u_{2,i}^t)$;

5: Calculate $G_{\mu_1}(x_t, y_t, u_{M_1}^t; \tilde{\xi}_M^t, \xi_{1,M}^t)$, $H_{\mu_2}(x_t, y_t, u_{M_2}^t; \tilde{\xi}_M^t, \xi_{2,M}^t)$;

6: Update $x_{t+1} = x_t - \eta_1 G_{\mu_1}(x_t, y_t, u_{M_1}^t; \tilde{\xi}_M^t, \xi_{1,M}^t)$;

7: Update $y_{t+1} = \mathcal{P}_{\mathcal{Y}}(y_t + \eta_y H_{\mu_2}(x_t, y_t, u_{M_2}^t; \tilde{\xi}_M^t, \xi_{2,M}^t))$;

8: **end for**

4.2 Convergence analysis of MD-ZO-SGDA

We now present the rigorous convergence analysis of MD-ZO-SGDA. Compared to the single-directional baseline established in Section 3, the utilization of m independent perturbation directions not only reduces the error in random directions but also reduces the sample noise due to finite samples by resampling ξ in each direction.

Next, we present the technical lemmas required for analyzing multi-directional estimators and conclude with the main convergence theorem for MD-ZO-SGDA.

Lemma 4. Suppose that Assumption 2 holds. Then, for any $x \in \mathbb{R}^{d_1}, y \in \mathcal{Y}$ and any $\mu_1, \mu_2 > 0$,

$$\mathbb{E}_{u_1 \sim \mathcal{U}_{Sp_1}} \left[\left\| G^M_{\mu_1}(x, y, u_1) \right\|^2 \right] \leq \frac{\mu_1^2}{2m} L_f^2 d_1^2 + \frac{6d_1}{m} \|\nabla_x F(x, y)\|^2, \tag{21}$$

$$\mathbb{E}_{u_2 \sim \mathcal{U}_{Sp_2}} \left[\left\| H^M_{\mu_2}(x, y, u_2) \right\|^2 \right] \leq \frac{\mu_2^2}{2m} L_f^2 d_2^2 + \frac{6d_2}{m} \|\nabla_y F(x, y)\|^2. \tag{22}$$

Proof. By the definition of $G^M_{\mu_1}(x, y, u_1)$, we have

$$\begin{aligned}
&\mathbb{E}_{u_1 \sim \mathcal{U}_1} \left[\|G^M_{\mu_1}(x, y, u_1)\|^2 \right] \\
&= \operatorname{tr} \left(\frac{d_1^2}{m} \mathbb{E}_{u_1} \left[\left(\frac{F(x + \mu_1 u_1, y) - F(x, y) - \mu_1 \nabla_x F(x, y)^\top u_1}{\mu_1} + \nabla_x F(x, y)^\top u_1 \right)^2 u_1 u_1^\top \right] \right) \\
&\leq \operatorname{tr} \left(\frac{d_1^2}{m} \mathbb{E}_{u_1} \left[2 \left(\frac{F(x + \mu_1 u_1, y) - F(x, y) - \mu_1 \nabla_x F(x, y)^\top u_1}{\mu_1} \right)^2 u_1 u_1^\top + 2 \left(\nabla_x F(x, y)^\top u_1 \right)^2 u_1 u_1^\top \right] \right) \\
&\leq \operatorname{tr} \left(\frac{2d_1^2}{m} \mathbb{E}_{u_1} \left[\frac{L_f^2 \mu_1^2 \|u_1\|^4}{4} u_1 u_1^\top + \left(\nabla_x F(x, y)^T u_1 \right)^2 u_1 u_1^\top \right] \right) \\
&= \operatorname{tr} \left(\frac{2d_1^2}{m} \left(\frac{L_f^2 \mu_1^2}{4d_1} I_{d_1} + \frac{\nabla_x F(x, y)^\top \nabla_x F(x, y)}{d_1(d_1 + 2)} I_{d_1} + \frac{2}{d_1(d_1 + 2)} \nabla_x F(x, y) \nabla_x F(x, y)^\top \right) \right) \\
&\leq \operatorname{tr} \left(\frac{2}{m} \left(\frac{L_f^2 \mu_1^2 d_1}{4} + \frac{3d_1}{d_1 + 2} \|\nabla_x F(x, y)\|^2 \right) I \right) \\
&\leq \frac{\mu_1^2}{2m} L_f^2 d_1^2 + \frac{6d_1}{m} \|\nabla_x F(x, y)\|^2,
\end{aligned}$$

where the second inequality is due to Assumption 3, the second equality follows from the fact that $\mathbb{E}[uu^\top] = \frac{1}{d} I_d$ and $\mathbb{E}[(a^\top u)^2 u u^\top] = \frac{\|a\|^2 I + 2aa^\top}{d(d+2)}$, the third inequality is due to the fact that $\nabla_x F(x, y) \nabla_x F(x, y)^\top \preceq \|\nabla_x F(x, y)\|^2 I_{d_1}$. Then, we complete the proof of (21).

Similarly, we can prove (22) through an analogous derivation. This proof is complete. □

Lemma 5. Suppose that Assumptions 2 and 3 hold. Then, for any $x \in \mathbb{R}^{d_1}$, $y \in \mathcal{Y}$,

$$\mathbb{E}_{u_1} \left[\mathbb{E}_{\tilde{\xi}_M, \xi_{1,M}} \left[\left\| G_{\mu_1}(x, y, u_{M_1}, \tilde{\xi}_M, \xi_{1,M}) \right\|^2 \right] \right] \leq \frac{3}{2m} \mu_1^2 d_1^2 L_f^2 + \frac{18 d_1}{m} \|\nabla_x F(x, y)\|^2 + \frac{6\sigma^2 d_1^2}{m \mu_1^2}, \tag{23}$$

$$\mathbb{E}_{u_2} \left[\mathbb{E}_{\tilde{\xi}_M, \xi_{2,M}} \left[\left\| H_{\mu_2}(x, y, u_{M_2}, \tilde{\xi}_M, \xi_{2,M}) \right\|^2 \right] \right] \leq \frac{3}{2m} \mu_2^2 d_2^2 L_f^2 + \frac{18 d_2}{m} \|\nabla_y F(x, y)\|^2 + \frac{6\sigma^2 d_2^2}{m \mu_2^2}, \tag{24}$$

where $\tilde{\xi}_M = \{\tilde{\xi}_i\}_{i=1}^m$, $\xi_{1,M} = \{\xi_{1,i}\}_{i=1}^m$, and $\xi_{2,M} = \{\xi_{2,i}\}_{i=1}^m$ denote the sets of sample realizations drawn independently from $D(x, y)$, $D(x + \mu_1 u_{1,i}, y)$, and $D(x, y + \mu_2 u_{2,i})$, respectively.

