Zeroth-Order Algorithms for Stochastic Nonconvex Minimax Problems with Improved Complexities

Zhongruo Wang [1] **Krishnakumar Balasubramanian** [2] **Shiqian Ma** [1] **Meisam Razaviyayn** [3]

Abstract

In this paper, we study zeroth-order algorithms for stochastic minimax optimization problems that are nonconvex in one variable and strongly-concave in the other variable. Such minimax optimization problems have attracted significant attention lately due to their applications in modern machine learning tasks. We design and analyze the Zeroth-Order Stochastic Gradient Descent Ascent (`ZO-SGDA`) algorithm, and provide improved results compared to existing works, in terms of oracle complexity. Next, we propose the Zeroth-Order Stochastic Gradient Descent Multi-Step Ascent (`ZO-SGDMSA`) algorithm that significantly improves the oracle complexity of `ZO-SGDA`. Numerical results are presented.

1. Introduction

Algorithms for solving optimization problems with only access to noisy evaluations of the objective function are called zeroth-order algorithms. Such zeroth-order optimization algorithms have been studied for decades in the optimization literature; see, for example, (Conn et al., 2009; Rios & Sahinidis, 2013; Audet & Hare, 2017) for a detailed overview of the existing approaches. Recently, the study of zeroth-order optimization algorithms has gained significant attention also in the machine learning literature, due to several motivating applications, for example, in designing black-box attacks to deep neural networks (Chen et al., 2017), hyperparameter tuning (Snoek et al., 2012), reinforcement learning (Moriarty et al., 1999; Salimans et al., 2017) and bandit convex optimization (Bubeck et al., 2017).

[1]Department of Mathematics, University of California, Davis, CA, USA [2]Department of Statistics, University of California, Davis, CA, USA [3]Departments of Industrial and Systems Engineering, University of Southern California, CA, USA. Correspondence to: Krishnakumar Balasubramanian <kbala@ucdavis.edu>, Shiqian Ma <sqma@ucdavis.edu>.

Presented at the workshop on "Beyond First-Order Methods in Machine Learning Systems" hosted by the 38th International Conference on Machine Learning (ICML).

In this work, we study zeroth-order algorithms for solving the following stochastic nonconvex minimax problems:

$$\min_{x\in\mathbb{R}^{d_1}} \max_{y\in\mathcal{Y}} f(x,y) = \mathsf{E}_{\xi\sim\mathcal{P}} F(x,y,\xi). \tag{1}$$

Here, $F(x,y,\xi)$ and hence $f(x,y)$ are assumed to be sufficiently smooth functions, $\mathcal{Y}\subset\mathbb{R}^{d_2}$ is a closed and convex constraint set, and $\mathcal{P}$ is a distribution characterizing the stochasticity in the problem. We allow for the function $f(\cdot,y)$ to be nonconvex for all $y\in\mathbb{R}^{d_2}$ but require $f(x,\cdot)$ to be strongly-concave for all $x\in\mathbb{R}^{d_1}$.

Our main motivation for studying zeroth-order algorithms for nonconvex minimax problems is its application in designing black-box attacks to deep neural networks. By now, it is well established that care must be taken when designing and training deep neural networks as it is possible to design adversarial examples that would make the deep network to misclassify, easily. Since the intriguing works of (Szegedy et al., 2013; Liu et al., 2017), the problem of designing such adversarial examples that transfer across multiple deep neural networks models, has been studied extensively. As the model architecture is unknown to the adversary, the problem could naturally be formulated as a minimax optimization problem under the availability of only noisy objective function evaluation. We refer the reader to (Liu et al., 2020) for details regarding such formulations. Moreover, we note that zeroth-order minimax optimization problems also arise in multi-agent reinforcement learning with bandit feedback (Wei et al., 2017; Zhang et al., 2019), robotics (Wang & Jegelka, 2017; Bogunovic et al., 2018) and distributionally robust optimization (Namkoong & Duchi, 2016).

