The Lay of the Land: a Brief Survey of Problem Understanding

Kent McClymont
Computer Science
University of Exeter, UK
K.McClymont@ex.ac.uk

David Walker
Computer Science
University of Exeter, UK
D.J.Walker@ex.ac.uk

Max Dupenois
Computer Science
University of Exeter, UK
M.P.Dupenois@ex.ac.uk

ABSTRACT

Optimisation research often concentrates on developing heuristic methods for solving a given optimisation problem, however a growing body of work surrounds the understanding and analysis of the problem itself to facilitate heuristic development and selection. We outline the broad themes that this field encompasses, presenting some of the techniques and approaches that have been taken, and present a sociometric view of selected publications in the field. We conclude by proposing a collaborative system to allow researchers working on problem understanding to more easily share results and work together.

Categories and Subject Descriptors

G.1.6 [**Optimization**]

General Terms

Theory.

Keywords

Problem understanding, optimisation problems, evolutionary algorithms.

1. INTRODUCTION

This brief literature survey aims to address and identify the general trends in optimisation research relating to problem understanding. The paper aims to highlight the primary areas of problem understanding research and provide a context from which more comprehensive surveys can be conducted. The survey explores papers published prior to 2012.

In this paper, the term *problem understanding* is defined as the identification, analysis, synthesis, classification and theoretical scrutiny of optimisation problems. The term heuristic is used generically to refer to an optimisation algorithm, such as a meta-heuristic.

In many cases, optimisation research is heavily biased towards the analysis and development of better heuristic methods for solving optimisation problems. As we show later, whilst the research community acknowledges the significance of the structure and characteristics of optimisation problems and the strong relationship between the problem and the efficacy of each method there has been comparatively little research conducted to analyse these problems independently of any one solving method.

The paper is organised as follows. An analysis of past and current research trends is conducted in Section 2 where key areas of interest are identified and potentially fruitful avenues of inquiry signposted. Section 3 provides an overview of the trends in problem understanding publications and identifies the Evolutionary Computation community as one of the key instigators of this research. Section 4 examines the sociometric structure of problem understanding research community by drawing on the author and reference information available for the 133 publications identified for this study (full reference data is available online at `http://algorithmica.org.uk`).

Finally, in Section 5 a case is made for the development of a collaborative system to foster new research specifically in the area of problem understanding. The proposed model is definitely "blue sky" but hopefully feasible for an emerging field of research. Indeed, we argue that future research will be greatly aided by adopting the proposed model (or something similar) at this early stage, allowing researchers easy access to what will undoubtedly be a growing set of resources and body of literature as well as enable fruitful discussion and regularly updated notes relating to potentially rewarding areas of study.

2. LITERATURE SURVEY

This section explores the key research themes identified in the research literature and discusses specific examples where appropriate. Due to the size constraints of this publication, only selected literature is referenced.

2.1 Research Themes

Table 1 provides a breakdown of the research themes identified in the surveyed publications. Despite the apparent disparity in the sources and connectivity of the field, there are some common threads which are more popular: methods or landscape analysis; construction of test problems; methods for determining problem difficulty; analysis of problem and heuristic interaction; construction and analysis of dynamic problems.

Suitable literature from all the identified themes are discussed below. Generally, the themes can be categorised into problem analysis, problem construction, problem ontologies and frameworks, and visualisation. It should be noted that the loose categorisation of research themes is supported by the body of literature rather than any specific classification scheme and could be improved upon with a more specific ontological study.

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA.

Research Theme	Number of Publications
Landscape Analysis	54
Test Problems	32
Problem Difficulty	21
Problem/Heuristic Interaction	16
Dynamic Problems	15
Dynamics	11
Theory	11
Classification	9
Combinatorial Encodings	9
Continuous Problems	8
Feature Definition	8
Problem Generators	6
Feature Identification	5
Constrained Problems	4
Neighbourhood Structures	4
Problem Relation	3
Framework	2
Taxonomy	2
Real-world Problems	1
Solution Encoding	1
Visualisation	1
Heterogeneous Encoding Problems	1

Table 1: The distribution of topics under the auspices of problem understanding. A clear majority of publications are concerned with landscape analysis and test problem design.

2.1.1 Landscape Analysis

Landscape Analysis and Feature Identification.

Landscape Analysis is easily the most studied area of research and covers a vast range of areas from NK Landscape Analysis [24] and [35] to Exploratory Landscape Analysis (ELA) [16]. The subject in general refers to methods for analysing the mapping in optimisation problems from the parameter (or decision) space to objective (or fitness) space.

