Recent Advances in Problem Understanding: Changes in the Landscape a Year On

Kent M[c]Clymont
Computer Science
University of Exeter
Exeter, UK
k.mcclymont@exeter.ac.uk

ABSTRACT

This paper provides an updated survey of new literature in, and related to, the field of problem understanding which has been published or made available since January 2012. The bibliographic information from the survey is available online at `http://bit.ly/ZWoY3X`. The survey covers work on the topics of: Benchmark Problems; Problem Decomposition & Multiobjectivisation; Landscape Analysis; Problem Difficulty; and Algorithm Selection & Performance Prediction. In addition, special attention is drawn to three recently published and excellent topic specific surveys. A side note is also made regarding the parallels between problem understanding, and specifically landscape analysis„ and the work of fitness landscape analysis in theoretical, conventional and evolutionary biology.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*

Keywords

Problem Understanding, Problem Analysis, Algorithm Analysis, Optimisation

1. INTRODUCTION

In this paper, I present a brief review of work that has recently been published or surfaced since our initial survey on work in the field of and relating to Problem Understanding: "The lay of the land" (26). That original survey was presented at the 1st Understanding Problems Workshop (GECCO-UP) (25) and GECCO 2013 and was intended as both a quick overview of work in the area and also as a position paper presenting one potential approach for researchers in the field to work together to more effectively explore this currently nebulous body of research. This paper continues the survey as part of that workshop series and aims to give an annual update on recent developments and publications in the area, highlighting recent advances and where appropriate flagging emerging directions of research.

In addition to those given in this paper, the references used in our first survey (26) have been made available in an publicly accessible reference repository (`http://bit.ly/ZWoY3X`). The online bibliography is maintained by this author but importantly is open for contributions from any user of Mendeley (which can be used for free).

The remainder of this paper is split into sections relating to different fields of problem understanding. Each section explores the most notable papers published or made available since (26), from approximately 2011 through to early 2013. The fields explored are: Benchmark Problems; Problem Decomposition & Multiobjectivisation; Landscape Analysis; Problem Difficulty; and Algorithm Selection & Performance Prediction. Of specific interest are three key and recent literature surveys. However, it should be noted that active research in problem understanding is not limited to those fields alone, only that they are the ones with recent publications and/or activity.

2. SURVEY

The survey below is split into subjects, although many of the papers can be seen to overlap these arbitrary labels. Some selected papers of note, when primarily introduced, are also annotated with some key words of interest [[*given in italics between double brackets*]].

2.1 Literature Surveys

Before diving into the body of new practical an theoretical research in problem understanding, it is important to make note of some key publications in the field - namely those that provide good surveys of a specific subject and which provide a guiding light for research in those individual areas.

Two key surveys have been published on the topic of (fitness) landscape analysis: (31) and (24). These two papers provide an excellent reference resource for researchers in the field, giving a wealth of useful references as well as outlining more strict definitions and descriptions of the basics of landscape analysis. Furthermore, the surveys reflect on past and recent work in the area and suggest some future directions for researchers in the field.

(24) provides a thorough and important survey of the different landscape analysis approaches and problem difficulty measures. In addition to the survey, the paper concludes that there needs to be a shift in focus from predicting prob-

GECCO'13 Companion, July 6–10, 2013, Amsterdam, The Netherlands.

lem difficulty to characterising problems and selecting based on algorithm suitability. A view also held by this author.

Similarly, the survey presented by (31) provides an in-depth and detailed description of landscape analysis, methods and basic principles. The survey provides a fundamental literary reference material which outlines and defines much of basic knowledge and areas of research in landscape analysis. The paper is a key resource for any researcher in landscape analysis and provides an excellent reference point for key terms and definitions which creates a shared understanding for future work, such as (4).

Another recent survey of note is presented by (36) which provides an in-depth review of dynamic optimisation problems. Continuing work on the analysis of dynamic optimisation problems (35), (36) explore, through the literature review, the effect of representations and variation operators on dynamic optimisation problems The review makes specific note of problem features which are unique to dynamic optimisation problems, drawing specific attention to these features.

The importance of these survey cannot be overstated. Future work should be conducted to continue this kind of focused survey to help shape the field, focus research and highlight under referenced work of importance and high impact to the subject. Similar work, although less recent, includes (30) survey on the various fitness functions in evolutionary robotics literature and an earlier survey on fitness approximation and meta-modelling (18) which has a direct relevance to many problem difficulty measures.