Recently, there has been an ever-growing interest in analyzing first-order algorithms for the case of nonconvex-concave objective and nonconvex-nonconcave objectives, motivated by its applications in training generative adversarial networks (Goodfellow et al., 2014), AUC maximization (Ying et al., 2016), designing fair classifiers (Agarwal et al., 2018), robust learning systems (Madry et al., 2017) fair machine learning (Zhang et al., 2018; Xu et al., 2018; Baharlouei et al., 2019), and reinforcement learning (Pfau & Vinyals, 2016; Dai et al., 2018; Neyman et al., 2003; Filar & Vrieze, 2012). Specifically, (Lu et al., 2019; Rafique et al., 2018; Nouiehed et al., 2019; Sanjabi et al., 2018; Lin et al., 2020; Thekumparampil et al., 2019), proposed and analyzed variants of gradient descent ascent for nonconvex-concave objectives. Very recently, under a stronger mean-squared Lipschitz gradient assumption, (Luo et al., 2020) obtained improved complexity for stochastic nonconvex-concave objectives. Furthermore, (Daskalakis et al., 2018; Daskalakis & Panageas, 2018; Hsieh et al., 2018; Mertikopoulos et al., 2018; Piliouras & Schulman, 2018; Gidel et al., 2018; Oliehoek et al., 2018; Jin et al., 2019; Flokas et al., 2019) studied general nonconvex-nonconcave objectives. Compared to first-order algorithms, zeroth-order algorithms for minimax optimization problems are underdeveloped. Motivated by the need for robustness in optimization, (Menickelly & Wild, 2018) proposed derivative-free algorithms for saddle-point optimization. However, they do not provide non-asymptotic oracle complexity analysis. Bayesian optimization algorithms and evolutionary algorithms were proposed in (Bogunovic et al., 2018; Picheny et al., 2019) and (Bertsimas & Nohadani, 2010; Al-Dujaili et al., 2018) respectively for minimax optimization, targeting robust optimization and learning applications. The above works do not provide any oracle complexity analysis. Recently, (Roy et al., 2019) studied zeroth-order Frank-Wolfe algorithms for strongly-convex and strongly-concave constrained saddle-point optimization problems and provided non-asymptotic oracle complexity analysis. Furthermore, (Liu et al., 2020) studied zeroth-order algorithms for nonconvex-concave minimax problems, similar to our setting. More recently, (Anagnostidis et al., 2021) proposed a stochastic direct search method for (1) under the assumption of the Polyak-Łojasiewicz (PL) condition. (Xu et al., 2021) and (Huang et al., 2020) also studied zeorth-order methods for (1), where they required mean-squared smoothnesss assumption, which is stronger than our assumptions.

Our Contributions. The contributions of this paper lie in two folds. First, we propose a zeroth-order stochastic gradient descent ascent algorithm (ZO−SGDA) for solving (1) and analyze its oracle complexity. Second, we propose a novel zeroth-order stochastic gradient descent multi-step ascent (ZO−SGDMSA) algorithm, which is motivated by (Nouiehed et al., 2019). This algorithm performs multiple steps of gradient ascent followed by one single step of gradient descent in each iteration. Its oracle complexity is significantly better than that of ZO−SGDA in terms of condition number dependency. The oracle complexity of both algorithms is better than (Liu et al., 2020). A detailed comparison of our results to existing results is provided in Table 1.

2. Preliminaries

The following assumptions are made throughout the paper.

Assumption 2.1. *The objective function $f(x,y)$ and the constraint set $\mathcal{Y}$ have the following properties:*

(i). $f(x,y)$ is continuously differentiable in x and y, and $f(\cdot,y)$ is nonconvex for all $y\in\mathcal{Y}$ and $f(x,\cdot)$ is τ-strongly concave for all $x\in\mathbb{R}^{d_1}$.

(ii). Function $g(x):=\max_{y\in\mathcal{Y}} f(x,y)$ is lower bounded. We use L_g to denote the Lipschitz constant of g.

(iii). When viewed as a function in $\mathbb{R}^{d_1+d_2}$, $f(x,y)$ is ℓ-gradient Lipschitz. We use $\kappa:=\ell/\tau$ to denote the problem condition number throughout this paper.

(iv). The constraint set $\mathcal{Y}\subset\mathbb{R}^{d_2}$ is bounded and convex, with diameter $D>0$.

We also make the following standard assumptions on the stochastic zeroth-order oracle (Nesterov & Spokoiny, 2017; Ghadimi & Lan, 2013; Balasubramanian & Ghadimi, 2019).