Pelikan [24] presents an empirical study of NK landscapes with an additional focus on different measures of problem difficulty of fitness landscapes. Another example of NK Landscape analysis can be found in [35] which uses experimental set-based fitness landscape analysis to better characterise the "geometry" of the problem to enhance set-based multiobjective search approaches.

Many landscape analysis methods presented in the literature are empirical in nature, such as ELA [16], and use data collected through experimental runs to search, identify and quantify the features and characteristics of a problem. For example, [19] uses statistical measures to characterise the global topographical properties of black-box continuous optimisation problems. In addition, most landscape analysis studies are closely tied to specific feature definitions, such as [30] which discusses a means of sampling and analysing k-satisfiability problems to measure landscape ruggedness using the autocorrelation function.

There are numerous other papers discussing landscape analysis – far too many to list here. Generally, the work relates to empirical methods for exploring and sampling problems to discover specific features and make specific measurements. A more detailed survey on this literature would be well placed to identify key methods, features and measurements. In addition, while approaches such as ELA aptly demonstrate their worth, it would be interesting to explore a more general framework for landscape analysis which can encompass a wider variety of approaches that can cover the spectrum of problem encoding schemes – a more holistic approach to landscape analysis.

Theory, Feature Definition and Neighbourhood Structures.

The theoretical foundations of problem understanding are mostly discussed in the context of wider studies or proposals for new test problem suites. For example, in [6], the theory of parameter linkages is considered by adapting existing test problems (such as the DTLZ problems [5]) and examining the effects of introducing these linkages. In the same manner, the survey of test problems conducted in [10] defines a number of interesting problem features which are then used to support a newly proposed test problem suite.

Indeed, as in [10], feature definitions appear in a number of different studies not directly focused on an analysis of the efficacy of the proposed features. For example, in addition to proposing the ELA method, [16] also defined a wide set of features relevant to the black-box optimization bench-marking (BBOB) optimisation problems [8].

Neighbourhood structures are also studied in a number of papers. For example, [24] makes a note of the neighbourhood structures in NK landscapes and conducts a study of shuffled NK landscapes with nearest-neighbor interactions. In contrast, [13] examines the impact of the global structure of a problem and how that impacts on search performance. Importantly, [13] provides an examination of the "big valley" structure of many test problems, which is not always the case for real-world problems. Hauschild [9] shows how the theory of neighbourhood structures can be used to inform problem difficulty measures and argues that previous difficulty measures were in fact using the wrong neighbourhood structures which contributed to their poor performance.

It should be noted that neighbourhood structures is a term used loosely in the literature. In some contexts, it is used to refer to topological structures that are specific to the problem and independent of any heuristic, such as the "big valley" identified in [13]. The term is also used to refer to the dynamic neighbourhoods dictated by a specific algorithm's interaction with a specific problem [4].

There are range of problem features already defined in the literature as well as definitions for a variety of neighbourhood structures. These theoretical foundations are very important to the field of problem understanding and further research is needed to unify these features and generalise across the whole family of optimisation problems to provide a strong base on which experimental studies can be grounded.

Problem Difficulty.

In the context of optimisation literature, problem difficulty usually refers to the prediction of an optimisation problem's hardness or difficulty to solve. The term also covers online measurements of difficulty based on an algorithm's performance, rather than predictive measures.

There are many examples of problem difficulty measures, such as [12]. Using Fitness-Probability Cloud (FPC), [12] defines a measure for problem difficulty called Accumulated Escape Probability (AEP). FPC is a measure of "evolvability" on a problem and provides the basis for the numerical AEP predictive problem difficulty measure for EAs. Similarly, [9] focuses on problem difficulty for EAs. Hauschild [9] builds on neighbourhood structures and uses a measure of parameter linkage to create a measure of problem difficulty in the context of EAs.

Although a considerable proportion of problem understanding literature is dedicated to neighbourhood structures and problem difficulty, most of these developments are made in the context of and tailored to EA methods. These methods should be made more general and applicable to problems independent of any heuristic methodology. Characterisation of a problem and its difficulty could be used as a profile of the problem and then matched with heuristics which are known to work well on those types of problem.

Problem/Heuristic Interaction Dynamics.

Another area related to problem difficulty and neighbourhood structures is the study of the interaction between problems and heuristics. Note: the study of problem/heuristic interaction dynamics is not the same as the study of the dynamics of dynamic optimisation problems.

Studies such as [14] and [25] use offline methods to analyse the interaction of heuristic performance on specific problems, such as random decomposable problems, as a proxy for the interaction and efficacy of each heuristic to the problem. This style of experimentation is also used in studies like [20] which examine interactions between problems and/or heuristics with specific features, such as a multi-funnel problems.