Proof. By the definition of the multi-directional estimator (12), we have

$$\mathbb{E}_{u_1, \tilde{\xi}_M, \xi_{1,M}} \left[\|G_{\mu_1}(x, y, u_{M_1}, \tilde{\xi}_M, \xi_{1,M})\|^2 \right]$$

$$
\begin{aligned}
&= \mathbb{E}_{u_1,\xi_1,\tilde{\xi}}\left[\left\|\frac{d_1}{m}\sum_{i=1}^{m}\frac{F(x+\mu_1 u_{1,i},y)-F(x,y)}{\mu_1}u_{1,i}\right.\right.\\
&\quad\left.\left.+\frac{d_1}{m}\sum_{i=1}^{m}\frac{f(x+\mu_1 u_{1,i},y,\xi_{1,i})-F(x+\mu_1 u_{1,i},y)}{\mu_1}u_{1,i}+\frac{d_1}{m}\sum_{i=1}^{m}\frac{F(x,y)-f(x,y,\tilde{\xi}_i)}{\mu_1}u_{1,i}\right\|^2\right]\\
&\le 3\mathbb{E}_{u_1}\left\|\frac{d_1}{m}\sum_{i=1}^{m}\frac{F(x+\mu_1 u_{1,i},y)-F(x,y)}{\mu_1}u_{1,i}\right\|^2\\
&\quad+\frac{3d_1^2}{m^2\mu_1^2}\mathbb{E}_{u_1,\xi_1}\left\|\sum_{i=1}^{m}\Delta f(x+\mu_1 u_{1,i},y,\xi_{1,i})u_{1,i}\right\|^2+\frac{3d_1^2}{m^2\mu_1^2}\mathbb{E}_{u_1,\tilde{\xi}}\left\|\sum_{i=1}^{m}\Delta f(x,y,\tilde{\xi}_i)u_{1,i}\right\|^2\\
&\le \mathbb{E}_{u_1}\left[\left(\frac{3L_f^2\mu_1^2d_1^2}{2m}+\frac{18d_1}{m}\|\nabla_x F(x,y)\|^2\right)+\frac{3d_1^2}{m^2\mu_1^2}\mathbb{E}_{\xi_1}\left[\sum_{i=1}^{m}\|\Delta f(x+\mu_1 u_{1,i},y,\xi_{1,i})u_{1,i}\|^2\right]\right.\\
&\quad\left.+\frac{3d_1^2}{m^2\mu_1^2}\mathbb{E}_{\tilde{\xi}}\left[\sum_{i=1}^{m}\|\Delta f(x,y,\tilde{\xi}_i)u_{1,i}\|^2\right]\right]\\
&\le \left(\frac{3\mu_1^2L_f^2d_1}{2m}+\frac{18d_1}{m}\|\nabla_x F(x,y)\|^2\right)+\frac{3d_1^2}{m^2\mu_1^2}\sum_{i=1}^{m}\mathbb{E}\left[\sigma^2\|u_{1,i}\|^2\right]+\frac{3d_1^2}{m^2\mu_1^2}\sum_{i=1}^{m}\mathbb{E}\left[\sigma^2\|u_{1,i}\|^2\right]\\
&= \frac{3\mu_1^2d_1^2L_f^2}{2m}+\frac{18d_1^2}{m(d_1+2)}\|\nabla_x F(x,y)\|^2+\frac{6\sigma^2d_1^2}{m\mu_1^2}\\
&\le \frac{3}{2m}\mu_1^2d_1^2L_f^2+\frac{18d_1}{m}\|\nabla_x F(x,y)\|^2+\frac{6\sigma^2d_1^2}{m\mu_1^2},
\end{aligned}
$$

where $\Delta f(x,y,\xi) := f(x,y,\xi)-F(x,y)$. The second inequality is due to the fact that, for $i\neq k$,

$$
\mathbb{E}_{\xi_{1,i},\xi_{1,k}}\left[\left(f(x+\mu u_{1,i},y,\xi_{1,i})-F(x+\mu u_{1,i},y)\right)\left(f(x+\mu u_{1,k},y,\xi_{1,k})-F(x+\mu u_{1,k},y)\right)u_{1,i}^\top u_{1,k}\right]=0,
$$

$$
\mathbb{E}_{\tilde{\xi}_i,\tilde{\xi}_k}\left[\left(f(x,y,\tilde{\xi}_i)-F(x,y)\right)\left(f(x,y,\tilde{\xi}_k)-F(x,y)\right)u_{1,i}^\top u_{1,k}\right]=0.
$$

Then, we complete the proof of (23).

Similarly, we can prove (24) through an analogous derivation. This proof is complete. □

Lemma 6. Let $\{(x_t,y_t)\}_{t=0}^T$ be the sequence generated by Algorithm 2, $\delta_t := \mathbb{E}[\|y^*(x_t)-y_t\|^2]$. Suppose that Assumptions 2 and 3 hold, and $\eta_y = \min(\frac{m}{36L_f d_2},\frac{4}{3\tau})$. Then

$$
\delta_t \le \gamma_2^t\delta_0+\frac{144d_1}{\eta_y\tau m}\kappa^2\eta_x^2\sum_{j=0}^{t-1}\gamma_2^{t-1-j}\mathbb{E}\left[\|\nabla\Phi(x_{t-1})\|^2\right]+C_2\sum_{j=0}^{t-1}\gamma_2^{t-1-j}, \tag{25}
$$

where

$$
\begin{aligned}
\gamma_2 &:= 1-\frac{\eta_y\tau}{4}+\frac{144d_1}{\eta_y\tau m}\kappa^2\eta_x^2L_f^2,\\
C_2 &:= \frac{4\mu_2^2L_f^2\eta_y}{\tau}+\frac{3\mu_2^2d_2^2L_f^2\eta_y^2}{m}+\frac{12\sigma^2d_2^2\eta_y^2}{\mu_2^2m}+\frac{6\kappa^2\eta_x^2\mu_1^2d_1^2L_f^2}{\eta_y\tau m}+\frac{24\kappa^2\eta_x^2\sigma^2d_1^2}{\eta_y\tau\mu_1^2m}.
\end{aligned}
$$