2.2 Benchmark Problems

Benchmark problems are a class of optimisation problems that are either artificially constructed or curated from real-world instances that are suitable for use as an analysis tool. Benchmark problems are primarily used to benchmark optimisation algorithm performance and provide a means of fairly (as far as is possible) comparing different optimisation techniques. Benchmarks are also used during the development of new algorithms and for tuning existing algorithms prior to application to more complex, real-world optimisation problems. Ideally, a benchmark problem is well understood with the primary features of the problem known. This should especially be the case for artificial benchmark problems, such as those built using benchmark toolkits and suites.

As is shown by (22), work on artificial benchmark problems and problem suites is still ongoing. (22) [[*Large Scale Optimisation, Benchmark Problems, Continuous Optimisation*]] present 15 new large-scale global optimization benchmark problems to extend the earlier 2010 CEC test problems. The paper extends an existing suite of benchmark problems, adding addition features such as non-uniformity, transforming functions and subcomponents. The work represents an update to existing work and maintains a focus on large-scale global optimisation problems. Like all other benchmark problem suites, there still remains to be seen a single unified tool which covers all varieties of optimisation problems. Although the utility and accessibility of such a suite is debatable.

(3) [[*Real-world Benchmarks, Benchmark Problems, Scheduling, Combinatorial Optimisation, Landscape Analysis*]] present another test suite, although in this case more closely related to real-world problems, containing 23 scheduling problem instances derived from real-world data collected for the preventive maintenance scheduling problem from the power industry. The paper firstly explores the key features of the problem instances (namely a highly non-linear function mapping with a rugged landscape) which is then complimented with an experiment to demonstrate that local search methods are equally effective at solving this problem as simple Evolutionary Algorithms. The paper provides a useful analysis of the benchmark problem and is a good example of how to present benchmark problems in the literature. Further work in the area of methods for encapsulating, analysing and presenting such real-world benchmark problems could greatly benefit the wider field of Evolutionary Computation and optimisation in general.

Expanding on previous benchmark problem suites (21), (37) [[*Benchmark Problems, Many-objective Optimisation*]] propose a new many-objective optimisation benchmark suite with "simple" and "complicated" Pareto sets. The problem suite provides a newly scalable benchmark toolkit which is designed with the specific intention of examining EA performance over increasing dimensionality and Pareto set and front complexity. The suite addresses many of the issues associated with early multi-objective benchmark suites - namely, their simple Pareto sets.

The above papers all represent three new benchmark suites, one related a specific real-world problem and the other two artificial problem toolkits. While not strictly academic, work is needed to draw together existing benchmark optimisation problems as a single, accessible resource to more effectively enable researchers to access these new and old research tools.

2.3 Problem Decomposition & Multiobjectivisation

Whilst technically two separate methods, problem decomposition and multiobjectivisation represent two approaches to manipulating problems to make them more solvable by optimisation algorithms. These methods are primarily addressed by the more traditional optimiser focused research community but also represent a bridge of sorts between better analysis and understanding of problems and the associated rewards in optimisation performance.

These two methods have been grouped here as a link can be seen between the two potentially opposite approaches – one breaking down and simplifying the problem while the other increases the objective space complexity to "unfold" overly condensed functions in order to make it easier to explore.

Optimisation problem decomposition (not factoring or subgoaling) is a long researched approach to optimisation which solves difficult problems by breaking them up into smaller, easier to solve problems which can be solved separately, either in parallel or sequentially. Often it is required that these problems be separable and can be solved independently. This technique has clear relevance to problem understanding research, requiring analytical methods for determining the extent to which a problem can be decomposed and how these sub-problem relates.

Indeed, many combinatorial problems can be decomposed into a series of sub-problems which represent the independent key components (or "elementary landscapes") that, when superimposed upon on another, create the landscape of the original problem. (11) [[*Elementary Landscapes, Problem Decomposition, Combinatorial Optimisation, Landscape Anal-*

ysis]] presents a 5-step method for decomposing any arbitrary combinatorial optimisation problem with symmetric underlying neighbourhoods which is demonstrated on the quadratic assignment problem. The authors also outline means of automating parts of the method to assist researchers. The method is known to be limited by a bounding (an unspecified small constant) on the number of elementary components which in turn limits the complexity of the problems to which it can be applied; such as NK landscapes. Extensions to this work which mitigate this constraint would clearly be a significant new contribution to the field. The authors also note planned future work in the area of applying the method *en masse* and producing an automated tool to facilitate use of the method.