Assumption 2.2. *For any $x\in\mathbb{R}^{d_1}$ and $y\in\mathcal{Y}$, the stochastic zeroth-order oracle outputs an estimator $F(x,y,\xi)$ of $f(x,y)$ such that $\mathsf{E}_\xi[F(x,y,\xi)]=f(x,y)$ and $\mathsf{E}_\xi[\nabla_x F(x,y,\xi)]=\nabla_x f(x,y)$, $\mathsf{E}_\xi[\nabla_y F(x,y,\xi)]=\nabla_y f(x,y)$, $\mathsf{E}_\xi[\|\nabla_x F(x,y,\xi)-\nabla_x f(x,y)\|_2^2]\le\sigma_1^2$, and $\mathsf{E}_\xi[\|\nabla_y F(x,y,\xi)-\nabla_y f(x,y)\|_2^2]\le\sigma_2^2$.*

2.1. Zeroth-order gradient estimator

We now discuss the idea of zeroth-order gradient estimator based on Gaussian Stein's identity (Nesterov & Spokoiny, 2017). We denote $\boldsymbol{u}_1\sim N(0,\mathbf{1}_{d_1})$, $\boldsymbol{u}_2\sim N(0,\mathbf{1}_{d_2})$, where $\mathbf{1}_d$ is $d\times d$ identity matrix. We define the Gaussian smoothed functions as $f_{\mu_1}(x,y):=\mathsf{E}_{\boldsymbol{u}_1,\xi}F(x+\mu_1\boldsymbol{u}_1,y,\xi)$, $f_{\mu_2}(x,y):=\mathsf{E}_{\boldsymbol{u}_2,\xi}F(x,y+\mu_2\boldsymbol{u}_2,\xi)$, and the zeroth-order stochastic gradient estimators as: $G_{\mu_1}(x,y,\boldsymbol{u}_1,\xi)=\frac{F(x+\mu_1\boldsymbol{u}_1,y,\xi)-F(x,y,\xi)}{\mu_1}\boldsymbol{u}_1$, $H_{\mu_2}(x,y,\boldsymbol{u}_2,\xi)=\frac{F(x,y+\mu_2\boldsymbol{u}_2,\xi)-F(x,y,\xi)}{\mu_2}\boldsymbol{u}_2$, where $\mu_1>0$ and $\mu_2>0$ are smoothing parameters. One can show that the zeroth-order gradient estimators provide unbiased estimates to the gradients of the Gaussian smoothed functions, i.e., $\mathsf{E}_{\boldsymbol{u}_1,\xi}G_{\mu_1}(x,y,\boldsymbol{u}_1,\xi)=\nabla_x f_{\mu_1}(x,y)$, and $\mathsf{E}_{\boldsymbol{u}_2,\xi}H_{\mu_2}(x,y,\boldsymbol{u}_2,\xi)=\nabla_y f_{\mu_2}(x,y)$.

2.2. Complexity Measure

Recall that the minimax problem (1) is equivalent to the following `argmin`-type minimization problem: $\min_x\{g(x):=\max_y f(x,y)\}$. Following (Lin et al., 2020), we define the ϵ-stationary point of (1) as follows.

Definition 2.1. *We call $\bar{x}$ an ϵ-stationary point of (1) if $\mathsf{E}[\|\nabla g(\bar{x})\|_2^2]\le\epsilon^2$.*

3. Zeroth-order Algorithms for Stochastic Minimax Problems

Our ZO−SGDA algorithm for solving (1) is presented in Algorithm 1, which is similar to the first-order approach analyzed in (Lin et al., 2020) with a few crucial differences.

Algorithm	Order	Complexity	Objective	Constraint
GDmax ((Lin et al., 2020))	1st	$\mathcal{O}(\kappa^2\epsilon^{-2})$	NC-SC	U,C
SGDmax ((Lin et al., 2020))	1st	$\mathcal{O}(\kappa^3(\sigma_1^2+\sigma_2^2)\epsilon^{-4})$	NC-SC	U,C
Multi-step GDA((Nouiehed et al., 2019))	1st	$\tilde{\mathcal{O}}(\log(\epsilon^{-1})\epsilon^{-2})$	NC-PL	C,U
Multi-step GDA ((Nouiehed et al., 2019))	1st	$\tilde{\mathcal{O}}(\log(\epsilon^{-1})\epsilon^{-3.5})$	NC-C	C,C
ZO-min-max((Liu et al., 2020))	0th	$\tilde{\mathcal{O}}((d_1+d_2)\epsilon^{-6})$	NC-SC	C,C
`ZO-SGDA` (this work)	0th	$\mathcal{O}(\kappa^5(\sigma_1^2 d_1+\sigma_2^2 d_2)\epsilon^{-4})$	NC-SC	U,C
`ZO-SGDMSA` (this work)	0th	$\mathcal{O}(\kappa(d_1\sigma_1^2+\kappa d_2\sigma_2^2\log(\epsilon^{-1}))\epsilon^{-4})$	NC-SC	U,C

