Abstract

Solving Combinatorial Optimization(CO) problems over graphs finds applications in many areas such as planning, scheduling, manufacturing, resource allocation and transportation. Many combinatorial optimization problems are NP-hard making it challenging to develop efficient algorithms to solve them in practical settings. The traditional algorithmic development paradigm to solve CO problems heavily relies on leveraging domain-specific knowledge or past experiences with similar problems. Such a design choice often requires substantial time and effort to craft solutions tailored to specific problem contexts.

Our key insight is that real-world instances often involve solving diverse graph-based CO problems repeatedly. This motivates leveraging machine learning to discover problem-specific heuristics from data. By harnessing deep neural networks' capabilities, ML algorithms can learn problem-specific patterns, leading to potentially superior solutions, especially for complex problems beyond human design intuition.

In this work, we focus on the problem of learning to solve graph combinatorial optimization problems. We investigate learning to solve diverse NP-hard graph combinatorial problems in various areas such as routing, coverage, allocation, partitioning, hardware, and material science. We address several challenges inherent in this context, including the ability to generalize to graph sizes, distributions, nodes, and parameters specific to the combinatorial problem at hand that haven't been encountered before. We also investigate and address the challenges associated with learning to solve CO problems under the constraint of data scarcity. Furthermore, we investigate the ability for continuous adaptation within a lifelong learning framework where the training data emerges over time in a streaming fashion. Through multiple experiments, we showcase the effectiveness of the proposed methods, achieving high-quality results for different problems and scenarios.