Conversely, multiobjectivisation is the process by which adding artificial objective functions to single-objective or lower dimension problems increases the "evolveability" of the problem, making them easier to solve or improving search results. This relatively new field is receiving increasing attention in the literature and will no doubt continue to attract significant research articles on subject.

Again, not directly related to problem understanding, the method presented in (23) [[*Multiobjectivization, Optimisation*]] looks at artificially increasing the number of objectives in an optimisation problem to improve the search in some why, either making it more easy to solve or increasing solution diversity. Multiobjectivization, and to some extent Novelty Search, both illustrate a bridge between problem understanding and more traditional optimisation research.

Clearly multiobjectivization bares some similarities to problem decomposition, perhaps in the opposite direction however. Some analysis of the effect of artificial objectives on landscape features could both help to increase the impact of landscape analysis as an incorporated part of an optimiser as well as understand the effect of these artificial objectives.

Although not directly related to problem decomposition or multiobjectivization, building on earlier work (19), (20) [[*Novelty Search, Optimisation*]] explores the concept of rejecting traditional notions of objectivity and continues the study on optimising for "novelty" rather than objective performance. This work touches on the area of dynamically adjusting or re-framing optimisation problems to improve the quality of search. Again, while the work does not directly relate to problem understanding, the dynamics of the search for novelty could potentially provide some useful insights in to how optimisation problems could be made more "accessible" in terms of evolvability, relating to work such as (13).

These methods of manipulating, adjusting and re-framing optimisation problems each place a new lens through which to examine optimisation problems. To what extent is an optimisation problem "fixed" and to what extent can the landscape be manipulated for specific needs, such as making them easier to solve or more suitable for different optimisers. Can the landscape be manipulated at all, or are these methods simply translations of the same underlying landscape and to what extent are the underlying features retained in these transformations? These, and other questions, have a direct impact on landscape analysis methods, discussed below.

2.4 Landscape Analysis

In the context of this brief literature survey, landscape analysis covers all forms of landscape analysis, from fitness landscapes to phenotypic landscape topography, and is primarily concerned with methods for identifying and describing landscape features through sampling techniques and feature descriptors. Landscape analysis is by far the most extensively researched topic in the field of problem understanding. A number of important publications have been published in the last two years, including two survey papers which represent a significant milestone in development and maturing of the field. Landscape analysis is one of the core topics of problem understanding research and often forms the basis for research in problem difficulty assignment, algorithm selection and performance prediction; which are discussed below.

Using recently developed general landscape analysis techniques (31; 32), (34) [[*Landscape Analysis, Real-World Optimisation*]] analyses the well known vehicle routing problem. The paper provides one of the few practical applications of general landscape analysis methods to realistic optimisation problems. It should be noted that further work on the practical application of these general methods is sorely needed to better understand their practical limitations - although Pitzer et al. are likely already in the process of undertaking this daunting task!

Another recent example of practical applications of these methods is given in (12) [[*Landscape Analysis, Elementary landscapes*]]. Building on (11), (12) paper develops existing work on autocorrelation (coefficient and length) for the Quadratic Assignment Problem (QAP). The paper provides an "exact express" for the autocorrelation of QAP instances and, through extensive experiments, actually explores and characterises the "difficulty" of this problem class. Specifically, the paper provides autocorrelation values for the QAP instances in the QAPLIB database (7). The paper represents one in a very underdeveloped area of problem understanding - practical application of analysis methods on existing problems. Few full examples of applications of work in this field are actually provided in the literature. An industrious research would do well to provide a more comprehensive application and analysis of methods on existing problems to more fully describe existing problems and "fill out" this gap in the literature.

Indeed, (34) and (12) are only two of very few examples of practical applications of landscape analysis techniques. This failing in the literature is noted by (24) [[*Problem Difficulty, Landscape Analysis, Real-World Optimisation, Algorithm Selection, Literature Survey*]], who argue that despite the "large" number of landscape analysis methods presented in the literature, few have been used practically on benchmark or real-world optimisation problems. Whether that is because the methods are too expensive or limited, the paper argues that the methods need to be more frequently applied.