Table 1. Comparison of different algorithms. The first four algorithms are first-order, and the last three algorithms are zeroth-order. Complexity refers to calls to the gradient oracle for first-order algorithms and calls to the zeroth-order oracle for the zeroth-order algorithms. We use $\tilde{\mathcal{O}}$ to hide the κ dependency, as it was not explictly tracked and stated in the work of (Liu et al., 2020; Nouiehed et al., 2019). In the "Objective" column, "NC-SC" denotes the objective function is nonconvex for x and strongly concave for y. "C" means concave, and "PL" denotes PL condition. In the "Constraint" column, "C" denotes "constrained" and "U" denotes "unconstrained".

Specifically, we require a mini-batch gradient estimator with the batch size depending on the dimensionality of the problem and the noise variance parameter σ^2. The complexity result of Algorithm 1 is provided in Theorem 3.1.

Algorithm 1 Zeroth-Order Stochastic Gradient Descent Ascent (`ZO-SGDA`)

Initialization: (x_0, y_0), step sizes (η_1, η_2), iteration limit $S > 0$, smoothing parameters μ_1 and μ_2. Indices sets $\mathcal{M}_1$ and $\mathcal{M}_2$.

for $s = 0, \ldots, S$ **do**

Sample $\boldsymbol{u}_{1,i} \sim N(0, 1_{d_1})$ and compute

$$x_{s+1} \leftarrow x_s - \eta_1 \frac{1}{|\mathcal{M}_1|} \sum_{i\in\mathcal{M}_1} G_{\mu_1}(x_s, y_s, \boldsymbol{u}_{1,i}, \xi_i).$$

Sample $\boldsymbol{u}_{2,i} \sim N(0, 1_{d_2})$ and compute

$$y_{s+1} \leftarrow \text{Proj}_{\mathcal{Y}}\Big[y_s + \eta_2 \frac{1}{|\mathcal{M}_2|} \sum_{i\in\mathcal{M}_2} H_{\mu_2}(x_s, y_s, \boldsymbol{u}_{2,i}, \xi_i)\Big].$$

end for

Return $(x_1, y_1), \ldots, (x_S, y_S)$.

Theorem 3.1. *Let $\epsilon \in (0,1)$. Under Assumptions 2.1 and 2.2, by setting the parameters $\eta_1 := \frac{1}{4\times 12^4\kappa^2(\kappa+1)^2(\ell+1)}$, $\eta_2 := 1/(6\ell)$, $S := \mathcal{O}(\kappa^5\epsilon^{-2})$, $\mu_1 := \mathcal{O}(\epsilon d_1^{-3/2}\kappa^{-2})$, $\mu_2 := \mathcal{O}(\epsilon d_2^{-3/2}\kappa^{-2})$, $|\mathcal{M}_1| = 4(d_1+6)(\sigma_1^2+1)\epsilon^{-2}$, and $|\mathcal{M}_2| = 4(d_2+6)(\sigma_2^2+1)\epsilon^{-2}$, `ZO-SGDA` (Algorithm 1) returns iterates $(x_1, y_1), \ldots, (x_S, y_S)$ such that there exists an iterate which is an ϵ-stationary point for (1) as defined in Definition 2.1. That is, `ZO-SGDA` (Algorithm 1) returns iterates that satisfy $\min_{s\in\{1,\ldots,S\}} \mathbb{E}[\|\nabla g(x_s)\|_2^2] \le \epsilon^2$. Moreover, the total number of calls to the stochastic zeroth-order oracle, $K_{\mathcal{SZO}}$ is given by:*

$$K_{\mathcal{SZO}} = S(|\mathcal{M}_1| + |\mathcal{M}_2|) \sim \mathcal{O}\Big(\kappa^5(d_1\sigma_1^2 + d_2\sigma_2^2)\epsilon^{-4}\Big).$$

