Fast Evolutionary Neural Architecture Search with Predictor

A/Prof. Andy Song

RMIT, Australia

Abstract:

This talk introduces our work recently published in GECCO and TEVC. The key contribution is using predictors to speed up Neural architecture search (NAS) as NAS often requires a high computational cost to evaluate candidate networks from the search space. That cost could be mitigated by performance prediction so there is no need to evaluate every candidate network. However, training predictors often require a large number of evaluated architectures which may be difficult to obtain. We address this challenge by proposing a novel evolutionary-based NAS strategy, predictor-assisted evolutionary NAS (PRE-NAS) which can perform well even with an extremely small number of evaluated architectures. PRE-NAS leverages new evolutionary search strategies and integrates high-fidelity weight inheritance over generations. Unlike popular one-shot strategies, which may suffer from bias in the evaluation due to weight sharing, offspring candidates in PRE-NAS are topologically homogeneous. This circumvents bias and leads to more accurate predictions. Extensive experiments on the NAS-Bench-201 and DARTS search spaces show that PRE-NAS can outperform state-of-the-art NAS methods. PRE-NAS can achieve 2.40% and 24% test error rates on CIFAR-10 and ImageNet respectively. The search cost is only 0.6 days with a single GPU, surpassing most SOTA. In our further work, we developed a novel training-free metrics leveraging active learning, GCN, and a few frontier techniques for the evolutionary NAS. To achieve similar test errors on CIFAR-10 and ImageNet as PRE-NAS, we only need 29 minutes and 37 minutes on a single GPU.