Many real-world problems are "black-box" and cannot be easily analysed. Selecting for solving such methods is therefore a difficult task, especially given that in many cases the number of evaluations is severely limited. Landscape analysis provides a means of quickly and efficiently qualifying the features present in the problem in order to aid the process of tuning and selecting optimisation problems as well as gaining a better understanding of the problem being solved.

(2) [[*Landscape Analysis, Landscape Features, Algorithm Selection*]] presents a set of problem independent features that analyse the fitness landscape of black-box optimization

problems which can be used to aid the selection of the most optimal optimisation method. The paper proposes 10 features, divided into three categories: problem definition features, hill climbing features, and random point features. The method and features are demonstrated through empirical numerical experiments. Given other recent advancements in landscape analysis and feature descriptors, it would be interesting to see to what extent the proposed features can be integrated with other existing sets. As with all landscape feature descriptors, there still remains to be seen a comprehensive taxonomy of problems with associated features and also solvers that best solve problems with given features - this represents a significant body of work which can only truly be undertaken as a joint effort by the wider community.

Another paper on fitness landscape analysis, the work in (8) [[*Landscape Analysis, Landscape Features, Algorithm Selection*]] is framed within the context of better understanding the topography of a landscape will enable better algorithm selection and avoid the "lengthy" process of trial-and-error application too often associated with Evolutionary Algorithm selection and tuning. The paper analyses through experimental study a set of landscape feature descriptors and highlights the strong underpinning of many algorithm selection methods on landscape analysis and feature identification techniques.

While not entirely novel (see (28)), characterising continuous optimisation problems is relatively under-researched compared combinatorial problems. Breaking away from the trend for landscape analysis on combinatorial optimisation problems, (1) [[*Landscape Analysis, Heuristic-Problem Interactions, Continuous Problems*]] explore methods for characterising continuous optimisation problems. A method called *length scale* for characterising continuous optimisation problems. The technique measures the fitness distance of continuous optimisation problem mappings and applies it to the BBOB'10 benchmark problems.

Unlike many landscape analysis studies, rather than look at proposing new features, (38) [[*Landscape Analysis, Multi-objective optimisation*]] explore the the specific task locating a feature called "knees" in multi-objective Pareto fronts. The paper presents a method for finding "proper knees" in Pareto fronts which could be useful for landscape analysis techniques. However, in the paper itself, the authors use the approach to construct two new optimisation methods that exploit the presence of this feature in optimisation problems. The paper does show how new or alternative mathematical definitions of a known problem feature can be useful for improving algorithm performance. Conversely, the work could be considered as a tailored method for sampling the problem space in search of the specific feature. This raises the question: are feature specific sampling methods useful or needed in landscape analysis?

2.4.1 Links to Algorithm Selection

In addition to gaining a better understanding of a problem's features, landscape analysis is a core aspect of many algorithm selection methods. However, as the number of methods and diversity of problems massively increase, the task of algorithm selection is becoming ever more difficult. Using expected run-time cost predictions for each optimising algorithm, (6) [[*Landscape Analysis, Exploratory Landscape Analysis, Algorithm Selection, Algorithm Performance Prediction*]] presents a method of systematic sampling which identifies "low-level" problem features taken from (27) and uses one-sided support vector regression to learn which algorithms are best suited to solve which problems. The method is demonstrated using the BBOB optimisation problems (5; 16). The authors suggesting an interesting idea, that given the knowledge of the low-level features provided in work such as (27), these features could be used to generate artificial benchmark problems which a desired mix of features. Indeed, such an approach would surely represent the natural joining of the currently disconnected fields of problem analysis and synthesis.

2.4.2 Biology

A parallel to fitness landscape analysis in Evolutionary Computation is given in Theoretical Biology (e.g., (40)). In the context of biology, significant efforts are being invested in developing methods for fitness landscape analysis. As (13) states, fitness landscapes are central in the theory of adaptation in the natural world. Work in this cousin subject can, on occasion, have direct impact on work in the EC community. For example, (13) propose a graph approach to analysing fitness landscapes that could potentially be applied directly to discretely encoded optimisation problems. It would be interesting to survey this larger but less directly applicable body of work and see to what extent methods developed there can be applied to optimisation and EC work. Landscape Analysis, Theoretical Biology, Biological Evolution

Another example of landscape analysis in the context of biological evolution, (14) [[*Landscape Analysis, Theoretical Biology, Biological Evolution*]] explores the strength of genetic interactions through mutation and crossover on artificial instances with varying degrees of "ruggedness" and the effect on evolutionary "accessibility". The concept of "evolutionary accessibility" is an interesting one which has many parallels to problem difficulty prediction in optimisation. In many cases, the "accessibility" of the optima for different types of optimisers could be used to describe the difficulty of the problem in the context of specific types of optimisers or conversely the suitability of different optimisers for a given problem.