We now show that the dependence of the complexity on the condition number κ could be reduced significantly (i.e., from κ^5 to κ^2) by making a simple modification to the `ZO-SGDA` algorithm. Specifically, we run T steps of the ascent part, for every descent step. The approach is presented formally in Algorithm 2 and the corresponding complexity results are provided in Theorem 3.2. The main idea behind running multiple ascent steps is to better approximate the maximum of the stongly-concave function in each step. Subsequently, picking the number of inner iterations T appropriately, helps us obtain improved dependence on κ while still maintaining the same dependency on ϵ. We emphasize that (Nouiehed et al., 2019) used the multi-step ascent approach to handle certain non-convex minimax optimization problems that satisfy the so-called Polyak-Łojasiewicz condition in the first-order setting.

Algorithm 2 Zeroth-Order Stochastic Gradient Multi-Step Descent (`ZO-SGDMSA`)

Initialization: (x_0, y_0), step sizes (η_1, η_2), iteration limits $S > 0$ and $T > 0$, smoothing parameters μ_1 and μ_2. Indices sets $\mathcal{M}_1$ and $\mathcal{M}_2$.

for $s = 1, \ldots, S$ **do**

Set $y_0(x_s) \leftarrow y_s$

for $t = 1, \ldots, T$ **do**

Sample $\boldsymbol{u}_{2,i} \sim N(0, 1_{d_2})$ and compute

$$y_t(x_s) \leftarrow \text{Proj}_{\mathcal{Y}}[y_{t-1}(x_s) + \eta_2 \frac{1}{|\mathcal{M}_2^t|} \sum_{i\in\mathcal{M}_2^t} H_{\mu_2}(x_s, y_{t-1}(x_s), \boldsymbol{u}_{2,i}, \xi_i)]$$

end for

$y_{s+1} \leftarrow y_T(x_s)$

$x_{s+1} \leftarrow x_s - \eta_1 \frac{1}{|\mathcal{M}_1^s|} \sum_{i\in\mathcal{M}_1^s} G_{\mu_1}(x_s, y_{s+1}, \boldsymbol{u}_{1,i}, \xi_i)$

with $\boldsymbol{u}_{1,i} \sim N(0, 1_{d_1})$

end for

Return $(x_1, y_1), \ldots, (x_S, y_S)$.

Theorem 3.2. *Let $\epsilon \in (0,1)$. Under Assumptions 2.1 and 2.2, by setting the parameters as $\eta_1 = 1/(12L_g)$, $\eta_2 = 1/(6\ell)$, $T = \mathcal{O}(\kappa \log(\epsilon^{-1}))$, $S = \mathcal{O}(\kappa\epsilon^{-2})$, $\mu_1 = \mathcal{O}(\epsilon d_1^{-3/2})$, $\mu_2 = \mathcal{O}(\kappa^{-1/2} d_2^{-3/2}\epsilon)$, $|\mathcal{M}_1| = 4(d_1+6)(\sigma_1^2+1)\epsilon^{-2}$, and $|\mathcal{M}_2| = 4(d_2+6)(\sigma_2^2+1)\epsilon^{-2}$, `ZO-SGDMSA` (Algorithm 2) returns iterates $(x_1,y_1),\ldots,(x_S,y_S)$ such that there exists an iterate which is an ϵ-stationary point for (1) as defined in Definition 2.1. That is, `ZO-SGDMSA` returns iterates satisfying $\min_{s\in\{1,\ldots,S\}} E[\|\nabla g(x_s)\|_2^2] \le \epsilon^2$. Moreover, the total number of calls to the zeroth-order oracle is given by:*

$$K_{SZO} = S|\mathcal{M}_1| + TS|\mathcal{M}_2| = \mathcal{O}\Big(\kappa\epsilon^{-4}(d_1\sigma_1^2 + \kappa d_2\sigma_2^2\log(\epsilon^{-1}))\Big).$$

