
Mining frequent patterns to drive the exploration of high-order neighborhoods

Florian Arnold^{*1}, Thibaut Vidal², Italo Gomes Santana², and Kenneth Sörensen¹

¹University of Antwerp – Belgium

²Pontifical Catholic University of Rio de Janeiro (PUC-Rio) – 1PUC-Rio – Pontificia Universidade Católica do Rio de Janeiro Caixa Postal 38097 Rio de Janeiro, RJ, Brazil, Brazil

Abstract

Pattern mining is a well-established technique to discover frequently-occurring substructures in datasets. These substructures can be used to infer general knowledge about the underlying data, e.g., "if a customer buys milk, he is likely to also buy bread and coffee". This idea can be easily translated into the routing domain to obtain statements such as "a visit to customer A is likely to be followed by visits to customers B, C and D in good vehicle routing solutions".

In this work, we use pattern mining to make a controlled exploration of high-order neighborhoods in a local search (LS). We maintain a limited number of most-frequent patterns in elite solutions. Then, during the LS, each pattern serves to define one move in which 1) incompatible edges are disconnected, 2) the edges defined by the pattern are reconnected, and 3) the remaining route fragments are optimally reconnected. Each such move is accepted only in case of improvement. This "pattern injection" local search (PILS) finds useful high-order local search moves which would not be detected with traditional operators. Our first experiments indicate that this technique can thereby complement and improve existing state-of-the-art heuristics with a limited implementation effort, and can potentially be applied to a wide range of combinatorial optimization problems.

Keywords: Vehicle routing, pattern, mining, local search, heuristics

*Speaker