Proof. By the definition of δ_t, we have

$$
\delta_t = \mathbb{E}\left[\|y^*(x_t)-y^*(x_{t-1})+y^*(x_{t-1})-y_t\|^2\right]
$$

$$
\begin{aligned}
&\leq (1+\lambda)\mathbb{E}\left[\|y^*(x_t)-y^*(x_{t-1})\|^2\right]+\left(1+\frac{1}{\lambda}\right)\mathbb{E}\left[\|y_t-y^*(x_{t-1})\|^2\right]\\
&\leq (1+\lambda)\kappa^2\eta_x^2\mathbb{E}\left[\|G'_{t-1}\|^2\right]+\left(1+\frac{1}{\lambda}\right)\mathbb{E}\left[\|y_t-y^*(x_{t-1})\|^2\right], \qquad (26)
\end{aligned}
$$

where $G'_{t-1} := G_{\mu_1}(x_{t-1}, y_{t-1}, u_{M_1}^{t-1}, \tilde{\xi}_M^{t-1}, \xi_{1,M}^{t-1})$, the first inequality holds for any $\lambda > 0$ due to Young's inequality, the second inequality follows from the κ-Lipschitz continuity of $y^*(\cdot)$ (see, e.g., [22, Lemma 16]) and primal update rule.

For the second term on the right-hand side of (26), we have

$$
\begin{aligned}
&\mathbb{E}\left[\|y_t-y^*(x_{t-1})\|^2\right]\\
&\leq \mathbb{E}\left[\|y_{t-1}+\eta_y H'_{t-1}-y^*(x_{t-1})\|^2\right]\\
&= \delta_{t-1}+2\eta_y\mathbb{E}\left[(y_{t-1}-y^*(x_{t-1}))^\top H'_{t-1}\right]+\eta_y^2\mathbb{E}\left[\|H'_{t-1}\|^2\right]\\
&= \delta_{t-1}+2\eta_y\mathbb{E}_{\zeta_{[0,t-2]}}\left[\mathbb{E}_{\zeta_{t-1}}\left[(y_{t-1}-y^*(x_{t-1}))^\top H'_{t-1}\mid \zeta_{[0,t-2]}\right]\right]+\eta_y^2\mathbb{E}\left[\|H'_{t-1}\|^2\right]\\
&= \delta_{t-1}+2\eta_y\mathbb{E}\left[(y_{t-1}-y^*(x_{t-1}))^\top(\nabla_y F(x_{t-1},y_{t-1})+B_{t-1})\right]+\eta_y^2\mathbb{E}[\|H'_{t-1}\|^2]\\
&\leq \delta_{t-1}-2\eta_y\mathbb{E}\left[(y^*(x_{t-1})-y_{t-1})^\top\nabla_y F(x_{t-1},y_{t-1})\right]+2\eta_y\mathbb{E}\left[\|y_{t-1}-y^*(x_{t-1})\|\cdot\|B_{t-1}\|\right]\\
&\quad+\eta_y^2\left(\frac{3}{2m}\mu_2^2d_2^2L_f^2+\frac{18d_2}{m}\|\nabla_y F(x_{t-1},y_{t-1})\|^2+\frac{6\sigma^2d_2^2}{\mu_2^2m}\right) \qquad (27)\\
&\leq \delta_{t-1}-2\eta_y\mathbb{E}\left[\frac{\tau}{2}\|y^*(x_{t-1})-y_{t-1}\|^2+\frac{1}{2L_f}\|\nabla_y F(x_{t-1},y_{t-1})\|^2\right]\\
&\quad+\frac{\eta_y\tau}{2}\delta_{t-1}+\frac{2\eta_y\mu_2^2L_f^2}{\tau}+\frac{36d_2}{m}\eta_y^2\mathbb{E}\left[\|\nabla_y F(x_{t-1},y_{t-1})\|^2\right]+\frac{3}{2m}\mu_2^2d_2^2L_f^2\eta_y^2+\frac{6\sigma^2d_2^2\eta_y^2}{\mu_2^2m}\\
&\leq \left(1-\frac{\eta_y\tau}{2}\right)\delta_{t-1}+\frac{2\eta_y\mu_2^2L_f^2}{\tau}+\frac{3\mu_2^2d_2^2L_f^2\eta_y^2}{2m}+\frac{6\sigma^2d_2^2\eta_y^2}{m\mu_2^2},
\end{aligned}
$$

where $H'_{t-1} := H_{\mu_2}(x_{t-1}, y_{t-1}, u_{M_2}^{t-1}; \tilde{\xi}_M^{t-1}, \xi_{2,M}^{t-1})$, $B_{t-1} := \nabla_y F_{\mu_2}(x_{t-1}, y_{t-1}) - \nabla_y F(x_{t-1}, y_{t-1})$ represents the smoothing bias bounded as $\|B_{t-1}\| \leq \mu_2 L_f$ by [1, (2.35)], and the third inequality follows from the co-coercivity of strongly concave functions.

Then,

$$
\begin{aligned}
\delta_t &\leq \left(1+\frac{1}{\lambda}\right)\left[\left(1-\frac{\eta_y\tau}{2}\right)\delta_{t-1}+\frac{2\eta_y\mu_2^2L_f^2}{\tau}+\frac{3\mu_2^2d_2^2L_f^2\eta_y^2}{2m}+\frac{6\sigma^2d_2^2\eta_y^2}{m\mu_2^2}\right]\\
&\quad+(1+\lambda)\kappa^2\eta_x^2\mathbb{E}\left[\frac{3}{2m}\mu_1^2d_1^2L_f^2+\frac{36d_1}{m}\|\nabla\Phi(x_{t-1})\|^2+\frac{36d_1}{m}L_f^2\|y^*(x_{t-1})-y_{t-1}\|^2+\frac{6\sigma^2d_1^2}{\mu_1^2b}\right]\\
&\leq \left(1-\frac{\eta_y\tau}{4}\right)\delta_{t-1}+\frac{4\eta_y\mu_2^2L_f^2}{\tau}+\frac{3\mu_2^2d_2^2L_f^2\eta_y^2}{m}+\frac{12\sigma^2d_2^2\eta_y^2}{\mu_2^2m}+\frac{6\kappa^2\eta_x^2\mu_1^2d_1^2L_f^2}{\eta_y\tau m}\\
&\quad+\frac{144d_1\kappa^2\eta_x^2}{\eta_y\tau m}\mathbb{E}\left[\|\nabla\Phi(x_{t-1})\|^2\right]+\frac{144d_1\kappa^2\eta_x^2L_f^2}{\eta_y\tau m}\delta_{t-1}+\frac{24\kappa^2\eta_x^2\sigma^2d_1^2}{\eta_y\tau\mu_1^2m}\\
&= \gamma_2\delta_{t-1}+\frac{144d_1}{\eta_y\tau m}\kappa^2\eta_x^2\mathbb{E}\left[\|\nabla\Phi(x_{t-1})\|^2\right]+C_2
\end{aligned}
$$

$$\leq \gamma_2^t D^2 + \frac{144 d_1}{\eta_y \tau m} \kappa^2 \eta_x^2 \sum_{j=0}^{t-1} \gamma_2^{t-1-j} \mathbb{E}\left[\|\nabla \Phi(x_j)\|^2\right] + C_2 \sum_{j=0}^{t-1} \gamma_2^{t-1-j}.$$

where the second inequality follows by choosing $\lambda = \frac{4-2\eta_y\tau}{\eta_y\tau}$ and the condition $\eta_y \leq \frac{4}{3\tau}$. The proof is complete. □