4. Numerical Results

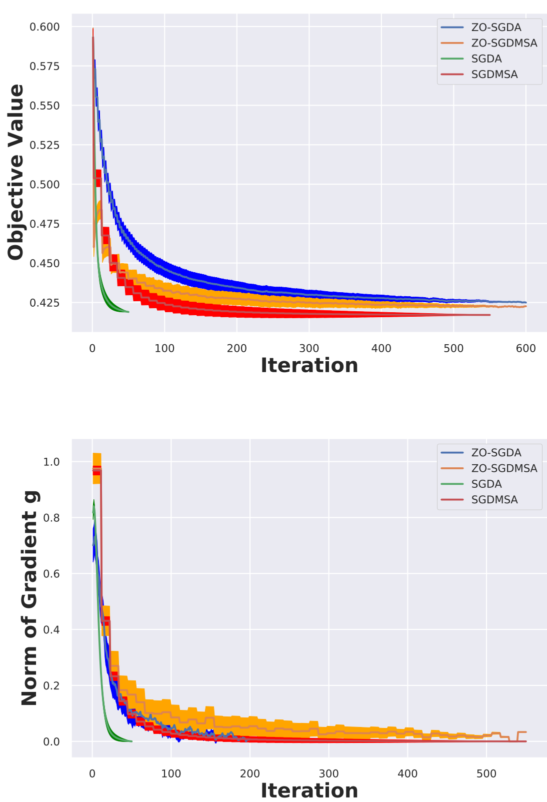

Figure 1. Performance of `ZO-SGDA` and `ZO-SGDMSA` in comparison to their first-order counterparts. Both figures are for the Colon Cancer dataset. The result corresponds to average over 500 trails. More results on other datasets are provided in Section D.

We now compare `ZO-SGDA` and `ZO-SGDMSA` with their first order methods (i.e., `SGDA` and `SGDMSA`) on the distributionally robust optimization problem (Namkoong & Duchi, 2016). For simplicity, we present the formulation of the problem in the finite-sum setting as:

$$\min_{x\in\mathbb{R}^d}\max_{y\in\mathcal{Y}} \sum_{i=1}^{n} y_i\ell_i(x) - r(y)$$

where $\mathcal{Y} = \{y \in \mathbb{R}^n; \sum_{i=1}^n y_i = 1, y_i \ge 0\}$ is the probability simplex; $r(y) = 10\sum_{i=1}^n (y_i - 1/n)^2$ is a divergence measure; $\ell_i(x) = f_1(f_2(x, s_i, z_i))$ where $f_1(x) = \log(1+x)$, $f_2(x) = \log\Big(1+\exp[-z_i(x^T s_i)]\Big)$, (s_i, z_i) is the feature and label pair of a sample i in the dataset. It is easy to see that the above problem is a nonconvex-strongly concave problem with $d_1 = d, d_2 = n$. For the tuning parameters, motivated by our theoretical results, we set the batch size $|\mathcal{M}_1| = d_1/\epsilon^2$ and $|\mathcal{M}_2| = d_2/\epsilon^2$ with $\epsilon = 0.01$. For `ZO-SGDA`, we choose $\eta_1 = \eta_2 = 0.01$. For `ZO-SGDMSA`, we choose $\eta_1 = 0.01$ and $\eta_2 = 0.001$. For `SGDA` and `SGDMSA`, we choose the same stepsize as `ZO-SGDA` and `ZO-SGDMSA` and set $|\mathcal{M}_1| = 1/\epsilon^2$ and $|\mathcal{M}_2| = 1/\epsilon^2$. We stop the iteration when $\|\nabla g(x)\|_2 \le \epsilon$, based on our theoretical analysis. We test our algorithm on datasets from UCI ML-repository (Dua & Graff, 2017) and LIBSVM (Chang & Lin, 2011). All the experiments were run on Google Colab Python 3.5 Notebook. We also remark that we cannot compare empirically to (Liu et al., 2020) as they consider constrained minimax optimization problems. In Figure 1, we plot the value of the objective versus iteration count and the value of gradient size versus iteration count, in the top and bottom rows respectively. From the results we find that the proposed zeroth-order methods perform favorably to their respective first-order counterparts in terms of both the objective value and the norm of the gradient of the function g, in terms of iteration count. It should be noted that to obtain this comparable behavior, the zeroth-order method uses a mini-batch of samples that depends on the dimension in each iteration (as expected), which results in the number of calls to the zeroth-order oracle of the order as illustrated in our theoretical results.

