

MBZUAI

Digital.Commons@MBZUAI

Machine Learning Faculty Publications

Scholarly Works

4-9-2021

Learning Sampling Policy for Faster Derivative Free Optimization

Zhou Zhai

Nanjing University of Information Science & Technology

Bin Gu

Mohamed bin Zayed University of Artificial Intelligence & JD Finance America Corporation, USA

Heng Huang

University of Pittsburgh & JD Finance America Corporation

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/mlfp>



Part of the [Computer Sciences Commons](#)

Preprint: arXiv

- Archived with thanks to arXiv
- Preprint License: [CC by](#)
- Uploaded 24 March 2022

Recommended Citation

Z. Zhai, B. Gu, and H. Huang, "Learning sampling policy for faster derivative free optimization," 2021, arXiv:2104.04405v1

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Machine Learning Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Learning Sampling Policy for Faster Derivative Free Optimization

Zhou Zhai, Bin Gu, and Heng Huang

Abstract—Zeroth-order (ZO, also known as derivative-free) methods, which estimate the gradient only by two function evaluations, have attracted much attention recently because of its broad applications in machine learning community. The two function evaluations are normally generated with random perturbations from standard Gaussian distribution. To speed up ZO methods, many methods, such as variance reduced stochastic ZO gradients and learning an adaptive Gaussian distribution, have recently been proposed to reduce the variances of ZO gradients. However, it is still an open problem whether there is a space to further improve the convergence of ZO methods. To explore this problem, in this paper, we propose a new reinforcement learning based ZO algorithm (ZO-RL) with learning the sampling policy for generating the perturbations in ZO optimization instead of using random sampling. To find the optimal policy, an actor-critic RL algorithm called deep deterministic policy gradient (DDPG) with two neural network function approximators is adopted. The learned sampling policy guides the perturbed points in the parameter space to estimate a more accurate ZO gradient. To the best of our knowledge, our ZO-RL is the first algorithm to learn the sampling policy using reinforcement learning for ZO optimization which is parallel to the existing methods. Especially, our ZO-RL can be combined with existing ZO algorithms that could further accelerate the algorithms. Experimental results for different ZO optimization problems show that our ZO-RL algorithm can effectively reduce the variances of ZO gradient by learning a sampling policy, and converge faster than existing ZO algorithms in different scenarios.

I. INTRODUCTION

Gradient based optimization is an important problem in machine learning. However, in many fields of science and engineering, explicit gradient information is difficult or even infeasible to obtain. Zeroth-order (ZO, also known as derivative-free) optimization has attracted an increasing amount of attention, where the optimizer is provided with only function values (zeroth-order information) instead of explicit gradients (first-order information). Specifically, the ZO optimization algorithms first generate perturbed vectors from a (standard) Gaussian distribution. Based on the sampled perturbed vectors, they query the corresponding function values. Then, they can approximate the gradient information based on the technique of finite difference [1]. ZO optimization can theoretically address a wide range of objectives and has been studied in

a large number of fields such as optimization, online learning and bioinformatics [2], [3], [4]. One of the most important applications of ZO optimization is to generate prediction-evasive adversarial examples in the black-box setting [5], [6], e.g., crafted images with imperceptible perturbations to deceive a well-trained image classifier into misclassification.

As mentioned above, standard ZO algorithm constructs a pseudo-gradient by uniformly and randomly sampling some perturbed directions from the standard Gaussian distribution [7], [8]. However, ZO algorithms often suffer from high variances of ZO gradient estimators, and in turn, hampered convergence rates. In the past few decades, many ZO methods have been proposed to overcome this problem to speed up the convergence of ZO optimization, which can be divided into two directions. The first direction is to borrow the gradient descent method used in the first-order algorithm to improve the parameter update rule in ZO optimization [9], [10], [11]. For example, [12] proposed two accelerated versions of ZO-SVRG, utilizing stochastic variance reduced gradient (SVRG) estimators to reduce the variance. [11] extended sign-based stochastic gradient descent (signSGD) method to ZO optimization, compressing the gradient with a single bit by sign operation to mitigate the negative effects of the extreme components of the gradient noise and speed up convergence rate. [5] extended SVRG under Gaussian smoothing to reduce the variances of both random samples and random query directions.

The second direction is to utilize the learned sampling distribution to sample perturbed vectors. Different from sampling perturbed vectors from a standard Gaussian distribution $\mathcal{N}(0, I_d)$, they consider generating perturbed vectors with non-isotropic Gaussian distribution $\mathcal{N}(0, \Sigma)$ such that the covariance matrix Σ may not be a scale of the identity matrix. For example, in a black-box adversarial attack task, there is usually a well-defined significant subspace that can be successfully attacked, and perturbed directions through this subspace leads to a faster convergence. Similar to this, the evolutionary strategies (ES) such as Natural ES [13], CMA-ES [14], and Guided ES [15] are proposed to guide the estimation of a descent direction that then can be passed to the ZO optimizer. [16] proposed the learning to learn framework to adaptively modify the search distribution with learned recurrent neural networks.

As discussed above, the existing ZO algorithms based on the parameter update rule (i.e., the first direction) can reduce the variance of ZO gradient, while they normally lead to a higher query complexity. The ZO algorithms based on the learned sampling distribution (i.e., the second direction) also

Z. Zhai is with School of Computer & Software, Nanjing University of Information Science & Technology, Nanjing, P.R.China (e-mail: zhouzhai@nuist.edu.cn).

