


Bridging Theory and Practice: One-step Full Gradient Can Suffice for Low-rank Fine-tuning in LLMs



Invited Speaker

Fanghui Liu

University of Warwick

Date: September 17, 2025 (Wednesday)

Time: 17:00 - 18:00 (Hong Kong Time)

Zoom Meeting: 923 9677 5958

Biography

Fanghui Liu is currently an assistant professor at University of Warwick, UK, a member of Centre for Discrete Mathematics and its Applications (DIMAP). His research focuses on the foundations of modern machine learning and efficient algorithm design. He was a recipient of AAAI'24 New Faculty Award and Rising Star in AI (KAUST 2023). He has co-founded the fine-tuning workshop at NeurIPS'24, and presented three tutorials at ISIT'24, CVPR'23, and ICASSP'23, respectively. He has served as an area chair of NeurIPS, ICLR and AISTATS, etc.

Abstract

In this talk, I will discuss how to improve the performance of Low-Rank Adaption (LoRA) for fine-tuning in large language models (LLMs) guided by our theory. Our theoretical results show that LoRA will align to the certain singular subspace of one-step gradient of full fine-tuning. Hence, the subspace alignment and generalization guarantees can be directly achieved by a well-designed spectral initialization strategy for both linear and nonlinear models. Our analysis leads to the LoRA-One algorithm, a theoretically grounded algorithm that achieves significant empirical improvement over vanilla LoRA and its variants on several benchmarks by fine-tuning Llama 2. Additionally, our theory also demonstrates that preconditioners improve convergence in the high-rank case, has independent interest for understanding matrix sensing and deep learning theory. Talk is based on <https://arxiv.org/abs/2502.01235> (ICML'25 oral) with Yuanhe Zhang and Yudong Chen.