2.5 Problem Difficulty

The subject of problem difficulty considers the task of assessing, assigning and predicting the inherent difficulty of optimising an optimisation problem. In effect, problem difficulty assessment methods grade optimisation problem hardness. However, recent trends are moving towards less arbitrary, optimiser independent measures to a more joined approach of problem difficulty with respect to different optimisation algorithm classes. Many modern difficulty measures take into account the pairing of problem features identified through landscape analysis and algorithm features identified through experimentation on benchmark problems. The natural progress of such techniques is to then consider the problem of algorithm selection and to what extent problem difficulty can be used to predict algorithm performance. Research on these specific topics are discussed in the following sub-section.

In the context of the combinatorial problems, (17) [[*Problem Difficulty, Combinatorial Optimization*]] describe the potential for the "granularity" of the objective function scales in combinatorial problems to effect the difficulty of a problem. Experimental analysis using well known EA optimisers is conducted which shows that the coarser granularity objectives reduces the rate of convergence for these conventional optimisers on low-objective problems. Interestingly, the results suggest an inverse correlation between granularity and number of objectives, where decreasing the granularity of a objective increases the performance of the EAs as the number of objectives increase (in terms on many objectives). It is these kinds of relations which are important to problem analysis and understanding how features, such as objective function granularity, affects the performance of algorithms and dynamics of problem-optimiser interactions. Further studies of this kind would provide a significantly improved knowledge base on the interactions between problem features and the effect on the evolvability or difficulty of a problem.

In the context of problem neighbours and locality, (15) [[*Problem Difficulty, Genetic Programming*]] explore the effect of locality in representations and the effect on performance through "evolvability" (closely linked to "accessibility" in evolutionary biology). The paper explores the effect of locality on GP and uses it as a difficulty indicator. The paper provides a comprehensive analysis of this approach to measuring difficulty, both on discrete and continuous problems. One measure, Def1, is shown through experiments as the most effective predictor of problem difficulty. While locality and neighbours have been explored extensively in the past (26), our understanding of neighbourhoods and locality in the context of heuristics is still limited. Is locality a property of a problem landscape, representation or heuristic?

A continuation of the work on elementary landscapes (11) and fitness-distance correlation (12) is presented in (9) [[*Elementary Landscapes, Problem Decomposition, Combinatorial Optimisation, Landscape Analysis, Problem Difficulty*]]. The paper presents a closed-form expressions for the fitness-distance correlation using problem decomposition through elementary landscapes. Interestingly, the theoretical work casts some doubt on the efficacy of fitness-distance correlation for quantifying problem difficulty (in the context of sample problems with binary string encodings) and challenges the current assumption that problem difficulty is primarily controlled by the higher order elementary components. No doubt, Chicano and Alba are continuing work in this area and future extensions and generalisations of this line of enquiry will be of great interest.

Again, reinforcing the strong links between subjects, (39) [[*Problem Difficulty, Algorithm Selection, Combinatorial Problems*]] identify large scale analysis of problem features as one of the key tasks required for effective algorithm selection, primarily as means of developing an extensive knowledge base for data-mining. Based on a survey of literature on landscape features for combinatorial optimisation problems, (39) explores how to select problem feature descriptors that most useful for problem difficulty prediction and algorithm selection. This work compliments much of the work on HeuristicLab and related landscape analysis techniques (32). However, care must be taken when undertaking such disconnected studies that work in the area does not become disjointed. Work is needed to agree, through consensus, a shared means of providing landscape feature descriptors, experimental results and so on; a topic discussed at great length at the first GECCO-UP workshop (25).