B. Gu is with the department of machine learning, Mohamed bin Zayed University of Artificial Intelligence, UAE, and with JD Finance America Corporation, Mountain View, CA 94043 USA (jsgubin@gmail.com).

H. Huang is with Department of Electrical & Computer Engineering, University of Pittsburgh, USA, and with JD Finance America Corporation (e-mail: heng.huang@pitt.edu).

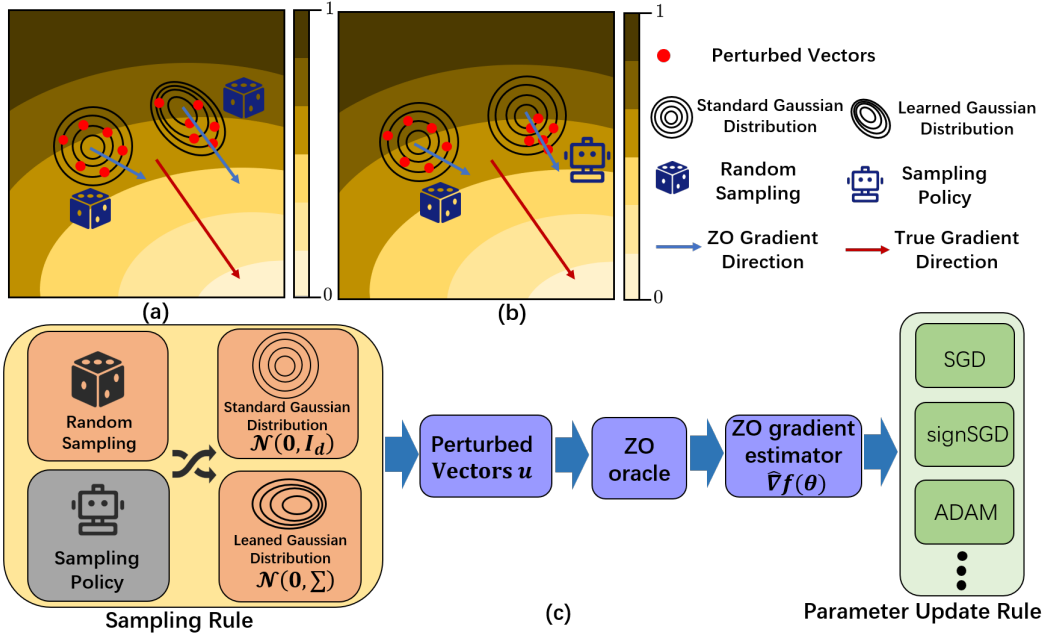


Fig. 1: (a) Comparison of the ZO gradient directions obtained by sampling perturbed vectors from the standard Gaussian distribution and a learned Gaussian distribution. (b) Comparison of the ZO gradient directions obtained by a random sampling policy and a learned sampling policy. (c) The architecture of ZO optimizer. We classify the existing ZO optimization algorithms into two categories (i.e., the sampling rule and the parameter update rule). For sampling rule, we can learn the rules for sampling distribution and policy. Learning the sampling policy is our main contribution in this paper and it is parallel to existing techniques. Importantly, our learned sampling policy can be combined with existing techniques to further accelerate ZO optimization.

can reduce the variance of ZO gradient, but they only care about the sampling distribution, and ignore the importance of sampling policy. We provide Fig 1.(a) to intuitively show the benefit of using the learned sampling distribution, and also provide Fig 1.(b) to intuitively demonstrate the benefit of using the learned sampling policy. Since learning the sampling distribution has been proved that it can improve the convergence of ZO optimization, naturally we have the following question:

Can we learn a sampling policy via using the generating perturbed vectors, instead of simply using random sampling policy, to further speed up the convergence of ZO optimization?

In this paper, we provide a positive answer to this question, meanwhile provide a feasible solution to achieve this goal.

Specifically, to solve this challenging problem, we propose a reinforcement learning based zeroth-order algorithm (ZO-RL) to learn the sampling policy in ZO optimization. Reinforcement learning (RL) is a self-adaptive model and has achieved appealing achievements in practical applications, including playing the “Atari games” [17], defeating professional Go players [18], and 3D manipulation of robots [19]. The policy gradient method is a frequently used algorithm in RL. Model-free policy gradient algorithms are divided into deterministic and stochastic policies. Compared with the stochastic policy gradient algorithm, the deterministic policy has the advantages of requiring less data to be sampled and high efficiency of the algorithm, and performs stably in a series of tasks with contin-

uous action space. Thus, to find the optimal policy, an actor-critic RL algorithm called deep deterministic policy gradient (DDPG) [20] with two neural network function approximators is adopted. The learned sampling policy of RL guide the optimizer to estimate more accurate ZO gradients in the parameter space. Especially, we can combine our ZO-RL with existing the ZO algorithms that utilize the improved parameter update rule and learned sampling distribution. Experimental results for different ZO optimization problems show that our ZO-RL algorithm can effectively reduce the variances of ZO gradient by learning the sampling policy, and converge faster than existing ZO algorithms in different scenarios.

Contributions. The main contributions of this paper are summarized as follows.

- 1) We propose to learn the sampling policy by reinforcement learning instead of using random sampling policy as in the traditional ZO algorithms to generate perturbed vectors. To the best of our knowledge, our ZO-RL is the first algorithm to learn the sampling policy for ZO optimization which is parallel to the existing ZO methods.
- 2) We conduct extensive experiments to show that our ZO-RL algorithm can effectively reduce the variances of ZO gradients by learning a sampling policy, and converge faster than existing ZO algorithms in different scenarios.