Theorem 2 (Convergence Results)**.** Let $\{(x_t, y_t)\}_{t=0}^T$ be the sequence generated by Algorithm 2, and $d := d_1 + d_2$. Suppose that Assumptions 1-4 hold. Then for any $\epsilon > 0$, the following statements hold.

(i) **Accuracy-Independent Constant Step-size:** Suppose that the parameters are chosen such that

(a) the smoothing parameters $\mu_1 = \Theta(\epsilon)$ and $\mu_2 = \Theta(\epsilon\kappa^{-1})$;

(b) the step-sizes $\eta_x = \Theta(\kappa^{-2})$ and $\eta_y = \Theta(\kappa^{-1})$ such that $\eta_x \leq \min\left(\frac{\eta_y\tau\sqrt{m}}{48\sqrt{6d_1}\kappa L_f}, \frac{m}{216 d_1 L_g}\right)$ and $\eta_y \leq \min\left(\frac{m}{36 L_f d_2}, \frac{4}{3\tau}\right)$;

(c) the batch size $m = \Theta(\epsilon^{-4}\kappa^2 d^2)$.

Then, MD-ZO-SGDA achieves an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-2})$ and a sample complexity of $\mathcal{O}(d^2\kappa^4\epsilon^{-6})$.

(ii) **Accuracy-Dependent Step-size:** Suppose that the parameters are chosen such that

(a) the smoothing parameters $\mu_1 = \Theta(\epsilon)$ and $\mu_2 = \Theta(\epsilon\kappa^{-1})$;

(b) the step-sizes $\eta_x = \Theta(\epsilon^3\kappa^{-2})$ and $\eta_y = \Theta(\epsilon^2\kappa^{-1})$ such that $\eta_x \leq \min\left(\frac{\eta_y\tau\sqrt{m}}{48\sqrt{6d_1}\kappa L_f}, \frac{m}{216 d_1 L_g}\right)$ and $\eta_y \leq \min\left(\frac{m}{36 L_f d_2}, \frac{4}{3\tau}\right)$;

(c) the batch size $b = \Theta(\epsilon^{-2}\kappa^2 d^2)$.

Then, MD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-5})$ and a sample complexity of $\mathcal{O}(d^2\kappa^4\epsilon^{-7})$.

Proof. By the same arguments as in Theorem 1, we obtain

$$\begin{aligned}
\sum_{t=0}^T \frac{\eta_x}{2}\|\nabla\Phi(x_t)\|^2 \leq \Phi(x_0) - \Phi(x^*) &+ \eta_x \sum_{t=0}^T \|\nabla_x F_{\mu_1}(x_t, y_t) - \nabla_x F(x_t, y_t)\|^2 \\
&+ \eta_x \sum_{t=0}^T \|\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)\|^2 - \eta_x \sum_{t=0}^T \nabla\Phi(x_t)^\top \Delta_t \\
&+ \frac{L}{2}\eta_x^2 \sum_{t=0}^T \|G_{\mu_1}(x_t, y_t, u_{M_1}^t, \tilde{\xi}_M^t, \xi_{1,M}^t)\|^2 \\
&\leq \Phi(x_0) - \Phi(x^*) + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^T \|y^*(x_t) - y_t\|^2
\end{aligned}$$

$$- \eta_x \sum_{t=0}^{T} \nabla\Phi(x_t)^\top \Delta_t + \frac{L}{2}\eta_x^2 \sum_{t=0}^{T} \|G_{\mu_1}(x_t, y_t, u_{M_1}^t, \tilde{\xi}_M^t, \xi_{1,M}^t)\|^2,$$

where $\Delta_t := G_{\mu_1}(x_t, y_t, u_{M_1}, \tilde{\xi}^M, \xi_1^M) - \nabla_x F_{\mu_1}(x_t, y_t)$.

Taking the expectation with respect to u_{M_1}, $\tilde{\xi}_M$, $\xi_{1,M}$ on the both sides of the above inequality, we have

$$\begin{aligned}
&\frac{\eta_x}{2} \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla\Phi(x_t)\|^2] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|y^*(x_t) - y_t\|^2] \\
&\quad - \eta_x \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\nabla\Phi(x_t)^\top \Delta_t] + \frac{9Ld_1}{m}\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla_x F(x_t, y_t)\|^2] \\
&\quad + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 m}(T+1) \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + \eta_x L_f^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|y^*(x_t) - y_t\|^2] \\
&\quad + \frac{18Ld_1}{m}\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla_x F(x_t, y_t) - \nabla\Phi(x_t)\|^2 + \|\nabla\Phi(x_t)\|^2] \\
&\quad + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 m}(T+1) \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + (\eta_x L_f^2 + \frac{18Ld_1}{m}\eta_x^2 L_f^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|y^*(x_t) - y_t\|^2] \\
&\quad + \frac{18Ld_1}{m}\eta_x^2 \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla\Phi(x_t)\|^2] + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 m}(T+1).
\end{aligned}$$

