Neurocomputing 205 (2016) 393–406

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Meta-learning to select the best meta-heuristic for the Traveling Salesman Problem: A comparison of meta-features

Jorge Kanda [a,*], Andre de Carvalho [b], Eduardo Hruschka [b], Carlos Soares [c], Pavel Brazdil [d]

[a] Instituto de Ciencias Exatas e Tecnologia, Universidade Federal do Amazonas Rua Nossa Senhora do Rosario, 3863, Tiradentes, 69103-128 Itacoatiara, Amazonas, Brazil
[b] Instituto de Ciencias Matematicas e de Computacao, Universidade de Sao Paulo Av. Trabalhador Sancarlense, 400, 13566-590 Sao Carlos, Sao Paulo, Brazil
[c] INESC TEC/Faculdade de Engenharia - Universidade do Porto, Porto, Portugal
[d] LIAAD-INESC Tec/Faculdade de Economia, Universidade do Porto, Porto, Portugal

ARTICLE INFO

Article history:
Received 15 October 2015
Received in revised form
16 February 2016
Accepted 7 April 2016
Communicated by K. Chan
Available online 12 May 2016

Keywords:
Meta-learning
Meta-features
Label ranking
Meta-heuristics
Traveling Salesman Problem

ABSTRACT

The Traveling Salesman Problem (TSP) is one of the most studied optimization problems. Various meta-heuristics (MHs) have been proposed and investigated on many instances of this problem. It is widely accepted that the best MH varies for different instances. Ideally, one should be able to recommend the best MHs for a new TSP instance without having to execute them. However, this is a very difficult task. We address this task by using a meta-learning approach based on label ranking algorithms. These algorithms build a mapping that relates the characteristics of those instances (i.e., the meta-features) with the relative performance (i.e., the ranking) of MHs, based on (meta-)data extracted from TSP instances that have been already solved by those MHs. The success of this approach depends on the quality of the meta-features that describe the instances. In this work, we investigate four different sets of meta-features based on different measurements of the properties of TSP instances: edge and vertex measures, complex network measures, properties from the MHs, and subsampling landmarkers properties. The models are investigated in four different TSP scenarios presenting symmetry and connection strength variations. The experimental results indicate that meta-learning models can accurately predict rankings of MHs for different TSP scenarios. Good solutions for the investigated TSP instances can be obtained from the prediction of rankings of MHs, regardless of the learning algorithm used at the meta-level. The experimental results also show that the definition of the set of meta-features has an important impact on the quality of the solutions obtained.

1. Introduction

The Traveling Salesman Problem (TSP) is one of the most intensively studied problems in combinatorial optimization and theoretical computer science. TSP has been used to represent applications from different domains, such as machine scheduling, DNA sequencing, transportation, and microchip manufacturing [1]. A TSP can be informally described as: given a set of cities and their respective pairwise distances, find the tour with the lowest possible cost that starts in one of the cities, visits all the other cities only once, and ends at the initial city [2]. The number of cities and how they are connected define different instances of the TSP. It is difficult to find a global optimal solution for a given TSP instance, since TSP belongs to the class of problems known as NP-hard [3]. Therefore, exhaustive search methods are not applicable to find the best solution for large TSP instances, due to their high computational cost. For example, there are approximately 1.22×10^{17} feasible solutions for a TSP with 20 cities. Thus, an exhaustive search to find a global optimum solution would take a long time.

The high computational cost involved in solving TSP problems can be significantly reduced by the use of Meta-Heuristics (MHs), which are often able to provide near-optimal solutions in reasonable time. MHs are high-level search strategies that guide the search to more promising regions of the solution space and try to escape from local optimal solutions [4]. Several MHs have been successfully used for TSP instances, including Tabu Search [5], Greedy Randomized Adaptive Search Procedure [6], Simulated Annealing [7], Genetic Algorithms [8], and Ant Colony Optimization [9].

In spite of their general success, different MHs have different biases, which make each of them more suitable for a particular class of instances [10]. Therefore, given a set of MHs and a new TSP instance to be solved, the best choice depends on the

* Corresponding author. Tel.: +55 92 35213519; fax: +55 92 35213603.
E-mail addresses: jkanda@ufam.edu.br (J. Kanda), andre@icmc.usp.br (A.d. Carvalho), erh@icmc.usp.br (E. Hruschka), csoares@fe.up.pt (C. Soares), pbrazdil@inescporto.pt (P. Brazdil).

http://dx.doi.org/10.1016/j.neucom.2016.04.027
0925-2312/

characteristics of the TSP instance. About a decade ago, the concept of hyper-heuristic was introduced as an alternative approach to the choice of MHs for solving optimization problems [11]. The key idea was to apply different MHs to different parts of the solution process according to the strength of each MH. Nevertheless, it is not possible to ensure the existence of interdependence between different parts of an optimization problem like the TSP. Furthermore, for each part of the problem, the most suitable MH must still be selected.