II. PRELIMINARIES

In this section, we give a brief review of zeroth-order optimization and reinforcement learning respectively.

A. Zeroth-Order Optimization

ZO optimization is widely used in the environments where gradient information is difficult or even impossible to obtain. For the loss function f with its parameter x , we can obtain its ZO gradient estimator by:

$$\hat{\nabla} f(x) = \frac{1}{\mu q} \sum_{i=1}^q [f(x + \mu u_i) - f(x)] u_i \quad (1)$$

where $\mu > 0$ is the smoothing parameter, $\{u_i\}$ are the random query directions drawn from the standard Gaussian distribution, and q is the number of sampled query directions. We summarized the standard ZO optimization algorithm in Algorithm 1. Fig 1.(c) shows the architecture of ZO optimizer. The perturbed vectors (query directions) are first generated by the sampling rule, then passed to ZO Oracle to calculate the ZO gradient estimator. Based on the ZO gradients, we can update solution based on the parameter update rule.

The high variance of ZO gradient estimators hinders the convergence speed of ZO algorithms due to the random sampling perturbed vectors. Thus, the choice of sampling policy and sampling distribution determines the performance of the ZO optimization algorithm. There has been a research trend to propose sampling the perturbed vectors from some non-isotropic Gaussian distribution instead of standard Gaussian sampling with an isotropic covariance for random query directions. They consider sampling perturbed vectors by $u_i \sim \mathcal{N}(0, \Sigma)$ such that the covariance Σ may not be a scale of the identity matrix [21], [16], [15]. By learning the significant sampling distribution, more accurate gradient estimators can be obtained for a fixed query budget, which can improve the convergence rate of the ZO optimization task. However, improving the performance of ZO optimization algorithms by learning sampling policy is still a vacancy in the literature. To overcome this problem, in this paper, we use a policy search approach in reinforcement learning to learn a sampling policy, and then use it to generate perturbed vectors to obtain more accurate gradient estimators, instead of plainly using the random sampling policy to generate perturbed vectors.

Algorithm 1 Zeroth-Order (ZO) Optimization Algorithm

Input: Hyper-parameter μ , q and η .

Output: $x \in \mathbb{R}^d$

- 1: **for** $k = 0$ to $K - 1$ **do**
 - 2: Sampling q perturbed vectors from the standard Gaussian distribution $u_i \sim \mathcal{N}(0, I_d)$.
 - 3: Calculating the ZO gradient estimator $\hat{\nabla} f(x_k)$.
 - 4: Obtain the next update $x_{k+1} = x_k - \eta \hat{\nabla} f(x_k)$.
 - 5: **end for**
-

B. Reinforcement Learning

RL can be modeled as a Markov decision process (MDP) with a four-tuple of (S, A, P, R) , where S means a set of states, A denotes a set of actions, P represents a transition probability function $p(s_{t+1}|s_t, a_t)$, and R is a reward function $R : S \times A \rightarrow \mathbb{R}$. An episode of task denotes that an agent and an environment interact with each other at discrete

time steps $t = 0, 1, 2, \dots, T$. The agent chooses an action $a_t \in A$ according to its policy under the state $s_t \in S$ of the environment. If the agent takes a certain action a_t , the environment translates its state from s_t to s_{t+1} responding to the action and the agent also obtains a reward $r_t \in \mathbb{R}$. The agent's objective is to learn an optimal policy π^* so as to maximize the expected accumulative rewards. Let $\eta(\pi)$ denote the expected cumulative reward:

$$\eta(\pi) = \mathbb{E}_{s_0, a_0, \dots} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right] \quad (2)$$

where $\gamma \in [0, 1]$ is a discounting factor. The deterministic policy is defined as:

$$\pi : S \rightarrow A \quad (3)$$

which is a mapping from the state space to the action space. Let $\pi(a_t|s_t)$ represent the conditional probability that the agent takes action a_t given the states s_t .

To apply the above RL framework to ZO optimization, we first need define some features to represent the state of ZO gradient, a set of actions to represent the sampling rule, and a reward function. Then, we should choose an appropriate algorithm to find the optimal policy π^* to maximize our expected cumulative reward. In the next section, we will discuss these in detail.

III. LEARNING SAMPLING POLICY IN ZEROTH-ORDER OPTIMIZATION

Considering that the learned sampling policy may perform better in ZO optimization compared to random sampling, we learn the sampling policy π in ZO optimization. We observe that the execution of ZO optimization algorithm can be viewed as the execution of a fixed policy in a MDP: the state consists of the current function value and ZO gradients evaluated at the current and past function values, the action is the step vector that is used to update the current parameter, and the transition probability is partially characterized by the parameter update formulation. The policy that is executed corresponds precisely to the choice of π used by the ZO optimization algorithm. Thus, we use the RL to learn the policy π . For this purpose, we need to define the reward function that should reward those policies that exhibit good behavior during execution. Since the performance metric of interest to the ZO optimization algorithm is the speed of convergence, the reward function should reward policies that converge quickly. To this end, assuming the goal is to minimize the objective function, we define the reward in a given state as the value of the ZO gradient. This will encourage the policy to reach the minimum of the objective function as soon as possible. Therefore, the learned sampling policy by maximizing expected cumulative reward can effectively reduce the variances of ZO gradient.