5. Conclusions

In this paper, we analyzed zeroth-order algorithms for stochastic nonconvex minimax optimization problems. Specifically, we considered two types of algorithms: standard single-step gradient descent ascent algorithm and a modified version with multiple ascent steps following each descent step. We obtain oracle complexities for both algorithms that match the performance of comparable first-order algorithms, up to unavoidable dimensionality factors.

Acknowledgements. This work was supported in part by NSF grants DMS-1953210, CCF-2007797, and UC Davis CeDAR Innovative Data Science Seed Funding Program.

References

Agarwal, A., Beygelzimer, A., Dudik, M., Langford, J., and Wallach, H. A reductions approach to fair classification. In *International Conference on Machine Learning*, pp. 60–69, 2018.

Al-Dujaili, A., Srikant, S., Hemberg, E., and O'Reilly, U.-M. On the application of danskin's theorem to derivative-free minimax optimization. *arXiv preprint arXiv:1805.06322*, 2018.

Anagnostidis, S., Lucchi, A., and Diouane, Y. Direct-search methods for a class of non-convex min-max games. In *AISTATS*, 2021.

Audet, C. and Hare, W. *Derivative-Free and Blackbox Optimization*. Springer, 2017.

Baharlouei, S., Nouiehed, M., and Razaviyayn, M. Rényi fair inference. *International Conference on Learning Representation*, 2019.

Balasubramanian, K. and Ghadimi, S. Zeroth-order (non)-convex stochastic optimization via conditional gradient and gradient updates. In *Advances in Neural Information Processing Systems*, pp. 3455–3464, 2018.

Balasubramanian, K. and Ghadimi, S. Zeroth-order non-convex stochastic optimization: Handling constraints, high-dimensionality, and saddle-points. *arXiv preprint arXiv:1809.06474*, 2019.

Bertsimas, D. and Nohadani, O. Robust optimization with simulated annealing. *Journal of Global Optimization*, 48 (2):323–334, 2010.

Bogunovic, I., Scarlett, J., Jegelka, S., and Cevher, V. Adversarially robust optimization with gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 5760–5770, 2018.

Bubeck, S., Lee, Y. T., and Eldan, R. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 72–85. ACM, 2017.

Chang, C.-C. and Lin, C.-J. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

Chen, P.-Y., Zhang, H., Sharma, Y., Yi, J., and Hsieh, C.-J. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 15–26. ACM, 2017.

Conn, A., Scheinberg, K., and Vicente, L. *Introduction to derivative-free optimization*, volume 8. Siam, 2009.

Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. SBEED: Convergent reinforcement learning with nonlinear function approximation. *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.

Daskalakis, C. and Panageas, I. The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems*, pp. 9236–9246, 2018.

Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. Training GANs with optimism. *International Conference on Learning Representations (ICLR)*, 2018.

Dua, D. and Graff, C. UCI machine learning repository, 2017. URL http://archive.ics.uci.edu/ml.

Filar, J. and Vrieze, K. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.

Flokas, L., Vlatakis-Gkaragkounis, E.-V., and Piliouras, G. Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games. *arXiv preprint arXiv:1910.13010*, 2019.

Ghadimi, S. and Lan, G. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23, 2013.

Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.

Hsieh, Y.-P., Liu, C., and Cevher, V. Finding mixed nash equilibria of generative adversarial networks. *arXiv preprint arXiv:1811.02002*, 2018.

Huang, F., Gao, S., Pei, J., and Huang, H. Accelerated zeroth-order and first-order momentum methods from mini to minimax optimization. *https://arxiv.org/pdf/2008.08170.pdf*, 2020.

Jin, C., Netrapalli, P., and Jordan, M. I. Minmax optimization: Stable limit points of gradient descent ascent are locally optimal. *arXiv preprint arXiv:1902.00618*, 2019.

Lin, T., Jin, C., and Jordan, M. I. On gradient descent ascent for nonconvex-concave minimax problems. *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.

Liu, S., Lu, S., Chen, X., Feng, Y., Xu, K., Al-Dujaili, A., Hong, M., and Obelilly, U.-M. Min-max optimization without gradients: Convergence and applications to adversarial ml. *Proceedings of the 37 th International Conference on Machine Learning (ICML)*, 2020.

Liu, Y., Chen, X., Liu, C., and Song, D. Delving into transferable adversarial examples and black-box attacks. *International Conference on Representation Learning*, 2017.