Rearranging the terms in the above inequality,

$$\begin{aligned}
&(\frac{\eta_x}{2} - \frac{18Ld_1}{m}\eta_x^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla\Phi(x_t)\|^2] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + (\eta_x L_f^2 + \frac{18Ld_1}{m}\eta_x^2 L_f^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|y^*(x_t) - y_t\|^2] \\
&\quad + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 m}(T+1) \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) \\
&\quad + (\eta_x L_f^2 + \frac{18Ld_1}{m}\eta_x^2 L_f^2) \left(D^2 \sum_{t=0}^{T} \gamma_2^t + \frac{144 d_1}{\eta_y \tau m}\kappa^2 \eta_x^2 \sum_{t=0}^{T} \sum_{j=0}^{t-1} \gamma_2^{t-1-j} \mathbb{E}[\|\nabla\Phi(x_j)\|^2] + C_2 \sum_{t=0}^{T} \sum_{j=0}^{t-1} \gamma_2^{t-1-j} \right) \\
&\quad + \frac{3L\eta_x^2\mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2\sigma^2 d_1^2}{\mu_1^2 m}(T+1)
\end{aligned}$$

$$
\begin{aligned}
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) \\
&\quad + (\eta_x L_f^2 + \frac{18Ld_1}{m}\eta_x^2 L_f^2)\left(\frac{8D^2}{\eta_y \tau} + \frac{1152 d_1}{\eta_y^2 \tau^2 m}\kappa^2 \eta_x^2 \sum_{t=0}^{T} \mathbb{E}[\|\nabla \Phi(x_j)\|^2] + \frac{8C_2}{\eta_y \tau}(T+1)\right) \\
&\quad + \frac{3L\eta_x^2 \mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2 \sigma^2 d_1^2}{\mu_1^2 m}(T+1),
\end{aligned}
$$

where the third inequality follows from $\sum_{t=0}^{\infty} \gamma_2^t \leq \frac{8}{\eta_y \tau}$.

By further rearranging the terms,

$$
\begin{aligned}
&(\frac{\eta_x}{2} - \frac{18Ld_1}{m}\eta_x^2 - \frac{1152 d_1}{\eta_y^2 \tau^2 m}\kappa^2 \eta_x^3 L_f^2 - \frac{20736}{\eta_y^2 \tau^2 m^2}\kappa^2 \eta_x^4 L d_1^2 L_f^2) \sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla \Phi(x_t)\|^2] \\
&\leq \mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \eta_x \mu_1^2 L_f^2 (T+1) + (\eta_x L_f^2 + \frac{18Ld_1}{m}\eta_x^2 L_f^2)\left(\frac{8D^2}{\eta_y \tau} + \frac{8C_2}{\eta_y \tau}(T+1)\right) \\
&\quad + \frac{3L\eta_x^2 \mu_1^2 L_f^2 d_1^2}{4m}(T+1) + \frac{3L\eta_x^2 \sigma^2 d_1^2}{\mu_1^2 m}(T+1).
\end{aligned}
$$

Dividing both sides of the above inequality by

$$
(\frac{\eta_x}{2} - \frac{18Ld_1}{m}\eta_x^2 - \frac{1152 d_1}{\eta_y^2 \tau^2 m}\kappa^2 \eta_x^3 L_f^2 - \frac{20736}{\eta_y^2 \tau^2 m^2}\kappa^2 \eta_x^4 L d_1^2 L_f^2)(T+1),
$$

and applying the parameter constraint $\eta_x \leq \min\left(\frac{\eta_y \tau \sqrt{m}}{48\sqrt{6d_1}\kappa L_f}, \frac{m}{216 d_1 L}\right)$, we have

$$
\begin{aligned}
&\frac{1}{T+1}\sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}[\|\nabla \Phi(x_t)\|^2] \\
&\leq \frac{4}{(T+1)\eta_x}\mathbb{E}_{\zeta[0,T]}[\Phi(x_0) - \Phi(x^*)] + \frac{(L_f^2 + \frac{18Ld_1}{m}\eta_x L_f^2)32D^2}{(T+1)\eta_y \tau} + 4\mu_1^2 L_f^2 \\
&\quad + (L_f^2 + \frac{18Ld_1}{m}\eta_x L_f^2)\frac{32C_2}{\eta_y \tau} + \frac{3L\eta_x \mu_1^2 L_f^2 d_1^2}{m} + \frac{12L\eta_x \sigma^2 d_1^2}{\mu_1^2 m},
\end{aligned}
$$

which implies

$$
\begin{aligned}
&\frac{1}{T+1}\sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}\left[\|\nabla \Phi(x_t)\|^2\right] \\
&= \mathcal{O}\left((T+1)^{-1}\eta_x^{-1}\right) + \mathcal{O}\left(\kappa(T+1)^{-1}\cdot \eta_y^{-1}\right) + \mathcal{O}\left(d_1 \kappa^2 (T+1)^{-1}\eta_y^{-1} m^{-1}\eta_x\right) + \mathcal{O}\left(\mu_1^2\right) \\
&\quad + \mathcal{O}\left(\mu_2^2 \kappa^2\right) + \mathcal{O}\left(\mu_2^2 d_2^2 \eta_y \kappa m^{-1}\right) + \mathcal{O}\left(d_2^2 \eta_y \kappa \mu_2^{-2} m^{-1}\right) + \mathcal{O}\left(\kappa^4 \eta_x^2 \mu_1^2 d_1^2 \eta_y^{-2} m^{-1}\right) + \mathcal{O}\left(\kappa^4 \eta_x^2 d_1^2 \eta_y^{-2} \mu_1^{-2} m^{-1}\right) \\
&\quad + \mathcal{O}\left(d_1 \eta_x \mu_2^2 d_2^2 \kappa^3 m^{-1}\right) + \mathcal{O}\left(d_1 \eta_x \mu_2^2 d_2^2 \eta_y \kappa^2 m^{-2}\right) + \mathcal{O}\left(d_1 \eta_x d_2^2 \eta_y \kappa^2 \mu_2^{-2} m^{-2}\right) + \mathcal{O}\left(\kappa^5 \eta_x^3 \mu_1^2 d_1^3 \eta_y^{-2} m^{-2}\right) \\
&\quad + \mathcal{O}\left(\kappa^5 \eta_x^3 d_1^3 \eta_y^{-2} \mu_1^{-2} m^{-2}\right) + \mathcal{O}\left(\eta_x \kappa \mu_1^2 d_1^2 m^{-1}\right) + \mathcal{O}\left(\eta_x \kappa d_1^2 \mu_1^{-2} m^{-1}\right).
\end{aligned}
$$

For the accuracy-independent constant step-sizes setting, $\mu_1 = \Theta(\epsilon)$, $\mu_2 = \Theta(\kappa^{-1}\epsilon)$, $\eta_x = \Theta(\kappa^{-2})$, $\eta_y = \Theta(\kappa^{-1})$, $m = \Theta(d^2 \kappa^2 \epsilon^{-4})$, we have

$$
\frac{1}{T+1}\sum_{t=0}^{T} \mathbb{E}_{\zeta[0,T]}\left[\|\nabla \Phi(x_t)\|^2\right] \leq \mathcal{O}(\epsilon^2), \tag{28}
$$

if $T = \Theta(\kappa^2\epsilon^{-2})$. This means that, MD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-2})$, and a sample complexity of $\mathcal{O}(d^2\kappa^4\epsilon^{-6})$.