In the following, we first introduce the principle of our ZO-RL algorithm. Then, we introduce the network structure and the batch normalization technique which are used in our ZO-RL algorithm.

A. Principle of our ZO-RL Algorithm

Since the action space is continuous in ZO optimization, we use the deterministic policy. Compared with the stochastic policy gradient algorithm, the deterministic policy has the advantages of requiring less data to be sampled and high efficiency of the algorithm, and performs stably in a series of tasks with continuous action space. Thus, to find the optimal policy, we use deep deterministic policy gradient (DDPG) to learn query policy π . DDPG is an actor-critic and model-free algorithm for RL over continuous action spaces and output deterministic actions in a stochastic environment to maximize cumulative rewards.

The DDPG has two neural network function approximators. One is called the actor network $\pi(s|\theta^\pi)$ with weights θ^π which learns a policy function of mapping a state to a deterministic action. The other is called the critic network $Q(s, a|\theta^Q)$ with weights θ^Q , which can learn a state-action value function and its input consists of a state and action. The actor-critic methods combine the advantages of critic-based and actor-based methods. These methods estimate the parameters of the critic and the actor simultaneously. The critic learns a value function, which is used to measure whether the current action is improved compared with the policy's default behavior. The parameters of the actor's policy are updated in a direction advised by the critic evaluation. The parameters of two structures are simultaneously estimated to find out the optimal policy. In addition, DDPG creates a copy of the actor and critic networks, $Q'(s, a|\theta^{Q'})$ and $\pi'(s|\theta^{\pi'})$ respectively, that are used for calculating the target values. The weights of these target networks are then updated by making them slowly track the learned networks: $\theta' \rightarrow \tau\theta' + (1-\tau)\theta$ with $\tau \ll 1$. This means that the target values are constrained to change slowly, greatly improving the stability of learning. This simple change moves the relatively unstable problem of learning the action-value function closer to the case of supervised learning. The updates at each iteration contain the critic update and actor update. Updating the critic is to minimize a squared-error loss L :

$$\min_{\theta^Q} L = \frac{1}{N} \sum_{i=1}^N (y_i - Q(s_i, a_i|\theta^Q))^2 \quad (4)$$

where N is the mini batch size, $Q(\cdot|\theta^Q)$ represents a parameterized state-action value function, and y_i represents the TD target denoted as

$$y_i = r_i + \gamma Q'(s_{i+1}, \pi'(s_{i+1}|\theta^{\pi'})|\theta^{Q'}) \quad (5)$$

Updating the actor can help maximize the cumulative reward using a sampled policy gradient:

$$\nabla_{\theta^\pi} J = \frac{1}{N} \sum_{i=1}^N \nabla_a Q(s, a|\theta^Q)|_{s=s_i, a=\pi(s_i)} \nabla_{\theta^\pi} \pi(s|\theta^\pi)|_{s_i} \quad (6)$$

where J and $\pi(\cdot|\theta^\pi)$ represent the expected cumulative reward and parameterized policy function respectively. We summarize our ZO-RL optimization algorithm in Algorithm 2.

Algorithm 2 Reinforcement Learning Based Zeroth-Order (ZO-RL) Algorithm

Input: Randomly initialize critic network $Q(s, a|\theta^Q)$ and actor $\pi(s|\theta^\pi)$. Hyper-parameter μ , q and η .

Output: $x \in \mathbb{R}^d$

- 1: Using DDPG to learn optimal sampling policy π^* from the standard Gaussian distribution.
 - 2: **for** $k = 0$ to $K - 1$ **do**
 - 3: Select action $a_k = \pi^*(s_k|\theta^{\pi^*})$ according to the sampling policy π^* .
 - 4: Calculating the ZO gradient estimator $\hat{\nabla} f(x_k)$.
 - 5: Obtain the next update $x_{k+1} = x_k - \eta \hat{\nabla} f(x_k)$.
 - 6: **end for**
-

B. Network Structure

The choice of the structure of the critic and actor nets is important because they are not only function approximators but also part of the feature learning. We choose the convolutional neural network (CNN) [22] both for the critic net and the actor net. CNN has a large variety of applications in image-classification, video-recognition, and also “natural language processing”. Generally, CNN consists of three types of layers: the convolutional, pooling and multilayer perceptron (MLP). A convolutional and pooling layer are structurally successive. In the “convolution layer”, the convolution operation is carried out, results are passed to the pooling layer. In the pooling layer, both the “number of parameters” and “the spatial size of representation” are reduced. In the last convolutional or pooling layer, the data forms a one-dimensional vector and is connected to an MLP. In other word, convolutional and pooling layers perform an implicit feature extraction, and MLP performs a traditional classifier. The input matrix can be viewed as a one-dimensional image with n channels, that is, each technical indicator represents a channel. For the critic network, the input is a state s_t and action a_t , and its output is state-action value function $Q(s_t, a_t|\theta^Q)$. For the actor network, the input is the state s_t and the output is the probabilities of taking actions.

C. Batch Normalization

The parameters of the optimizer have different descent rates in different dimensions and the range may be different in different environments. This may make it difficult for the network to learn efficiently and find hyper-parameters that generalize the scale of state values in different environments. One way to address this issue is to manually scale features so that they are in a similar range across environments and units. We address this problem by adapting one of the latest techniques in deep learning, called batch normalization [23]. This technique normalizes each dimension of a sample in a mini-batch to have unit mean and variance. In addition, batch normalization maintains a running average of the mean and variance to be used for normalization during testing. In deep networks, batch normalization is used to minimize the covariance bias during training, by ensuring that each layer receives whitened inputs.