2.6 Algorithm Selection & Performance Prediction

The problem of algorithm selection has long been considered in the literature. However, recently, renewed interest in the problem is resulting new papers on the topic. In simple terms, the algorithm selection problem is concerned with selecting the most appropriate algorithm for solving a given optimisation problem or problem class. This clearly has strong links with problem difficulty and problem-algorithm matching taxonomies. Algorithm selection is also closely related to algorithm performance predictors, which can be thought of as the counterpart to problem difficulty and explores methods for predicting how effectively an algorithm will solve a given problem. Good performance predictors can obviously be exploited by algorithm selection methods and so have been grouped here as literature on the two methods are often linked or combined.

There are a number of recent works in these two related fields, such as (10) [[*Algorithm Performance Prediction, Elementary Landscapes, Problem Decomposition, Landscape Analysis*]] which combines local optima networks and elementary landscapes (such as in (11)) to produce a method for predicting optimisation algorithm performance.

Meanwhile, in the context of continuous optimisation problems, (29) [[*Algorithm Performance Prediction, Algorithm Selection, Continuous Problems*]] explores algorithm selection through performance prediction models. The paper adapts a meta-learning framework (employing ANN regression models) and tailors it for application to continuous optimisation problems. The model ties the (assumed) independent parameters of the various parameterisations of CMA-ES and a set of problem landscape feature descriptors taken from the literature.

The approach represents a trend in algorithm selection for using machine learning methods to try and build associations and performance predictors based on empirical results in the form of historical data. This approach is admitted less strict than taxonomic approaches, but also provides a currently workable and more easily adapted approach to algorithm selection.

In contrast to the machine learning approach, by drawing on earlier work in landscape analysis, (33) [[*Algorithm Selection, Landscape Analysis*]] brings together landscape analysis and algorithm selection. The study examines the efficacy of these landscape analysis methods for algorithm selection, using a vector of landscape features and historical algorithm performance on these problems to predict future performance. Unlike more machine learning heavy techniques (29), this algorithm-feature matching technique provides a more "open-box" approach to algorithm selection which is often preferable for giving users confidence in the selection and importantly the reasoning behind the selection. This is especially important when considering very expensive problems or those that can only be optimised once.

3. SUMMARY

Evidently, there have been significant developments in the area of landscape analysis. The majority of recent literature is dedicated to this specific topic which highlights it's growing importance and researcher's awareness to potential developments in the area. Of note is the work conducted by (31) who have brought together through a detailed survey, the key works in the field. Similarly, (24) conduct a similarly useful survey in the field and also provide some interesting insights in to future directions of research in this area. The two papers present a significant contribution to landscape analysis and provide key references for future work.

Considering the literature found for this survey, many of which are presented in journals, it is clear that work in the area of problem understanding is receiving renewed attention by the (mostly EC) research community. Although the literature reference herein is limited to only a few fields of study, it should be noted that active research in problem understanding is not limited to those fields alone, only that they are the ones with recent publications and/or activity.

3.1 Looking Forward

Bibliographic Survey – following the second workshop on problem understanding, again held at GECCO, this author intends to continue to pursue the task of collecting and referencing new work in the field of problem understanding (send references to k.mcclymont@exeter.ac.uk or update the online bibliography direct at `http://bit.ly/ZWoY3X`).

Exeter benchmark problems – as part of the above survey, a significant body of work has been undertaken to identify many features, feature descriptors and recent advances benchmark suites. This has been conducted as part of an ongoing project to develop a general, comprehensive and robust benchmark problem toolkit which is aimed at not just providing a means of analysing optimisation methods but also assessing methods from problem understanding, such as landscape analysis. An early release of the Exeter toolkit (2.0) will be made available online at `http://www.kentm.co.uk/` prior to the GECCO'13 conference.

4. ACKNOWLEDGMENTS

This research was supported by an EPSRC project (Grant No. EP/K000519/1) in conjunction with Mouchel Ltd.

References

[1] Length scale for characterising continuous optimization problems. In C. A. Coello, V. Cutello, K. Deb, S. Forrest, G. Nicosia, and M. Pavone, editors, *Parallel Problem Solving from Nature - PPSN XII*, volume 7491 of *Lecture Notes in Computer Science*, pages 407–416. Springer Berlin Heidelberg, 2012.

[2] T. Abell, Y. Malitsky, and K. Tierney. Fitness landscape based features for exploiting black-box optimization problem structure. Technical report, Technical Report TR-2012-162, IT University of Copenhagen, 2012.

