
Enhancing Local Search Through Machine Learning: a Case Study on the Vehicle Routing Problem

Daniele Vigo^{*1}, Luca Accorsi¹, Michele Lombardi², and Michela Milano²

¹DEI, University of Bologna – Italy

²DISI, University of Bologna – Italy

Abstract

Local search is one of most popular optimization techniques used to effectively solve hard combinatorial optimization problems (COPs) and in particular large scale ones, for which exact solution approaches are prohibitive. Local search is indeed the main engine of all modern metaheuristics which perform hundreds of thousands of iterations exploring the neighborhoods of given solutions.

Practical algorithms typically exploit a relatively small set of simple local search operators and use them to optimize a solution until a local optimum is reached.

A common example of such approaches is the Variable Neighborhood Descent (VND), in which the neighborhood defined by each operator is completely examined before moving to the next operator.

Up to now, little attention has been given to the orderings in which operators are examined; common orderings usually consists in prefixed static or dynamic randomly ordered operators examination.

This work we study whether in a VND setting, using Machine Learning techniques to identify the most promising operator (or knowing that none of them will be able to improve a given solution) enables computational savings that could be used to perform a more fruitful search. In particular, as a case study, we focused on the (capacitated) vehicle routing problem, and we train an Artificial Neural Network for ranking operators at search time.

Keywords: Vehicle Routing, Machine Learning, Heuristics

*Speaker