Characterising Fitness Landscapes using Predictive Local Search

Marius Gheorghita
Swinburne University of Technology
Melbourne, Australia
mgheorghita@swin.edu.au

Irene Moser
Swinburne University of Technology
Melbourne, Australia
imoser@swin.edu.au

Aldeida Aleti
Monash University
Melbourne, Australia
aldeida.aleti@monash.edu

ABSTRACT

Search space characterisation is a field that strives to define properties of gradients with the general aim of finding the most suitable stochastic algorithms to solve the problems. Diagnostic Optimisation characterises the search landscape while the search progresses. In this work, we have improved Predictive Diagnostic Optimisation to reduce the cost of the local search by introducing a sampling procedure to explore the neighbourhood. The neigbhourhood is created by the swap operator and the sample size recorded during the search is shown to correlate with the known characteristics of the problems.

Categories and Subject Descriptors

D.2.8 [**Artificial Intelligence**]: Diagnostic Optimisation

1. INTRODUCTION

Local search is a successful class of approximation algorithms which has been shown to achieve near-optimal solutions to many difficult problems when coupled with a suitable global search [4]. In multimodal landscapes, where local search is applied, the solution obtained at the end is usually a local optimum. In recent work [3] it was shown that the quality of such local optima can be predicted with a certain accuracy and that this can help decide which of the initial solutions promise the largest gain from a local search. This work intends to reduce the complexity of exploring the neighbourhood in a way that is computationally more efficient than the initial Predictive Diagnostic Optimisation (PDO).

2. PREDICTIVE DIAGNOSTIC OPTIMISATION

The core technique of PDO is the prediction of the ultimate solution quality after a local search has been applied, solving the problem of choice of initial solution for the local search. The prediction is based on the fitness improvement ensuing from the first change made to the initial solution. This first step of the local search is used to project the final quality to be expected when no further improvement is possible. The ratio of the improvement achieved by the first step and the fitness improvement after the search stops forms a predictor. A predictor can then be matched to new initial solutions and their first moves. When, after the local optimisation, it is observed that none of the existing predictors was able to predict the ultimate locally optimal fitness, a new predictor is created based on this solution. As predictors are created dynamically whenever the existing predictors are unable to predict the quality of the optimum to a predefined margin of accuracy, the number of predictors created during the run can be used as an indicator of the homogeneity of the search space.

2.1 Steepest descent

The method used for exploring the neighbourhood created from the swap operator, steepest descent (SD), is an expensive technique that exhaustively explores the complete neighbourhood before making the move that leads to the best fitness improvement. It is an iterative improvement method in which, at each step, the current solution is improved by a predefined small modification. At each step, all possible modifications of the current solution are explored and the change with the biggest fitness improvement is made permanent. It ends when no further improving moves are possible. For most types of problems of realistic sizes, the resulting neighbourhood is prohibitively large. In this work, we decided to replace SD by a sampling procedure, computationally inexpensive, which explores the neighbourhood looking for an approximate best solution before making a next move.

2.2 Sampling

The primary goal is to reduce the cost of the local search and at the same time to preserve approximately the same prediction accuracy and fitness improvement as the exhaustive SD approach. Compared to SD, when we do sampling the neigbourhood size is reduced to the sample size, therefore the exploration is limited to the sample. Instead of searching the neighbourhood exhaustively, the sampling procedure selects the best solution found in a representative sample. A representative sample is one that finds an approximate best fitness improvement. It is expected that the sample size is influenced by the diversity of the gradients and to have more samples in a rugged landscape than in a homogeneous one since more exploration is needed to find the approximate best fitness improvement. Initially, $m = 30$ local search moves are sampled beginning from the same solution. Since the sampling size is determined using the average of the best fitnesses, lower values for the minimum size would increase the chances for the sample to stop pre-

GECCO'13 Companion, July 6–10, 2013, Amsterdam, The Netherlands.
ACM 978-1-4503-1964-5/13/07.

maturely. The sample size is recorded during the predictors creation and when they are used to make predictions. For testing whether the sample size correlates indeed with the ruggedness of the landscape, a second batch of experiments with $m = 50$ was performed as well.

For determining the approximate best improvement from the neighbourhood by sampling, a statistical significance test is proposed, where a dynamic stopping criterion based on accuracy monitoring determines when the sample reaches the confidence level for the neighbourhood. After the initial m samples, the average of the best fitnesses a_1 and the mean square of those best fitnesses, a_2 is computed for each new sample. The sampling increments k, the current sample size, until the predefined significance level is reached, computed according to Eq. 1.

$$wr = \frac{2Z}{\sqrt{k}} \frac{\sqrt{a_2 - a_1^2}}{a_1} \tag{1}$$

where Z-values for the 95%, 98% and 99% confidence intervals are 1.96, 2.33 and 2.58.