IV. COMBINING OUR ZO-RL WITH EXISTING ZO OPTIMIZATION ALGORITHMS

In this section, we discuss how to combine our ZO-RL with an existing ZO algorithm based on the parameter update rule or the learned sampling distribution. For example, [16] proposed a ZO optimization algorithm called ZO-LSTM, which replaces parameter update rule as well as guided sampling rule to sample the perturbed samples with learned recurrent neural networks (RNN). Especially, they updated the parameter through a Long Short-Term Memory (LSTM) network called UpdateRNN:

$$x_k = x_{k-1} + \text{UpdateRNN}(\hat{\nabla} f(x_k)) \quad (7)$$

where x_k is the optimizer parameter at iteration k . UpdateRNN can reduce the negative impact of high variances of ZO gradient estimator due to long-term dependence, in addition to learning to compute parameter updates adaptively by exploring the loss landscape. They use another LSTM network called QueryRNN to learn the sampling rules for distributions. They dynamically predict the covariance matrix Σ_k :

$$\Sigma_k = \text{QueryRNN}([\hat{\nabla} f(x_k), \Delta x_{k-1}]) \quad (8)$$

QueryRNN can increase the sampling probability in the direction of the bias of the estimated gradient or the parameter update of the previous iteration.

Although they considered both the sampling distribution and the parameter update rule, they still used random sampling for the perturbed samples. Using the learned sampling policy on the learned sampling distribution can further speed up the convergence of ZO optimization. Thus, we can combine our ZO-RL algorithm with ZO-LSTM algorithm. Especially, we first train the UpdateRNN using standard Gaussian random vectors as query directions. Then we freeze the parameters of the UpdateRNN and train the QueryRNN. Finally, we use the previous work as a warm start and use our ZO-RL in the pre-learning distribution to learn the sampling policy. In addition, other ZO algorithms based on parameter update rules such as [9], [10], can be directly combined with our ZO-RL algorithm.

V. EXPERIMENTS

In this section, we empirically demonstrate the superiority of our proposed ZO optimizer on a practical application (black-box adversarial attack on MNIST dataset) and a synthetic problem (non-convex optimization problems on benchmark datasets). To show the effectiveness of the learned sampling policy, we compare our ZO optimizer with existing ZO optimization algorithms. Specifically, we consider the convergence behavior and change in variances of ZO optimizer under three different parameter update settings. The three different parameter update settings are summarized as follows:

- 1) SGD [24]: A gradient descent algorithm based on the ZO gradient estimators.
- 2) signSGD [11]: A gradient descent algorithm based on the sign of the ZO gradient estimators.
- 3) ADAM [10]: A gradient descent algorithm based on adaptive estimates of lower-order moments.

We obtain ZO gradient estimator along sampled directions via ZO Oracle. Since our algorithm is the first one to learning the sampling policy, we compare the performance of the ZO gradient estimators sampled from the standard Gaussian distribution and two learned Gaussian distributions, *i.e.* using different covariance matrix Σ . In addition, we compare the algorithm of synchronous learning sampling and distribution policy by combining our algorithm with other algorithms. The five algorithms for calculating ZO gradient estimators are summarized as follows:

- 1) ZO-GS [25]: Randomly Sampling the perturbed vectors u_i from a standard Gaussian distribution.
- 2) ZO-LSTM [16]: They learned the Gaussian sampling rule and dynamically predicted the covariance matrix Σ for query directions with recurrent neural networks.
- 3) Guided ES [15]: They let the covariance matrix Σ be related with the recent history of surrogate gradients during optimization.
- 4) ZO-RL: Our proposed ZO algorithm learns the sampling policy through reinforcement learning.
- 5) ZO-RL-LSTM: Our proposed ZO algorithm combined with ZO-LSTM to learn sampling policy on a learned Gaussian distribution.

A. Implementation

For each task, we tune the hyper-parameters of baseline algorithms to report the best performance. We coarsely tune the constant δ on a logarithmic range $\{0.01; 0.1; 1; 10; 100; 1000\}$ and set the learning rate of baseline algorithms to $\eta = \delta/d$, where d is the dimension of dataset. For ADAM, we tune β_1 values over $\{0.9, 0.99\}$ and β_2 values over $\{0.99, 0.996, 0.999\}$. We set the smoothing parameter $\mu = 0.01$ in all experiments. To ensure fair comparison, all optimizers are using the same number of query directions to obtain ZO gradient estimator at each iteration.

B. Adversarial Attack to Black-box Models

We consider generating adversarial examples to attack black-box DNN image classifier and formulate it as a zeroth-order optimization problem. The targeted DNN image classifier $F(x) = [F_1, F_2, \dots, F_K]$ takes as input an image $x \in [0, 1]^d$ and outputs the prediction scores of K classes. Given an image $x_0 \in [0, 1]^d$ and its corresponding true label $t_0 \in [1, 2, \dots, K]$, an adversarial example x is visually similar to the original image x_0 but leads the targeted model F to make wrong prediction other than t_0 . The black-box attack loss is defined as:

$$\min_x \max \{F_{t_0}(x) - \max_{j \neq t_0} F_j(x), 0\} + c \|x - x_0\|_p \quad (9)$$