Lu, S., Tsaknakis, I., Hong, M., and Chen, Y. Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *arXiv preprint arXiv:1902.08294*, 2019.

Luo, L., Ye, H., and Zhang, T. Stochastic recursive gradient descent ascent for stochastic nonconvex-strongly-concave minimax problems. *arXiv preprint arXiv:2001.03724*, 2020.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Menickelly, M. and Wild, S. M. Derivative-free robust optimization by outer approximations. *Mathematical Programming*, pp. 1–37, 2018.

Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2703–2717. SIAM, 2018.

Moriarty, D. E., Schultz, A. C., and Grefenstette, J. J. Evolutionary algorithms for reinforcement learning. *Journal of Artificial Intelligence Research*, 11:241–276, 1999.

Namkoong, H. and Duchi, J. C. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in neural information processing systems*, pp. 2208–2216, 2016.

Nesterov, Y. and Spokoiny, V. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.

Nesterov, Y. E. *Introductory lectures on convex optimization: A basic course*. Applied Optimization. Kluwer Academic Publishers, Boston, MA, 2004. ISBN 1-4020-7553-7.

Neyman, A., Sorin, S., and Sorin, S. *Stochastic games and applications*, volume 570. Springer Science & Business Media, 2003.

Nouiehed, M., Sanjabi, M., Huang, T., Lee, J., and Razaviyayn, M. Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, pp. 14905–14916, 2019.

Oliehoek, F. A., Savani, R., Gallego, J., van der Pol, E., and Groß, R. Beyond local nash equilibria for adversarial networks. *arXiv preprint arXiv:1806.07268*, 2018.

Pfau, D. and Vinyals, O. Connecting generative adversarial networks and actor-critic methods. *arXiv preprint arXiv:1610.01945*, 2016.

Picheny, V., Binois, M., and Habbal, A. A bayesian optimization approach to find nash equilibria. *Journal of Global Optimization*, 73(1):171–192, 2019.

Piliouras, G. and Schulman, L. J. Learning dynamics and the co-evolution of competing sexual species. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

Rafique, H., Liu, M., Lin, Q., and Yang, T. Non-convex min-max optimization: Provable algorithms and applications in machine learning. *arXiv preprint arXiv:1810.02060*, 2018.

Rios, L. and Sahinidis, N. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3): 1247–1293, 2013.

Roy, A., Chen, Y., Balasubramanian, K., and Mohapatra, P. Online and bandit algorithms for nonstationary stochastic saddle-point optimization. *arXiv preprint arXiv:1912.01698*, 2019.

Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.

Sanjabi, M., Ba, J., Razaviyayn, M., and Lee, J. D. On the convergence and robustness of training gans with regularized optimal transport. In *Advances in Neural Information Processing Systems*, pp. 7091–7101, 2018.

Snoek, J., Larochelle, H., and Adams, R. P. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pp. 2951–2959, 2012.

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Thekumparampil, K., Jain, P., Netrapalli, P., and Oh, S. Efficient algorithms for smooth minimax optimization. In *Advances in Neural Information Processing Systems*, pp. 12659–12670, 2019.

Wang, Z. and Jegelka, S. Max-value entropy search for efficient bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3627–3635. JMLR. org, 2017.

Wei, C.-Y., Hong, Y.-T., and Lu, C.-J. Online reinforcement learning in stochastic games. In *Advances in Neural Information Processing Systems*, pp. 4987–4997, 2017.

Xu, D., Yuan, S., Zhang, L., and Wu, X. Fairgan: Fairness-aware generative adversarial networks. In *IEEE International Conference on Big Data (Big Data). IEEE*, pp. 570–575, 2018.

Xu, T., Zhe Wang, Z., Liang, Y., and Poor, H. V. Gradient free minimax optimization: Variance reduction and faster convergence. *https://arxiv.org/pdf/2006.09361.pdf*, 2021.

Ying, Y., Wen, L., and Lyu, S. Stochastic online auc maximization. In *Advances in neural information processing systems*, pp. 451–459, 2016.

Zhang, B. H., Lemoine, B., and Mitchell, M. Mitigating unwanted biases with adversarial learning. In *AAAI/ACM Conference on AI, Ethics, and Society. ACM*, pp. 335–340, 2018.

Zhang, K., Yang, Z., and Başar, T. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *arXiv preprint arXiv:1911.10635*, 2019.