Machine Learning Techniques for Solving Combinatorial Optimisation Problems

Deepak Ajwani[a], Paula Carroll[a], James Fitzpatrick[a], Saurabh Ray[b], Dena Tayebi[a]

[a]University College Dublin, Ireland
[b]New York University, Abu Dhabi, United Arab Emirates

In recent years, machine learning techniques are being increasingly used for solving combinatorial optimisation problems. This often requires a deep integration between techniques from optimisation literature, algorithm engineering and machine learning. For instance, while the optimisation and algorithmic literature guides the feature engineering in learning models, the learning models can guide crucial design steps in exact MILP solvers as well as heuristics.

Specifically, I would like to talk about the research done in my group on a range of fundamental combinatorial optimisation problems such as set cover, k-median, facility location, vehicle routing problems, Max Cut, Max Clique, Steiner tree etc. Firstly, I will describe a simple supervised learning framework called learning-to-prune that can be used to reduce the size of the problem instances. This enables the computation of high quality solutions to larger and harder instances of combinatorial optimisation problems. Later, I will present a heuristic based on reinforcement learning that provides close-to-optimal solutions on many NP-hard vehicle routing problem variants.

If time permits, I can also point out some opportunities in using machine learning for discovering combinatorial structures.