THE UNIVERSITY OF CHICAGO | DEPARTMENT OF STATISTICS

Master's Thesis Presentation

Kairun Zhang
Department of Statistics
The University of Chicago

"Learning to Optimize Zeroth-Order Perturbations for Fine-Tuning Large Language Models"

August 1, 2025 at 10:00 AM
Jones 304, 5747 S. Ellis Avenue

Abstract

Zeroth-order optimizers have recently emerged as a practical approach for fine-tuning large language models (LLMs), significantly reducing GPU memory consumption compared to traditional first-order methods. However, existing zeroth-order methods rely on hand-crafted and static sampling strategies that are not adaptable to the model-specific structure of foundation models. In this work, we propose ZO Fine-tuner, a learning-based zeroth-order optimizer that automatically learns efficient perturbation strategies through a compact and memory-efficient design. A key insight behind our approach is that only a small number of foundation models are widely adopted in practice—thus, optimizing the optimizer once for a given LLM and reusing it across diverse downstream tasks is both feasible and highly impactful. ZO Fine-tuner is designed to scale learning to learn (L2L) to the foundation model era by supporting one-time training per LLM with minimal overhead. Experiments on 4 LLMs and 7 datasets show that ZO Fine-tuner outperforms prior zeroth-order baselines in 82.1% of the task-model combinations compared to MeZO, thereby demonstrating strong performance and scalability for efficient LLM fine-tuning.