The first term is the attack loss which measures how successful the adversarial attack is and penalizes correct prediction by the targeted model. The second term is the distortion loss (p -norm of added perturbation) which enforces the perturbation added to be small and c is the regularization coefficient. In our experiment, we use ℓ_1 norm (*i.e.*, $p = 1$), and set $c = 0.1$ for MNIST attack task. Due to the black-box setting, one can

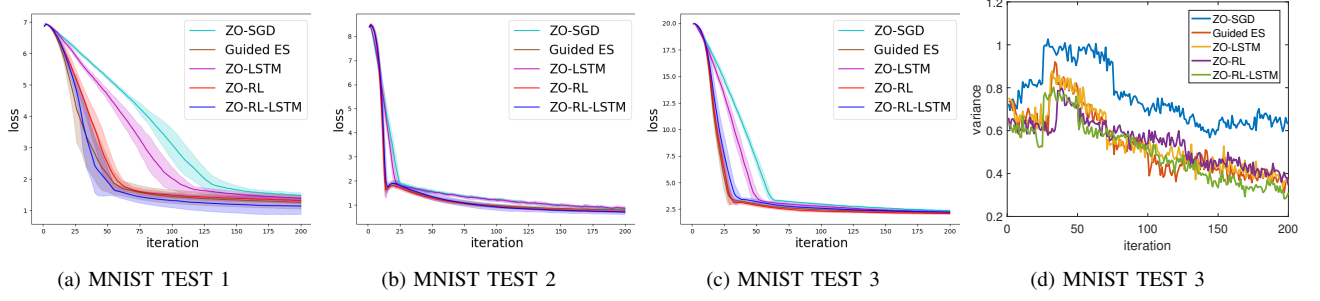


Fig. 2: Adversarial attack to black-box models in the SGD setting.

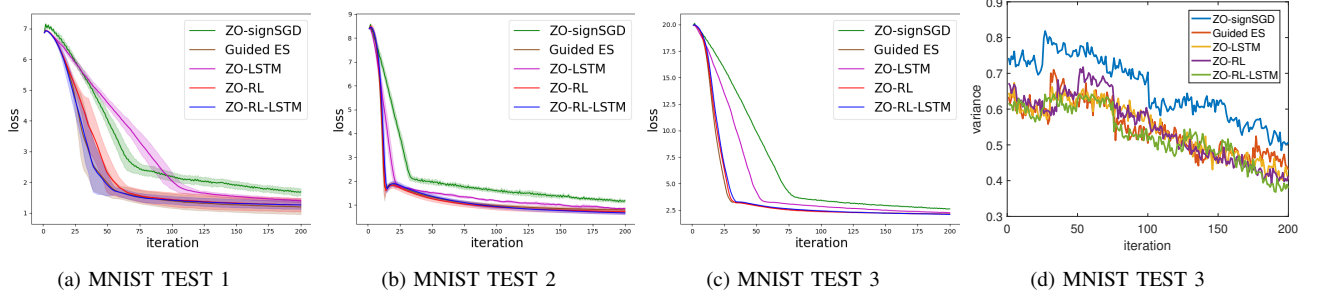


Fig. 3: Adversarial attack to black-box models in the signSGD setting.

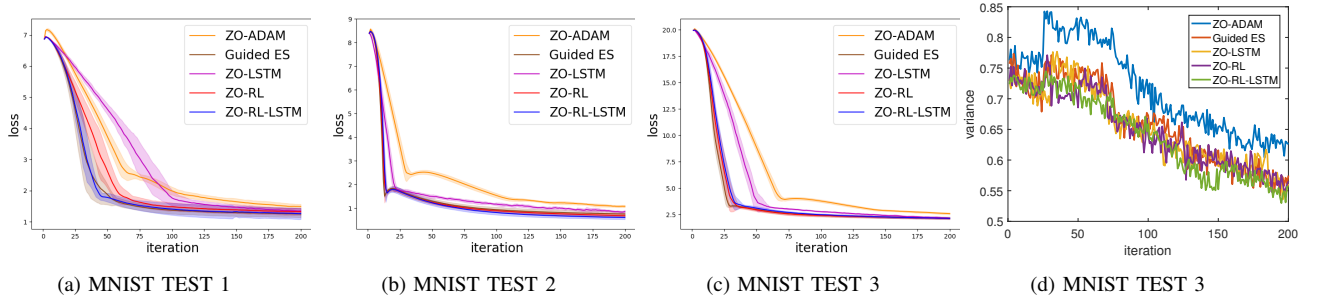


Fig. 4: Adversarial attack to black-box models in the ADAM setting.

only compute the function values of the above objective, which leads to ZO optimization problems [10]. Note that attacking each sample x_0 in the dataset corresponds to a particular ZO optimization problem instance, which motivates us to train a ZO optimizer offline with a small subset, and apply it to online attack to other samples with faster convergence (which means lower query complexity) and lower final loss (which means less distortion). We randomly select 50 images that are correctly classified by the targeted model in each test set to train the optimizer and select another 50 images to test the learned optimizer. The number of sampled query directions is set to $q = 20$ for MNIST, and the optimizer is allowed to run 200 steps.

C. Non-Convex Optimization Problems

We consider a binary classification problem with a non-convex least squared loss function $\min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n (y_i - 1/(1 + e^{-w^T x_i}))^2$. Here (x_i, y_i) is the i th data sample containing feature $x_i \in \mathbb{R}^d$ and label $y_i \in \{-1, 1\}$. We compare the algorithms on benchmark datasets (heat scale, german and

a9a¹). All the algorithms can only access to the oracle of function value evaluations. We use the same set of hyper-parameters for different datasets and repeated runs in the experiments. The number of query directions are set to $q = 20$. For each dataset, we repeat the experiment 10 times and report the average and the standard deviation. At each iteration of training, the optimizer is allowed to run 200 steps.