Mapping the characteristics of problem instances with the relative performance of a set of algorithms that can deal with these instances is a task that can be performed with meta-learning [12]. Meta-learning allows the selection of the most promising techniques using an inductive learning process. Recently, a meta-learning approach addressed the problem of recommending MHs for new TSP instances as a multilabel classification task [13]. However, when multiple MHs are recommended, the user is left without any guidance concerning which one should be initially tried. We address this problem by using a label ranking approach [14], which predicts a ranking of MH candidates, according to their expected performance for a new TSP instance. A ranking is more useful for the user, since it recommends the execution of the MHs in the predicted order until a satisfactory result is obtained.

Previous approaches have investigated only a single TSP scenario, usually containing only *symmetric and strongly connected* instances [15,13,16]. The TSP instance is symmetric if the cost of traveling from city i to the adjacent city j is equal to the cost of traveling from j to i; and it is asymmetric when the costs of traveling between cities i and j are different. The TSP instance is strongly connected when all cities are interconnected; and it is weakly connected if the TSP instance has at least one pair of cities that is not directly connected. In our study, we completed those earlier approaches by systematically investigating the four possible scenarios in terms of symmetry and connectedness: *symmetric and strongly connected*, *asymmetric and strongly connected*, *symmetric and weakly connected*, and *asymmetric and weakly connected*.

The main contributions of this work are:

- We show that meta-learning models can accurately predict rankings of MHs for different TSP scenarios.
- We investigate new meta-features for TSP to be applied to the meta-learning process.
- We offer evidence that good solutions for new TSP instances can be obtained from the prediction of rankings of MHs regardless of the learning algorithm used at the meta-level.

The remainder of this paper is organized as follows. A short description of the TSP is given in Section 2. The required background for meta-learning for algorithm selection and label ranking is provided in Section 3. Adaptations of some machine learning techniques to the label ranking problem are described in Section 4. Section 5 presents the sets of meta-features used to characterize the TSP instances. Practical application scenarios of interest are given in Section 6. Based on these scenarios, the experimental setting is detailed in Section 7, while the obtained results are reported and analyzed in Section 8. Finally, the main conclusions are presented in Section 9.

2. The Traveling Salesman Problem

Formally, the Traveling Salesman Problem (TSP) can be defined by means of a graph $G=(V,E)$, in which $V=\{v_1, v_2, \ldots, v_n\}$ is a set of vertices and $E=\{\langle v_i, v_j\rangle : v_i, v_j \in V\}$ is a set of edges. Each vertex $v_i \in V$ represents a city and each edge $\langle v_i, v_j\rangle \in E$ connects the vertices v_i and v_j. The cost associated with the edge $\langle v_i, v_j\rangle$ (i.e., the cost of traveling from city v_i to city v_j) is indicated by the value c_{ij}.

Given a TSP instance, the goal is to find the minimum value for Eq. (1):

$$\min \quad z = \sum_{i=1}^{n}\sum_{j=i}^{n} c_{ij}x_{ij} \tag{1}$$

subject to:

$$\sum_{i=1}^{n} x_{ij} = 1 \quad \forall j \in V \tag{2}$$

$$\sum_{j=1}^{n} x_{ij} = 1 \quad \forall i \in V \tag{3}$$

$$\sum_{i,j \in S} x_{ij} \le |S| - 1 \quad \forall S \subset V \tag{4}$$

$$x_{ij} \in \{0,1\} \quad \forall i \ne j \in V \tag{5}$$

If the tour traverses edge $\langle v_i, v_j\rangle$ then $x_{ij}=1$, and $x_{ij}=0$, otherwise.

The best solution for a TSP is given by the Hamiltonian cycle of minimum total cost. A cycle is Hamiltonian if all cities are visited only once and the route ends at the initial city [2]. Fig. 1 illustrates a TSP instance with six cities. The route given by the edges in bold indicates the optimal solution.

TSP instances can be characterized according to strength of their connections and to their symmetry. A TSP instance is strongly

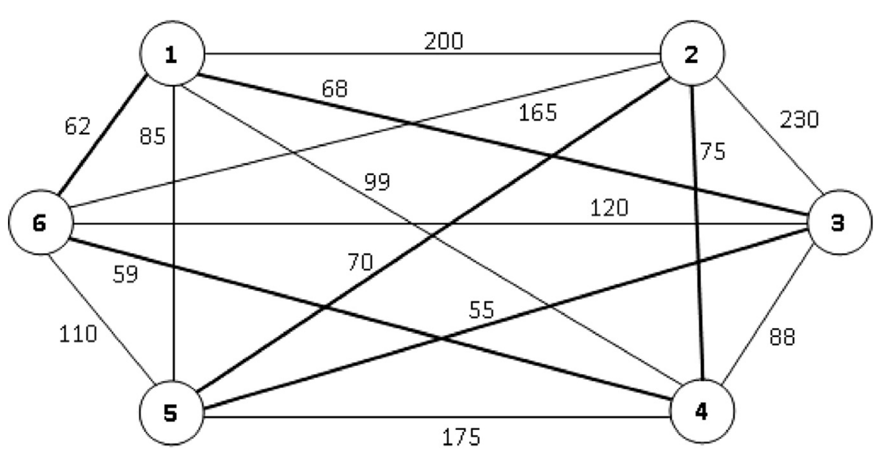

Fig. 1. A TSP instance and its optimal solution.