The *sample size index* is calculated as the neighbourhood sample size attained when the statistical significance level has been reached divided by the initial sample size m. The goal of sampling is to reduce the overhead of investigating all possible moves that can be made from one solution while losing as little as possible of the accuracy provided by SD. Intuitively, a higher significance level would allow the sample size to develop more, giving more chances to discover better fitness, but a tradeoff between sampling overhead and accuracy is desirable. For example, on Chr25a problem, 95% had the sample index 1.001 for the best fitness normalised to the global optimum at 0.66, while 98% had the index 1.24 for a normalised best fitness at 0.68. 99% had the index 2.42 for a normalised best fitness of 0.67, therefore we have setup for 98%.

2.3 Results

To investigate the effects of replacing SD with the sampling technique, a set of twelve instances of Quadratic Assignment Problem (QAP) from the QAPLIB collection [1], six instances of Linear Ordering Problem (LOP) available in the LOLIB benchmark library [2] and six instances of Flow Shop Scheduling Problem (FSSP) generated according to Taillard [5] were chosen for experimentation.

The values confirm that for instances like Nug20, Nug30, Tai100a, Tai100b with a homogeneous landscape, the sample size index stays at minimal values, 1 , which means that the minimum size sample is representative enough for exploring the neighbourhood. The sample size index shown in Table 1 correlates with the ruggedness of the landscape created, for hard instances like Chr20a, Chr25a and Ste36b showing that a larger sample is needed. If the landscape is rugged, the size of the sample is larger due to the diversity of gradients, and a larger sample promises a better solution to be discovered in the neighbourhood exploration. On FSSP instances, on $m = 30$ there is a slight difference in index for the less complex problems. In most cases, the LOP instances present a small index, more uniform for $m = 50$, except for hard problems where we obtain a bigger sample size index, e.g econ36 containing real-world data. Observing the overall distribution of the sample index values over the two studied cases, we conclude that starting sampling with bigger sizes will attenuate the differences in the index leaving only the very rugged landscapes with a significant index. From the diagnostic angle, a sample that can capture as much as possible the differences between the problems is preferable. To use this sample index as a metric for ruggedness, it is necessary to confirm these results with more tests which apply a diversity of neighbourhoods.

Table 1: Sample size index.

Problem		m=30	m=50
QAP	Chr20a	2.093	1.54
	Chr25a	1.528	1.24
	Kra30a	1	1
	Kra32	1	1
	Nug20	1.001	1
	Nug30	1	1
	Ste36a	1.064	1.02
	Ste36b	1.277	1.12
	Tai20a	1	1
	Tai30b	1.119	1.04
	Tai100a	1	1
	Tai100b	1	1
LOP	econ36	1.243	1.11
	p40.01	1.001	1
	be75eec	1.056	1.02
	sgb75.01	1.01	1
	t1d100.01	1	1
	be75eec150	1	1
FSSP	Tai20x10	1.002	1
	Tai20x20	1.002	1
	Tai50x10	1	1
	Tai50x20	1	1
	Tai100x10	1	1
	Tai100x20	1	1

3. CONCLUSIONS

In this paper, we have improved PDO, an optimisation approach which is based on predictive local search. The main improvement of the algorithm comprises the introduction of a sampling procedure that reduces the cost of the local search during the exploration of the neighbourhood. The information obtained during the search correlates with the difficulty of the problem and the values obtained for the metric vary significantly depending on the landscapes. These landscapes were investigated in the literature and our previous work and the observations coincide with the existing knowledge about the problem instances. An interesting direction for further research is to test different neighbourhoods for the consistency of the sample size index as a characterisation metric especially for unknown problems.

4. REFERENCES

[1] R. E. Burkard, S. Karisch, and F. Rendl. Qaplib-a quadratic assignment problem library. *European Journal of Operational Research*, 55(1):115–119, 1991.

[2] R. Marti, G. Reinelt, and A. Duarte. Optsicom project-library for linear ordering problem, 2009.

[3] I. Moser and M. Gheorghita. Combining search space diagnostics and optimisation. In *Congress on Computational Intelligence, Proceedings*, pages 897–904. IEEE, 2012.

[4] T. Stützle. *Local search algorithms for combinatorial problems - analysis, improvements, and new applications*, volume 220 of *DISKI*. Infix, 1999.

[5] E. D. Taillard. Scheduling instances, 2005.