D. Discussion and Analysis

Fig. 2 shows the black-box attack loss and variance versus iterations using different ZO optimizers in the SGD setting. Fig. 3 visualizes the black-box attack loss and variance versus iterations using different ZO optimizers in the signSGD setting. Fig. 4 plots the black-box attack loss and variance versus iterations using different ZO optimizers in the ADAM setting. The loss curves are averaged over 10 independent random trails and the shaded areas indicate the standard deviation.

¹<http://archive.ics.uci.edu/ml/datasets.html>

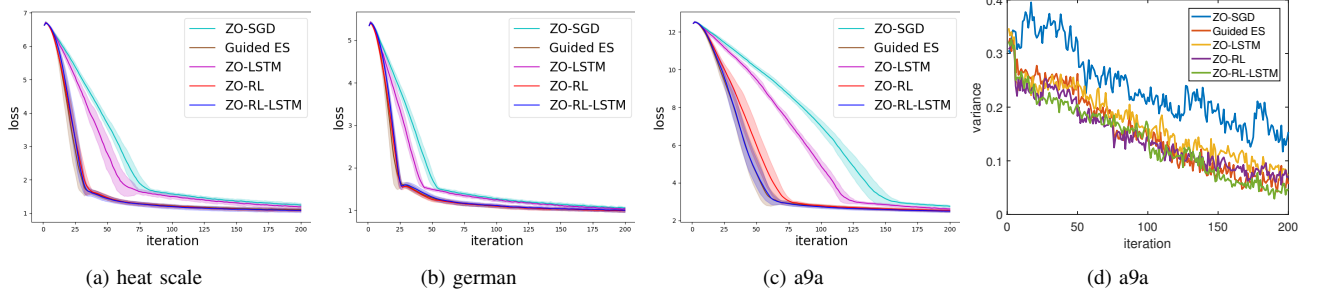


Fig. 5: Non-convex optimization problems in the SGD setting.

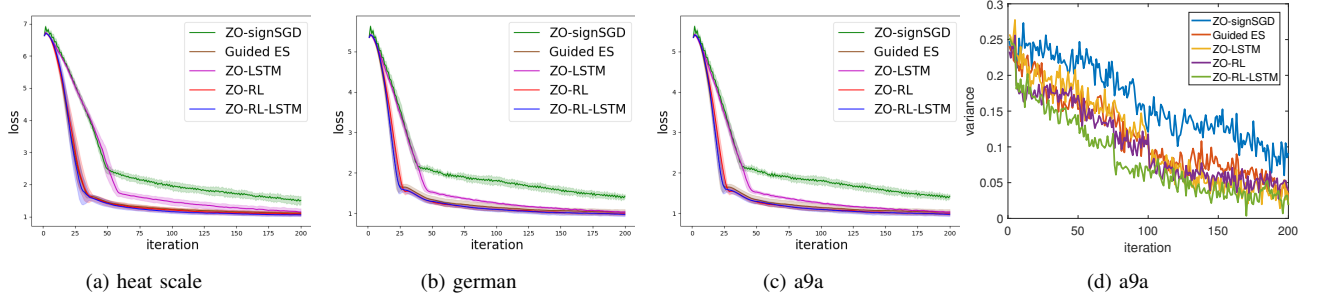


Fig. 6: Non-convex optimization problems in the signSGD setting.

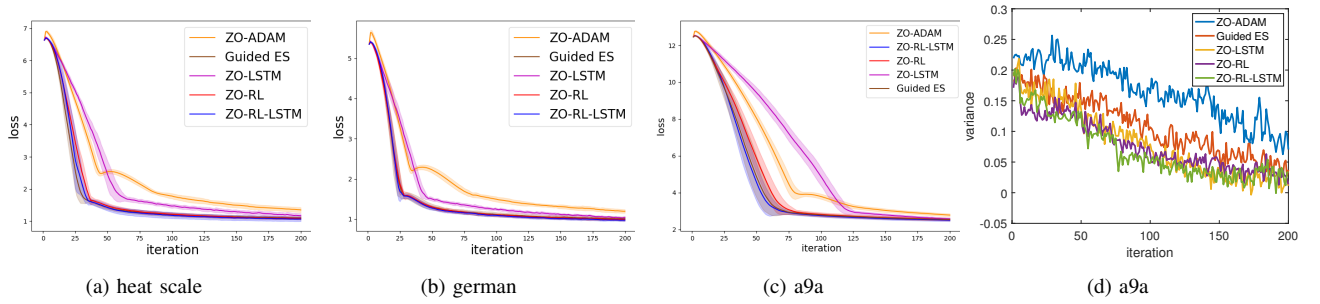


Fig. 7: Non-convex optimization problems in the ADAM setting.

The results clearly show that our ZO-RL has significant advantage over random sampling to sample perturbed vectors, and outperforms the ZO algorithms that use the learned sampling distribution most of the time. This is due to the fact that our ZO-RL learn the sampling policy instead of random sampling that can reduce the variances of ZO gradient.