On the other hand, for the accuracy-dependent step-sizes setting, $\mu_1 = \Theta(\epsilon)$, $\mu_2 = \Theta(\kappa^{-1}\epsilon)$, $\eta_x = \Theta(\kappa^{-2}\epsilon^3)$, $\eta_y = \Theta(\kappa^{-1}\epsilon^2)$, $m = \Theta(d^2\kappa^2\epsilon^{-2})$, we have

$$\frac{1}{T+1}\sum_{t=0}^{T}\mathbb{E}_{\zeta[0,T]}\left[\|\nabla\Phi(x_t)\|^2\right] \leq \mathcal{O}(\epsilon^2), \tag{29}$$

if $T = \Theta(\kappa^2\epsilon^{-5})$. This means that, MD-ZO-SGDA attains an iteration complexity of $\mathcal{O}(\kappa^2\epsilon^{-5})$, and a sample complexity of $\mathcal{O}(d^2\kappa^4\epsilon^{-7})$. □

Remark 2. A comparison between Theorem 1 and Theorem 2 reveals that MD-ZO-SGDA reduces the iteration complexity from $\mathcal{O}(d^3\kappa^3\epsilon^{-2})$ to $\mathcal{O}(\kappa^2\epsilon^{-2})$, and the sample complexity from $\mathcal{O}(d^3\kappa^5\epsilon^{-6})$ to $\mathcal{O}(d^2\kappa^4\epsilon^{-6})$. The error of the zeroth-order gradient estimator can be decomposed into two primary sources: the smoothing bias and the estimation variance. The smoothing bias is controlled by the parameters μ_1 and μ_2, our contribution focuses on reducing the estimation variance, which results from two different random sources: the directional variance (due to random perturbations) and the sampling variance (due to finite samples of ξ). Increasing the number of directions m yields a tighter upper bound on the estimation error because it addresses both issues simultaneously. Intuitively, each additional direction u^i involves a fresh resampling of ξ, which concurrently reduces the directional variance and the sampling variance. In contrast, increasing the batch size b for a single direction only reduces the sampling variance, and provides no relief for the error resulting from poor directional exploration. Specifically, MD-ZO-SGDA uses $m = \Theta(d^2\kappa^2\epsilon^{-4})$ to reduces the $\mathcal{O}(\frac{d^2}{m})$ directional variance. This allows for a dimension-independent step-size, leading to the $\mathcal{O}(\kappa^2\epsilon^{-2})$ iteration complexity, and an improved sample complexity of $\mathcal{O}(d^2\kappa^4\epsilon^{-6})$.

5 Numerical Experiments

In this section, we evaluate the performance of SD-ZO-SGDA and MD-ZO-SGDA across both a simple synthetic example and a real-world application.

5.1 A synthetic nonconvex-strongly concave minimax problem

Inspired by recent studies on alternating gradient descent ascent for nonconvex–strongly concave quadratic problems [2], we consider a multi-dimensional synthetic example to fully evaluate the performance of our sample-based zeroth-order algorithms. We incorporate a Lorentzian-type potential term to introduce a more complex, non-convex landscape. Specifically, the stochastic minimax problem is formulated as:

$$\min_{x\in\mathbb{R}^d}\max_{y\in\mathbb{R}^d}\mathbb{E}_{\xi\sim D(x,y)}[f(x,y,\xi)],$$

where the stochastic objective f and the decision-dependent distribution $D(x,y)$ are given by:

$$f(x,y,\xi) = -\frac{10}{1+0.5\|x\|^2} - y^\top Hy + \xi^\top(Hy), \quad \xi\sim\mathcal{N}(x+0.5y,\sigma^2 I_d).$$

The matrix H is diagonal, and the ratio of its maximum eigenvalue to the minimum eigenvalue defines the condition number κ of the strongly concave term. Under this setting, the resulting expected objective function is $F(x,y) = -\frac{10}{1+0.5\|x\|^2} - 0.5y^\top Hy + x^\top Hy$. It can be verified that $F(x,y)$ is non-convex in x and strongly concave in y. For any fixed x, the unique optimal response is $y^*(x) = x$. Due to the analytically tractable nature of this problem, the unique global saddle point is $(x^*, y^*) = (0,0)$. Therefore, we use the squared distance to the saddle point $\|z_t - z^*\|^2 = \|x_t - x^*\|^2 + \|y_t - y^*\|$ as the performance metric.

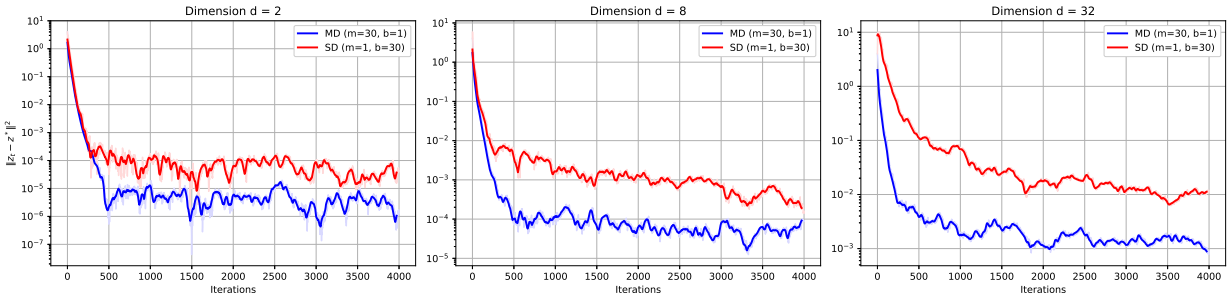

Figure 1: Convergence comparison between MD-ZO-SGDA and SD-ZO-SGDA across different dimensions ($d \in |\{2, 8, 32\}$). Under a fixed sampling budget per iteration, MD-ZO-SGDA consistently achieves a lower error floor than SD-ZO-SGDA.