[3] A. Almakhlafi and J. Knowles. Benchmarks for maintenance scheduling problems in power generation. In *Evolutionary Computation (CEC), 2012 IEEE Congress on*, pages 1–8, 2012.

[4] D. Ashlock and S. McNicholas. Fitness landscapes of evolved apoptotic cellular automata. *Evolutionary Computation, IEEE Transactions on*, 17(2):198–212, 2013.

[5] A. Auger, S. Finck, N. Hansen, R. Ros, et al. Bbob 2009: Comparison tables of all algorithms on all noiseless functions. Technical Report RT-388, INRIA, 2010.

[6] B. Bischl, O. Mersmann, H. Trautmann, and M. Preuß. Algorithm selection based on exploratory landscape analysis and cost-sensitive learning. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, GECCO '12, pages 313–320, New York, NY, USA, 2012. ACM.

[7] R. Burkard, S. Karisch, and F. Rendl. Qaplib âĂŞ a quadratic assignment problem library. *Journal of Global Optimization*, 10(4):391–403, 1997.

[8] P. Caamano, F. Bellas, J. Becerra, V. Diaz, and R. Duro. Experimental analysis of the relevance of fitness landscape topographical characterization. In *Evolutionary Computation (CEC), 2012 IEEE Congress on*, pages 1–8, 2012.

[9] F. Chicano and E. Alba. Exact computation of the fitness-distance correlation for pseudoboolean functions with one global optimum. In J.-K. Hao and M. Middendorf, editors, *Evolutionary Computation in Combinatorial Optimization*, volume 7245 of *Lecture Notes in Computer Science*, pages 111–123. Springer Berlin Heidelberg, 2012.

[10] F. Chicano, F. Daolio, G. Ochoa, S. Vérel, M. Tomassini, and E. Alba. Local optima networks, landscape autocorrelation and heuristic search performance. In *Parallel Problem Solving from Nature-PPSN XII*, pages 337–347. Springer, 2012.

[11] F. Chicano, G. Luque, and E. Alba. A methodology to find the elementary landscape decomposition of combinatorial optimization problems. *Evolutionary Computation*, 19(4):597–637, 2011.

[12] F. Chicano, G. Luque, and E. Alba. Autocorrelation measures for the quadratic assignment problem. *Applied Mathematics Letters*, 25(4):698–705, 2012.

[13] K. Crona, D. Greene, and M. Barlow. The peaks and geometry of fitness landscapes. *Journal of Theoretical Biology*, 317(0):1–10, 2013.

[14] J. Franke and J. Krug. Evolutionary accessibility in tunably rugged fitness landscapes. *Journal of Statistical Physics*, 148(4):705–722, 2012.

[15] E. Galván-López, J. Mcdermott, M. O'Neill, and A. Brabazon. Defining locality as a problem difficulty measure in genetic programming. *Genetic Programming and Evolvable Machines*, 12(4):365–401, 2011.

[16] N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization bench-marking 2009: Experimental setup. Technical Report RR-6828, INRIA, 2009.

[17] H. Ishibuchi, M. Yamane, and Y. Nojima. Difficulty in evolutionary multiobjective optimization of discrete objective functions with different granularities. In R. C. Purshouse, P. J. Fleming, C. M. Fonseca, S. Greco, and J. Shaw, editors, *Evolutionary Multi-Criterion Optimization*, volume 7811 of *Lecture Notes in Computer Science*, pages 230–245. Springer Berlin Heidelberg, 2013.

[18] Y. Jin. A comprehensive survey of fitness approximation in evolutionary computation. *Soft Computing*, 9(1):3–12, 2005.

[19] J. Lehman and K. O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. 2008.

[20] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.

[21] H. Li and Q. Zhang. Multiobjective optimization problems with complicated pareto sets, moea/d and nsga-ii. *Evolutionary Computation, IEEE Transactions on*, 13(2):284–302, 2009.

[22] X. Li, K. Tang, M. N. Omidvar, Z. Yang, K. Qin, and H. China. Benchmark functions for the cec'2013 special session and competition on large-scale global optimization. Technical report, Evolutionary Computation and Machine Learning Group, RMIT University, 2013.

[23] D. Lochtefeld and F. Ciarallo. Multiobjectivization via helper-objectives with the tunable objectives problem. *Evolutionary Computation, IEEE Transactions on*, 16(3):373–390, 2012.

