Towards Efficient Neural Zeroth-Order Optimization Algorithms

Akhil Nadigatla

27 August 2021

Abstract

We consider the problem of zeroth-order optimization (ZOO), where our goal is to minimize an unknown function $f : \mathcal{X} \to \mathbb{R}$ given noisy zeroth-order access to f; that is, f is accessible only through noisy point evaluations. ZOO has a wide variety of applications in computer science, machine learning, and various engineering disciplines. For example, in machine learning, ZOO techniques are often used for hyperparameter tuning, where we need to tune several knobs of a statistical model and identify the configuration of hyperparameters which leads to the best model. Owing to its importance, ZOO has received a lot of attention from various research communities and, over the years, numerous techniques have been proposed for this problem. However, existing ZOO techniques have one or more of the following drawbacks: (a) they make restrictive structural assumptions on f, such as linearity or convexity, which rarely hold in practice, (b) they are computationally expensive and do not scale well to high-dimensional problems, and (c) they make too many queries to the zeroth-order oracle, which is prohibitive in applications such as hyperparameter tuning, where each function evaluation is expensive. Designing computationally efficient ZOO techniques that scale well to high-dimensional problems and work for general non-convex functions is very much an open problem.

In this work, we aim to make progress on this problem by designing efficient ZOO techniques for the setting where f is parameterized by a neural network. There are two main reasons for working with neural networks: (a) they are non-parametric function classes which can approximate any function to arbitrary precision. Moreover, they seem to be able to approximate functions occurring in many real world applications very well, and (b) we have good frameworks such as PyTorch and Tensorflow that we can utilize to efficiently implement our neural network-based ZOO techniques.

1 Problem

Consider the problem of minimizing a (potentially non-convex) function f over a (potentially non-convex) set $\mathcal{X} \subseteq \mathbb{R}^d$

$$f^* = \min_{x \in \mathcal{X}} f(x).$$

Suppose that we only have access to the function via a noisy zeroth-order/black-box oracle, which outputs a noisy estimate of the function when queried at any $x \in \mathcal{X}$

$$y = f(x) + \xi,$$

where ξ is a mean-zero random variable. The goal in zeroth order optimization (ZOO) is to find an approximate minimizer of f while making as few queries to the oracle as possible. Any ZOO algorithm makes a sequence of queries $x_1, x_2, \ldots, x_T$, and outputs some point $\hat{x}_T$ as a minimizer of f.

Such ZOO problems naturally arise in a number of fields. For example, in machine learning, ZOO techniques are often used for hyperparameter tuning, where we need to tune several knobs of a statistical model and identify the configuration of hyperparameters which leads to the best model [Snoek et al., 2012]. ZOO problems also arise in robust machine learning, where an adversary tries to make imperceptible changes/perturbations to the inputs of a neural network (NN) with the goal of making the network misclassify its inputs [Bhagoji et al., 2018, Liu et al., 2020]. Defending against such adversarial attacks often requires ZOO techniques which can efficiently identify the worst possible perturbation for any given input. These worst perturbations are subsequently used to train a robust NN. Another application of ZOO is in engineering design where these techniques help expedite the search for promising designs [Forrester et al., 2008].

2 Research Plan

2.1 Background Reading

Over the years, numerous techniques have been proposed for zeroth-order optimization. Below, we present some of the works that are relevant to this project and those that we intend to read. These works can be broadly classified into three categories. One category makes structural assumptions on the loss function f and while the other does not. The final category, which falls in between these two categories, assumes that f can be parameterized as a neural network.

1. **Structural Assumptions.** Works which assume f is a linear function of x: [Filippi et al., 2010, Kveton et al., 2020]. Works which assume f is a convex function of x: [Agarwal et al., 2011, Belloni et al., 2015]. Works which assume f is strongly convex: [Shamir, 2013]. Works which assume f is convex and smooth: [Bach and Perchet, 2016, Balasubramanian and Ghadimi, 2021].

2. **No Structural Assumptions.** Works that fall in this category do not make any assumptions (or make minimal assumptions such as continuity) on f. Some of these include: [Srinivas et al., 2012, Kleinberg et al., 2008, Bubeck et al., 2009].

3. **Neural Bandit Algorithms.** A recent line of work assumes that f is a neural network [Riquelme et al., 2018, Zhou et al., 2020, Zhang et al., 2020]. While this might look like a structural assumption, it actually is not because neural networks have the power to arbitrarily approximate any continuous function.

We believe reading these papers can help us develop a deeper understanding of the challenges involved in ZOO and help us develop new ZOO techniques for neural networks.