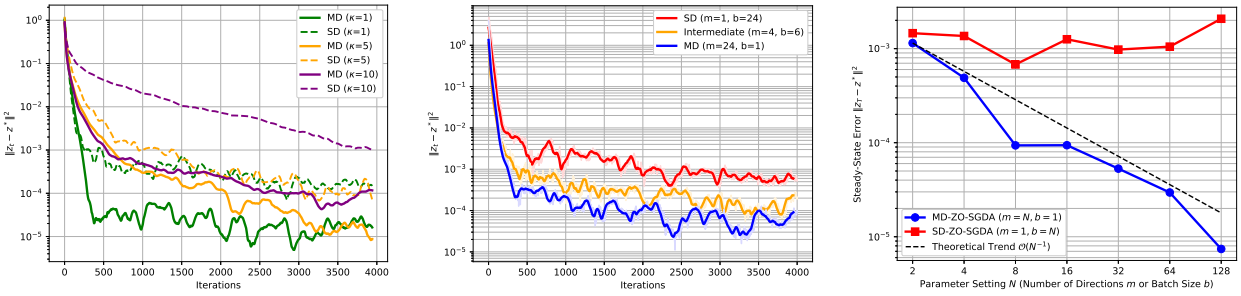

(a) Impact of Condition Number κ on Convergence

(b) Ablation: Trade-off between m and b

(c) Accuracy Dependence Verification

Figure 2: Comparisons of MD-ZO-SGDA and SD-ZO-SGDA. (a) Convergence curves under different condition numbers $\kappa \in \{1, 5, 10\}$. (b) Ablation study on the combinations of directions m and batch size b under a fixed budget. (c) The steady-state error versus the budget parameter N. The dashed line represents the theoretical $\mathcal{O}(N^{-1})$ convergence trend.

As illustrated in Figure 1 and Figure 2a, the steady-state error floors for both algorithms rise as d and κ increase. This upward trend aligns with the polynomial dimension dependency relationship established in our complexity theorem. However, MD-ZO-SGDA consistently maintains a significant accuracy advantage over SD-ZO-SGDA. Figure 2b demonstrates that under a fixed sampling budget ($N = 24$), utilizing more directions leads to better convergence. To further verify the theoretical relationship between the target accuracy and the budget N (where we set $N = m$ for MD-ZO-SGDA and $N = b$ for SD-ZO-SGDA), we plot their steady-state errors in Figure 2c. In this log-log plot, the steady-state error of MD-ZO-SGDA decreases linearly which perfectly matches our theoretical $\mathcal{O}(N^{-1})$ rate. However, the error of SD-ZO-SGDA stops decreasing and hits an error

floor, even if we use a very large batch size b. This result is highly consistent with our theoretical analysis in Remark 2. For SD-ZO-SGDA, the directional variance term does not have a $\frac{1}{b}$ factor. Therefore, increasing b only reduces the sampling noise but cannot reduce the directional error. MD-ZO-SGDA solves this problem by using m, directions, which introduces a $\frac{1}{m}$ factor to reduce both types of variance at the same time.

5.2 Problem Formulation: Competitive EV Charging

In this subsection, we evaluate the performance of the proposed algorithms on a competitive electric vehicle (EV) charging problem, a benchmark task for decision-dependent minimax optimization formulated by [33]. We compare our methods against their established baselines: EPD, SEPD, and a one-point zeroth-order algorithm (DFO). We consider a market with two competing charging service providers operating across $n = 3$ distinct zones. Each provider aims to maximize their respective profit by adjusting the price differentials $x, y \in \mathbb{R}^n$. Following the model in [19], the profit function for provider one is defined as:

$$u_1(x, a) = \langle a + r, x + p \rangle - \|\Gamma_1 x\|^2,$$

where a represents the service demand, r is the utility vector, p is the baseline price, and Γ_1 is a diagonal matrix representing the cost of maintaining service quality. The demand a is decision-dependent and follows a linear best-response model:

$$a = a_0 + A_1 x + A_2 y, \quad b = b_0 + B_1 x + B_2 y,$$

where $a_0 \sim D_1$ and $b_0 \sim D_2$ are stochastic components. The objective is to find the saddle point of the relative profit, leading to the following stochastic minimax problem:

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \mathbb{E}_{(a,b) \sim \mathcal{D}(x,y)} \left[\|\Gamma_1 x\|^2 - \|\Gamma_2 y\|^2 - \langle a + r, x \rangle + \langle b + r, y \rangle \right],$$

where $\mathcal{X} = \mathcal{Y} = [-1, 2]^3$.

Our simulation utilizes real-world demand data processed from [11]. We focus on a specific time window (e.g., 12:00 PM to 1:00 PM) to observe peak-hour dynamics. The data are normalized by subtracting the mean and scaling by the variance.

For the algorithm parameters, we set the service quality coefficients $\gamma_{j,i} = 1$ and the utility values $r_i = 0$ for all stations. The elasticity matrices are configured as $(A_1)_{i,j} = -0.3\delta_{i,j}$ and $(A_2)_{i,j} = 0.3\delta_{i,j}$, reflecting the sensitivity of demand to price changes.

We execute the algorithms for $T = 4,000$ iterations. Figure 3 illustrates the daily demand fluctuations for the three charging stations over a year, highlighting the inherent stochasticity of the environment. Figure 4 displays the convergence behavior in terms of the squared distance to the optimal saddle point $\|z_t - z^*\|^2$. The results demonstrate that our proposed algorithm effectively reduces the estimation error despite the decision-dependent nature of the distributions. Notably, during the initial phase of optimization, both SD-ZO-SGDA and MD-ZO-SGDA exhibit a significantly more rapid descent in the squared gradient norm compared to DFO. Specifically, our proposed methods approach a stable error regime within fewer iterations, escaping the initial

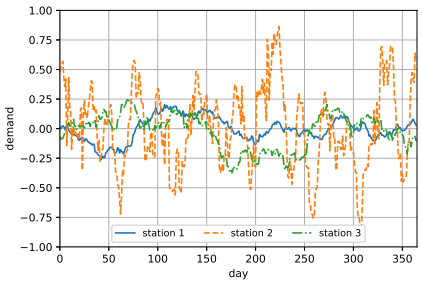

Figure 3: Demand change between 12-1 pm

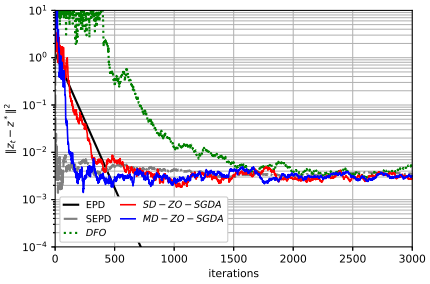

Figure 4: Algorithm Error per iteration

high-error state more efficiently than the DFO. Once reaching the steady state, the performance gap between the algorithms diminishes. Both SD-ZO-SGDA and MD-ZO-SGDA settle at an error floor that is comparable to, and slightly lower than, that of the DFO baseline. These observations suggest that while all three algorithms eventually reach a similar level of precision, the proposed zeroth-order SGDA frameworks offer a clear advantage in terms of initial convergence rate.

References

[1] Albert S Berahas, Liyuan Cao, Krzysztof Choromanski, and Katya Scheinberg. A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Foundations of Computational Mathematics*, 22(2):507–560, 2022.