Fig. 5 draws the non-convex least squared loss and variance versus iterations using different ZO optimizers in the SGD setting. Fig. 6 demonstrates the non-convex least squared loss and variance versus iterations using different ZO optimizers in the signSGD setting. Fig. 7 illustrates the non-convex least squared loss and variance versus iterations using different ZO optimizers in the ADAM setting. The loss curves are averaged over 10 independent random trails and the shaded areas indicate the standard deviation. Our ZO-RL has a definite advantage in sampling the perturbed vectors compared with random sampling, and our ZO-RL-LSTM can always obtain the best results by combining learned sampling policy and sampling distribution. The results clearly demonstrate that our ZO-RL leads to much faster convergence and lower final loss

under different parameter update settings compared to existing ZO algorithms. The results also show that our ZO-RL can effectively reduce the variances of ZO gradient by learning the sampling policy that maximizes expected cumulative reward.

VI. CONCLUSION

We proposed a new reinforcement learning based sampling policy for generating the perturbations in ZO optimization instead of using the existing random sampling. The learned sampling policy guides the perturbation (direction) in the parameter space to estimate a more accurate ZO gradient. To the best of our knowledge, our ZO-RL is the first algorithm to learn the sampling policy via reinforcement learning for ZO optimization which is parallel to the existing methods. Especially, our ZO-RL can be combined with the existing ZO algorithms that could further accelerate them. Experimental results of solving different ZO optimization problems show that our ZO-RL algorithm effectively reduces the variances of ZO gradient by learning the sampling policy, and converges faster than existing ZO algorithms in different scenarios.

REFERENCES

- [1] X. Chen, S. Liu, K. Xu, X. Li, X. Lin, M. Hong, and D. Cox, “Zo-adam: Zeroth-order adaptive momentum method for black-box optimization,” *arXiv preprint arXiv:1910.06513*, 2019.
- [2] P. Koch, O. Golovidov, S. Gardner, B. Wujek, J. Griffin, and Y. Xu, “Autotune: A derivative-free optimization framework for hyperparameter tuning,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 443–452.
- [3] H. Mania, A. Guy, and B. Recht, “Simple random search of static linear policies is competitive for reinforcement learning,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 1805–1814.
- [4] A. Vemula, W. Sun, and J. Bagnell, “Contrasting exploration in parameter and action space: A zeroth-order optimization perspective,” in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 2926–2935.
- [5] L. Liu, M. Cheng, C.-J. Hsieh, and D. Tao, “Stochastic zeroth-order optimization via variance reduction method,” *arXiv preprint arXiv:1805.11811*, 2018.
- [6] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017, pp. 506–519.
- [7] Y. Nesterov and V. Spokoiny, “Random gradient-free minimization of convex functions,” *Foundations of Computational Mathematics*, vol. 17, no. 2, pp. 527–566, 2017.
- [8] J. C. Duchi, P. L. Bartlett, and M. J. Wainwright, “Randomized smoothing for stochastic optimization,” *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 674–701, 2012.
- [9] X. Lian, H. Zhang, C.-J. Hsieh, Y. Huang, and J. Liu, “A comprehensive linear speedup analysis for asynchronous stochastic parallel optimization from zeroth-order to first-order,” *Advances in Neural Information Processing Systems*, vol. 29, pp. 3054–3062, 2016.
- [10] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, “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*, 2017, pp. 15–26.
- [11] S. Liu, P.-Y. Chen, X. Chen, and M. Hong, “signsgd via zeroth-order oracle,” in *International Conference on Learning Representations*, 2018.
- [12] S. Liu, B. Kailkhura, P.-Y. Chen, P. Ting, S. Chang, and L. Amini, “Zeroth-order stochastic variance reduction for nonconvex optimization,” *Advances in Neural Information Processing Systems*, vol. 31, pp. 3727–3737, 2018.
- [13] D. Wierstra, T. Schaul, J. Peters, and J. Schmidhuber, “Natural evolution strategies,” in *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, 2008, pp. 3381–3387.
- [14] N. Hansen, “The cma evolution strategy: a comparing review,” in *Towards a new evolutionary computation*. Springer, 2006, pp. 75–102.
- [15] N. Maheswaranathan, L. Metz, G. Tucker, D. Choi, and J. Sohl-Dickstein, “Guided evolutionary strategies: Augmenting random search with surrogate gradients,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 4264–4273.
- [16] Y. Ruan, Y. Xiong, S. Reddi, S. Kumar, and C.-J. Hsieh, “Learning to learn by zeroth-order oracle,” *arXiv preprint arXiv:1910.09464*, 2019.
- [17] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [18] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, “Mastering the game of go with deep neural networks and tree search,” *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [19] S. Gu, E. Holly, T. Lillicrap, and S. Levine, “Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396.
- [20] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015.
- [21] H. Ye, Z. Huang, C. Fang, C. J. Li, and T. Zhang, “Hessian-aware zeroth-order optimization for black-box adversarial attack,” *arXiv preprint arXiv:1812.11377*, 2018.
- [22] O. B. Sezer and A. M. Ozbayoglu, “Algorithmic financial trading with deep convolutional neural networks: Time series to image conversion approach,” *Applied Soft Computing*, vol. 70, pp. 525–538, 2018.
- [23] S. Santurkar, D. Tsipras, A. Ilyas, and A. Madry, “How does batch normalization help optimization?” *arXiv preprint arXiv:1805.11604*, 2018.
- [24] S. Ghadimi and G. Lan, “Stochastic first-and zeroth-order methods for nonconvex stochastic programming,” *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2341–2368, 2013.
- [25] J.-K. Wang, X. Li, and P. Li, “Zeroth order optimization by a mixture of evolution strategies,” 2019.