2.2 Intended Contributions

Existing ZOO techniques have one or more of the following downsides: (1) they make restrictive structural assumptions about the objective function f such as linearity or convexity, which often don't hold in practice, (2) they are computational expensive and do not scale well to high-dimensional problems, and (2) they make too many queries to the zeroth-order oracle, which can be prohibitive in applications where each function evaluation is expensive.

Therefore, we hope to further the discussion in this space by designing efficient ZOO techniques for the case where f is parametrized by NNs. Concretely, our aim is to arrive at ZOO techniques which (1) make no assumptions about the underlying objective function f, (2) require fewer queries to the zeroth-order oracle (at least by a constant factor), and (3) scale to high-dimensional problems and spare little computation time between queries.

2.3 Expected Results

The performance of any ZOO technique is typically measured using one of the following optimality criteria

$$f(\hat{x}_T) - f^* \quad \text{(Simple Regret)}$$
$$\frac{1}{T}\sum_{t=1}^{T}(f(x_t) - f^*) \quad \text{(Cumulative Regret)}.$$

Our goal is to develop practical ZOO techniques for neural networks which achieve $O(\mathrm{poly}(d)T^{-1/2})$ simple or cumulative regret after T calls to the zeroth-order oracle.

2.4 Timeline

Time	Event(s)
September 1 - 30	Perform background reading.
October 1 - 31	Propose potential algorithms. Setup code-base for running experiments.
November 1 - 30	Experiment with proposed algorithms and understand their drawbacks (and iterate). If there are no drawbacks, show that proposed algorithms provably work (i.e. derive regret bounds and/or rates of convergence).
December 1 - 14	Prepare for mid-thesis check-in presentation.
January 1 - February 20	By now, we should hopefully have a candidate algorithm which is computationally efficient and outputs a good minimizer. The main focus now will be to theoretically understand the algorithm. In parallel, perform thorough empirical evaluation.
February 21	Start writing thesis.
February 21 - April 21	Finish the theoretical parts and aim to submit the work to ML conferences such as NeurIPS.
April 21 - May 12	Prepare for Meeting of the Minds poster presentation. Prepare for final thesis presentation.

3 Research Advisors

My research advisor for this thesis project will be Professor Pradeep Ravikumar, an associate professor who also leads the Foundations of Statistical Machine Learning group in the Machine Learning Department.

I will also be working closely with Arun Sai Suggala, one of Prof. Pradeep's former PhD students who is now a research scientist at Google.

References

A. Agarwal, D. P. Foster, D. J. Hsu, S. M. Kakade, and A. Rakhlin. Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems*, pages 1035–1043, 2011.

F. Bach and V. Perchet. Highly-smooth zero-th order online optimization. In *Conference on Learning Theory*, pages 257–283. PMLR, 2016.

K. Balasubramanian and S. Ghadimi. Zeroth-order nonconvex stochastic optimization: Handling constraints, high dimensionality, and saddle points. *Foundations of Computational Mathematics*, pages 1–42, 2021.

A. Belloni, T. Liang, H. Narayanan, and A. Rakhlin. Escaping the local minima via simulated annealing: Optimization of approximately convex functions, 2015.

A. N. Bhagoji, W. He, B. Li, and D. Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–169, 2018.

S. Bubeck, G. Stoltz, C. Szepesvári, and R. Munos. Online optimization in x-armed bandits. In *Advances in Neural Information Processing Systems*, pages 201–208, 2009.

S. Filippi, O. Cappe, A. Garivier, and C. Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, pages 586–594, 2010.

A. Forrester, A. Sobester, and A. Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 681–690, 2008.

B. Kveton, M. Zaheer, C. Szepesvari, L. Li, M. Ghavamzadeh, and C. Boutilier. Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076, 2020.

S. Liu, P.-Y. Chen, B. Kailkhura, G. Zhang, A. Hero, and P. K. Varshney. A primer on zeroth-order optimization in signal processing and machine learning, 2020.

C. Riquelme, G. Tucker, and J. Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. *arXiv preprint arXiv:1802.09127*, 2018.

O. Shamir. On the complexity of bandit and derivative-free stochastic convex optimization. In *Conference on Learning Theory*, pages 3–24. PMLR, 2013.

J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012.

N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012. ISSN 1557-9654. doi: 10.1109/tit.2011.2182033. URL http://dx.doi.org/10.1109/TIT.2011.2182033.

W. Zhang, D. Zhou, L. Li, and Q. Gu. Neural thompson sampling. In *International Conference on Learning Representations*, 2020.

D. Zhou, L. Li, and Q. Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.