[2] Radu Ioan Boţ and Axel Böhm. Alternating proximal-gradient steps for (stochastic) nonconvex-concave minimax problems. *SIAM Journal on Optimization*, 33(3):1884–1913, 2023.

[3] Yair Carmon, Yujia Jin, Aaron Sidford, and Kevin Tian. Variance reduction for matrix games. *Advances in Neural Information Processing Systems*, 32, 2019.

[4] Yatong Chen, Wei Tang, Chien-Ju Ho, and Yang Liu. Performative prediction with bandit feedback: Learning through reparameterization. *arXiv preprint arXiv:2305.01094*, 2023.

[5] Wang Chi Cheung, David Simchi-Levi, and He Wang. Dynamic pricing and demand learning with limited price experimentation. *Operations Research*, 65(6):1722–1731, 2017.

[6] William L Cooper, Tito Homem-de Mello, and Anton J Kleywegt. Models of the spiral-down effect in revenue management. *Operations research*, 54(5):968–987, 2006.

[7] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018.

[8] Xiang Gao, Bo Jiang, and Shuzhong Zhang. On the information-adaptive variants of the admm: an iteration complexity perspective. *Journal of Scientific Computing*, 76(1):327–363, 2018.

[9] Yan Gao and Yongchao Liu. Adaptive stochastic gradient descent ascent algorithm for nonconvex minimax problems with decision-dependent distributions. *arXiv preprint arXiv:2509.11018*, 2025.

[10] Yan Gao, Yongchao Liu, and Zili Luo. Trust region algorithm for stochastic minimax problems with decision-dependent distributions. *arXiv preprint arXiv:2509.12868*, 2025.

[11] Madeline Gilleran, Eric Bonnema, Jason Woods, Partha Mishra, Ian Doebber, Chad Hunter, Matt Mitchell, and Margaret Mann. Impact of electric vehicle charging on the power demand of retail buildings. *Advances in Applied Energy*, 4:100062, 2021.

[12] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[13] Yuya Hikima and Akiko Takeda. Zeroth-order methods for nonconvex stochastic problems with decision-dependent distributions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 17195–17203, 2025.

[14] Yuya Hikima and Akiko Takeda. Zeroth-order gradient estimators for stochastic problems with decision-dependent distributions. *arXiv preprint arXiv:2510.24929*, 2025.

[15] Feihu Huang, Shangqian Gao, Jian Pei, and Heng Huang. Accelerated zeroth-order and first-order momentum methods from mini to minimax optimization. *Journal of Machine Learning Research*, 23(36):1–70, 2022.

[16] Juan Inga and Erwin Sacoto-Cabrera. Credit default risk analysis using machine learning algorithms with hyperparameter optimization. In *International Conference on Science, Technology and Innovation for Society*, pages 81–95. Springer, 2022.

[17] Kaiyi Ji, Zhe Wang, Yi Zhou, and Yingbin Liang. Improved zeroth-order variance reduced algorithms and analysis for nonconvex optimization. In *International conference on machine learning*, pages 3100–3109. PMLR, 2019.

[18] Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally robust optimization. *Advances in neural information processing systems*, 33: 8847–8860, 2020.

[19] Chaojie Li, Zhaoyang Dong, Guo Chen, Bo Zhou, Jingqi Zhang, and Xinghuo Yu. Data-driven planning of electric vehicle charging infrastructure: A case study of sydney, australia. *IEEE Transactions on Smart Grid*, 12(4):3289–3304, 2021.

[20] Thomas Liebig, Nico Piatkowski, Christian Bockermann, and Katharina Morik. Dynamic route planning with real-time traffic predictions. *Information Systems*, 64:258–265, 2017.

[21] Tianyi Lin, Chi Jin, and Michael I Jordan. Two-timescale gradient descent ascent algorithms for nonconvex minimax optimization. *Journal of Machine Learning Research*, 26(11):1–45, 2025.

[22] Haitong Liu, Qiang Li, and Hoi-To Wai. Two-timescale derivative free optimization for performative prediction with markovian data. *arXiv preprint arXiv:2310.05792*, 2023.

[23] Sijia Liu, Songtao Lu, Xiangyi Chen, Yao Feng, Kaidi Xu, Abdullah Al-Dujaili, Mingyi Hong, and Una-May O'Reilly. Min-max optimization without gradients: Convergence and applications to black-box evasion and poisoning attacks. In *International conference on machine learning*, pages 6282–6293. PMLR, 2020.

[24] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. *Advances in neural information processing systems*, 29, 2016.

[25] Adhyyan Narang, Evan Faulkner, Dmitriy Drusvyatskiy, Maryam Fazel, and Lillian J Ratliff. Multiplayer performative prediction: Learning in decision-dependent games. *Journal of Machine Learning Research*, 24(202):1–56, 2023.

[26] John Nash. Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, pages 128–140, 1953.

[27] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609. PMLR, 2020.

[28] Mitas Ray, Lillian J Ratliff, Dmitriy Drusvyatskiy, and Maryam Fazel. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8081–8088, 2022.

[29] Natasha Robinson and Nidhi Sindhwani. Loan default prediction using machine learning. In *2024 11th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions)(ICRITO)*, pages 1–5. IEEE, 2024.

[30] Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.

[31] Hoi-To Wai, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong. Multi-agent reinforcement learning via double averaging primal-dual optimization. *Advances in neural information processing systems*, 31, 2018.

[32] Zhongruo Wang, Krishnakumar Balasubramanian, Shiqian Ma, and Meisam Razaviyayn. Zeroth-order algorithms for nonconvex minimax problems with improved complexities. *arXiv preprint arXiv:2001.07819*, 2020.

[33] Killian Wood and Emiliano Dall'Anese. Stochastic saddle point problems with decision-dependent distributions. *SIAM Journal on Optimization*, 33(3):1943–1967, 2023.

[34] Killian Wood, Ahmed S Zamzam, and Emiliano Dall'Anese. Solving decision-dependent games by learning from feedback. *IEEE Open Journal of Control Systems*, 3:295–309, 2024.

[35] Zi Xu, Zi-Qi Wang, Jun-Lin Wang, and Yu-Hong Dai. Zeroth-order alternating gradient descent ascent algorithms for a class of nonconvex-nonconcave minimax problems. *Journal of Machine Learning Research*, 24(313):1–25, 2023.

[36] Landi Zhu, Mert Gürbüzbalaban, and Andrzej Ruszczyński. Distributionally robust learning with weakly convex losses: Convergence rates and finite-sample guarantees. *arXiv preprint arXiv:2301.06619*, 2023.