

Reinforcement Learning for Large-Scale Communications and Computer Infrastructures

Kin K. Leung

EEE and Computing Departments
Imperial College, London

Web: www.commsp.ee.ic.ac.uk/~kkleung/

Abstract

Among many machine-learning techniques, reinforcement learning (RL) has been widely applied to many domains. Deep RL techniques have been used to solve routing, service-placement and power allocation problems on the software defined networks (SDN) as well as in the software defined coalitions (SDC) developed in the DAIS ITA Program <https://dais-ita.org/pub>.

This seminar begins with a brief introduction to RL for communications and computer systems. For illustration purposes, the speaker presents use of RL to train and obtain a smart policy for domain controllers to synchronize with each other in order to maximize performance gains in distributed SDN. Evaluation results show that the RL policy significantly outperforms other algorithms for inter-domain routing tasks.

As reflected in the above work, a challenging issue for deep RL is the huge state and action spaces, which increase the model complexity and training time beyond practical feasibility. Various techniques have been proposed to address the issue. In particular, we present here a method to decouple actions from the state space for the value-function learning process and a relatively simple transition model is learned to determine the action that causes the associated state transition. Numerical results show that the state-action separable RL can greatly reduce training time without noticeable performance degradation.

The speaker will conclude by highlighting the open issues for use of RL for control of large-scaled communications and computer systems.