[24] K. M. Malan and A. P. Engelbrecht. A survey of techniques for characterising fitness landscapes and some possible ways forward. *Information Sciences*, 2013. To Appear.

[25] K. McClymont and E. Keedwell. 1st understanding problems workshop (gecco-up). In T. Soule, editor, *GECCO Companion '12: Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference companion*, New York, NY, USA, 2012. ACM. 910122.

[26] K. McClymont, D. Walker, and M. Dupenois. The lay of the land: a brief survey of problem understanding. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference companion*, GECCO Companion '12, pages 425–432, New York, NY, USA, 2012. ACM.

[27] O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, and G. Rudolph. Exploratory landscape analysis. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation. GECCO âĂŹ11*, pages 829–836. ACM, New York, NY, USA, 2011.

[28] C. Müller and I. F. Sbalzarini. Global characterization of the cec 2005 fitness landscapes using fitness-distance analysis. In *Applications of Evolutionary Computation*, volume 6624 of *Lecture Notes in Computer Science*, pages 294–303. Springer Berlin Heidelberg, 2011.

[29] M. A. Muñoz, M. Kirley, and S. K. Halgamuge. A meta-learning prediction model of algorithm performance for continuous optimization problems. In *Parallel Problem Solving from Nature-PPSN XII*, pages 226–235. Springer, 2012.

[30] A. L. Nelson, G. J. Barlow, and L. Doitsidis. Fitness functions in evolutionary robotics: A survey and analysis. *Robot. Auton. Syst.*, 57(4):345–370, 2009.

[31] E. Pitzer and M. Affenzeller. A comprehensive survey on fitness landscape analysis. In *Recent Advances in Intelligent Engineering Systems*, volume 378 of *Studies in Computational Intelligence*, pages 161–191. Springer Berlin Heidelberg, 2012.

[32] E. Pitzer, M. Affenzeller, A. Beham, and S. Wagner. Comprehensive and automatic fitness landscape analysis using heuristiclab. In *Computer Aided Systems Theory âĂŞ EUROCAST 2011*, volume 6927 of *Lecture Notes in Computer Science*, pages 424–431. Springer Berlin Heidelberg, 2012.

[33] E. Pitzer, A. Beham, and M. Affenzeller. Automatic algorithm selection for the quadratic assignment problem using fitness landscape analysis. In *Evolutionary Computation in Combinatorial Optimization*, volume 7832 of *Lecture Notes in Computer Science*, pages 109–120. Springer Berlin Heidelberg, 2013.

[34] E. Pitzer, S. Vonolfen, A. Beham, M. Affenzeller, V. Bolshakov, and G. Merkuryeva. Structural analysis of vehicle routing problems using general fitness landscape analysis and problem specific measures. In *14th International Asia Pacific Conference on Computer Aided System Theory*, pages 36–38, 2012.

[35] P. Rohlfshagen and X. Yao. Dynamic combinatorial optimisation problems: an analysis of the subset sum problem. *Soft Computing*, 15(9):1723–1734, 2011.

[36] P. Rohlfshagen and X. Yao. Dynamic combinatorial optimization problems: A fitness landscape analysis. In *Metaheuristics for Dynamic Optimization*, volume 433 of *Studies in Computational Intelligence*, pages 79–97. Springer Berlin Heidelberg, 2013.

[37] D. Saxena, Q. Zhang, J. a. A. Duro, and A. Tiwari. Framework for many-objective test problems with both simple and complicated pareto-set shapes. In *Evolutionary Multi-Criterion Optimization*, volume 6576 of *Lecture Notes in Computer Science*, pages 197–211. Springer Berlin Heidelberg, 2011.

[38] P. K. Shukla, M. A. Braun, and H. Schmeck. Theory and algorithms for finding knees. In *Evolutionary Multi-Criterion Optimization*, volume 7811 of *Lecture Notes in Computer Science*, pages 156–170. Springer Berlin Heidelberg, 2013.

[39] K. Smith-Miles and L. Lopes. Measuring instance difficulty for combinatorial optimization problems. *Computers & Operations Research*, 39(5):875–889, 2012.

[40] I. G. Szendro, M. F. Schenk, J. Franke, J. Krug, and J. A. G. M. de Visser. Quantitative analyses of empirical fitness landscapes. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01005, 2013.