SMART Fine-tuning Factor Augmented Neural Lasso

Jianqing Fan

Princeton University

ABSTRACT

Fine-tuning is a widely used strategy for adapting pre-trained models to new tasks, yet its methodology and theoretical properties in high-dimensional nonparametric settings with variable selection have not yet been developed. We propose a source-model-augmented residual tuning (SMART) framework, which incorporates the pre-trained source model into the target learner and estimates only the residual target-specific component. The approach is widely applicable, from parametric and sparse models to neural networks and blackbox machine learning models. We focus on the development of fine-tuning factor-augmented neural Lasso, resulting in SMART-FAN-Lasso. This transfer-learning framework for high-dimensional nonparametric regression with variable selection simultaneously handles covariate and posterior shifts. We use a low-rank factor structure to manage high-dimensional dependent covariates and a residual tuning decomposition in which the target function is expressed as a function of source model and other target-specific variables, thereby reducing the effective complexity of the target task. We derive minimax-optimal excess risk bounds for SMART-FAN-Lasso, characterizing the precise conditions, in terms of relative sample sizes and function complexities, under which fine-tuning yields statistical acceleration over single-task learning. Extensive numerical experiments across diverse covariate- and posterior-shift scenarios demonstrate that SMART-FAN-Lasso consistently outperforms standard baselines and achieves near-oracle performance even under severe target sample size constraints, empirically validating the derived rates.

(Joint work with Jinhang Chai, Cheng Gao, and Qishuo Yin)