Information Sciences 372 (2016) 162–180

Contents lists available at ScienceDirect

Information Sciences

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

Investigating the correlation between indicators of predictive diagnostic optimisation and search result quality

I. Moser[a,*], Marius Gheorghita[a], Aldeida Aleti[b]

[a] Swinburne University, Melbourne, Australia
[b] Monash University, Melbourne, Australia

ARTICLE INFO

Article history:
Received 23 October 2015
Revised 5 August 2016
Accepted 8 August 2016
Available online 8 August 2016

Keywords:
Search space characterisation
Local search
Prediction
Combinatorial optimisation

ABSTRACT

Combinatorially complex problems are often optimised with heuristic solvers which generally provide acceptable results but no indication as to how the quality achieved compares to the best possible. In previous work we have introduced Predictive Diagnostic Optimisation (PDO), a heuristic based on local search that provides information about the search space structure through a set of indicators whilst searching for the optimal solution. PDO can collect useful information about the search process, such as the variation in the number of steps needed to locally optimise a random solution and the error between the expected and actual qualities of the local optimum, known as the prediction error. Given previous experimental results on the quadratic assignment problem, it appears that a high prediction error coincides with lower search quality and vice versa. This work confirms this assumption with the help of two additional problems but also shows that the reliability of the prediction error is challenged by structural properties that lead to a homogeneity of the optima basins. Conversely, a high variation in the number of steps that lead to the local optima increases the reliability of the prediction error as an indicator of search quality.

1. Introduction

Combinatorial optimisation problems often have too many potential solutions to be solved with an exact approach. It is customary to apply heuristic search methods such as Evolutionary Algorithms, Swarm-inspired methods or GRASP algorithms. A considerable body of research has been devoted to comparative studies of such algorithms when applied to particular problems. Such experimentation often lacks scientific rigour [3] and does not succeed in uncovering the underlying causes of different algorithm performances [20]. In the absence of knowledge about the interaction between problem and solver, it is impossible to extract generalisable results that are helpful in judging the potential of an algorithm for solving a new problem.

It is customary to cite Wolpert and Macready [60] in this context to underline the argument that different heuristics (and their comparisons) are necessary because algorithms that perform well on one type of problem perform badly on others: '[NFL theorems] ... demonstrate that if an algorithm performs well on a certain class of problems then it necessarily pays for that with degraded performance on the set of *all remaining* problems.' (Emphasis added.) Since the field of heuristic algorithms for combinatorial optimisation generally only considers one specific class of problems - multimodal problems with

* Corresponding author.
E-mail address: imoser@ict.swin.edu.au (I. Moser).

http://dx.doi.org/10.1016/j.ins.2016.08.021
0020-0255/

gradients a heuristic can follow - Wolpert and Macready's theorem cannot be assumed to have any bearing. Their study emphasises: 'In particular, if an algorithm performs better than random search on some class of problems then it must perform *worse than random search* (sic) on the remaining problems.' If this was applicable to the subset of problems researchers generally quote it for, we would find that, say, Genetic algorithms perform exceedingly well on bin packing problems while they are unusable with quadratic assignment problems. This is evidently not the case - generally, in experimental studies, all heuristics applied outperform random search.

Even so, in many cases even small improvements one heuristic achieves over another can be relevant to a practitioner. The practice of tuning an algorithm and comparing it with many existing algorithms is wasteful - given a new problem, the same procedure has to be repeated. Moreover, the subset of instances used for the comparison may not be representative of the actual instances in need of a solution [45]. Recognising these problems, some researchers study the properties of problems to establish features which render them particularly suitable for a search by a specific algorithm.

Initially, researchers examined properties of individual problems in an attempt to find the features that make them difficult to solve, such as the variation of distance matrices for TSP [43] or QAP [33]. More recently, authors have attempted to mine a number of properties for a diverse set of problem instances and subsequently extract the most expressive ones using a statistical technique [29,45,46]

In fitness landscape analysis, an active research area since the early 1990s, researchers interpret properties of the problem's fitness landscape which arises from the search operator used to improve solutions incrementally. In many cases, the search operator is applied randomly to establish the continuity of fitness development in the absence of a search bias, as does autocorrelation [1]. Other approaches measure landscape properties in connection with a given heuristic, such as the negative slope coefficient [52], which follows solutions while they are being optimised by an Evolutionary algorithm (EA). In all of these approaches, the principal goal is to determine which of a variety of algorithms are most likely to produce the best outcome.

In previous work [34], we have introduced Predictive Diagnostic Optimisation (PDO), an iterated local search method designed to estimate the difficulty of the search as the optimisation progresses, with the goal of providing information about the relative quality of the solutions achieved by the search. Although the PDO indicators have shown a good correlation with result quality on many QAP instances, the method has not been tested on other problems. The current work remedies this, providing an interpretation of a number of indicators collected during the search on QAP, the linear ordering problem (LOP) and the permutation flow-shop scheduling problem (PFSP), which are described in Section 4.1.

After a review of the major approaches in fitness landscape characterisation in Section 2, we provide a brief overview of the PDO approach in Section 3. Section 4 explains how the experiments were conducted on a representative instance set of the problems. The results section correlates the three PDO indicators with the optimal quality found, including three additional indicators which add detail to the characterisation results.

2. Fitness landscape characterisation

The concept of fitness landscape relates the topological description of the search space of a problem with the dynamics of search algorithms. Formally, a fitness landscape is represented by a triplet $\{S, N, F\}$ where [40,47,49]:

(i) S is the set of potential solutions for the given problem, also known as the *search space*.
(ii) N is neighbourhood relation, defined as a subset of $S \times S$. Two solutions are neighbours (i.e. $\langle s, s'\rangle \in N$) if it is possible to reach one solution by applying a search operator to the other.
(iii) F: $S \rightarrow \Re$ is the fitness or objective function that assigns a real value to each solution, which defines its relative quality.

As the neighbourhood of a solution depends on the operator employed, different operators give rise to different landscapes. Finding the best suited heuristic for a problem can also be seen as the problem of finding a heuristic able to create a landscape with the most helpful gradients to exploit, which arise from the changes in fitness incurred from each incremental change of the solution at hand.

2.1. Features of fitness landscapes

As contemporary heuristics attempt to exploit gradients towards optima, from the point of view of the search and using the terminology for minimisation, local optima form basins of attraction which heuristics are expected to find. Considering the search space S and a neighbourhood relation N, a *local optimum* is considered a solution $s \in S$ such that for all solutions in the neighbourhood $\forall \langle s, s'\rangle \in N$, $F(s) < F(s')$, considering a minimisation problem. A *global optimum* is a solution $s \in S$ such that for all solutions in the search space $\forall s' \in S$, $F(s) < F(s')$.

2.2. Modality

Unimodal fitness landscapes have a single optimum which can be reached from every solution to the problem simply by applying a local search algorithm. The algorithm makes improving changes to a given solution until it is optimal. Problem difficulty arises from multimodality, where the search landscape has interrupted gradients which end at numerous