Another common approach is to use a combination of landscape analysis and performance measurement to determine the efficacy (or interaction dynamics) of a specific class of algorithm on each problem. For example, [17] uses landscape analysis to explore the interaction between Memetic Algorithms and NK landscapes.

The subject of problem/heuristic interaction dynamics is at the very “edge” of problem understanding and bridges the gap between problem and heuristic research. Consequently, the research in this area is more difficult to identify, often appearing as footnotes in heuristic research, and requires a more detailed literature survey to fully understand the state of current thinking in this area.

2.1.2 Problem Construction

Test Problems and Problem Generators.

Test problems are used extensively in heuristic research to benchmark and demonstrate the performance characteristics of each proposed method. The current trend in test problem research is to provide ‘toolkits’ or ‘suites’ of test problems manually designed with specific features, such as the well known DTLZ [5] and WFG [10] toolkits. Indeed, a large number of test suites have been presented in the literature, such as the DTLZ [5] and WFG [10], BBOB [8], Exeter [15] toolkits. Other sources included competitions like the CEC test problems [31].

Another approach to constructing test problems is by creating problem generators which are capable of generating large numbers of problems with varying features. Some early examples include [28] and [18] and more recently [26]. Problem generators are widely used in dynamic optimisation problem research.

In addition to problem suites and generators, a wide number of individual test problems can be found in the literature. Generally these problems are used to highlight specific cases. For example, noting the general use of “big valley” structures in earlier test problem suites, [13] examines the change in the performance of Evolutionary Algorithm (EA) when a “two funnel” structure is used and demonstrates the EA’s poor performance on these problems.

The creation of test problems is still ongoing with recent examples such as [15] and [11] which look at heterogeneous encodings and complicated Pareto sets respectively. These efforts should be united with existing toolkits and incorporate features identified in studies like [13]. A combined research effort to create a larger encompassing test suite would be beneficial to the field and provide a more up-to-date baseline with which to operate.

Dynamic Problems.

Rohlfshagen and Yao [27] extensively analyse a specific dynamic optimisation problem and examine the correlation between problem parameters and the dynamism of the problem optimum. The paper provides a good study of dynamic optimisation problems and addresses some interesting theoretical aspects, such as representation, which are often overlooked in heuristic focused studies.

Younes et al [36] address the problem of creating generalised benchmarks for dynamic optimisation problems noting, for example, that such benchmarks are often restricted to a certain problem type and do not generalise well to other types.

There is a significant body of research for dynamic problems and it is one of the better explored areas of research which is more connected and focused on “problems” than any other. Interested readers should look at existing summary and introduction papers such as [7].

Encodings (Combinatorial, Continuous, Heterogeneous, Constrained).

As with many of the research themes already identified, work on the analysis of different encodings is scarce, with different test problems providing the largest source of encoding specific analysis. For example, continuous test problems such as DTLZ [5] and WFG [10] do not specifically discuss the merits or effects of using a continuous encoding but do provide analysis of scaling and bias in the search space.

McClymont et al [15] present a preliminary toolkit for heterogeneous (mixed encoding) test problems and identify the encoding as an important aspect of optimisation research. Anecdotally, the issue of encoding and representation is widely recognised in the optimisation research community, frequently remarked upon and often associated with the phrase “Representation, representation, representation” (e.g., in Kenneth A. De Jong’s tutorial “Evolutionary Computation: A Unified Approach” at CEC’10; and in Toby Walsh’s talk at CP2001 Post Conference Workshop on Modelling and Problem Formulation; etc.). However, no clearly identifiable study was found that focused specifically on generally analysing the effect of different representations, rather studies that focused on new, novel and better representations for specific cases, such as [22] or examination of specific representations for specific problems, such as [27], were found.

Much work is needed to better understand the effects of encodings on different problem types and better define how the choice of encoding representation can impact on the difficulty of a problem. Specific studies like [27] illustrate how this type of analysis can be achieved and benefits that can be received by better understanding this aspect of optimisation.

2.1.3 Problem Ontologies and Frameworks

Perhaps one the of the least well recognised fields of research within optimisation, ontological studies specifically related problems play a key role in providing a framework for organising research efforts as well as providing clear definition of sub-fields and terminology.

Classifications.

A number of problem classifications have been presented in the literature however most have focused on test problems, rather than real-world problems. [23] presents an early classification of test problems, outlining the types of problem by the desired accuracy or tolerance given to the quality of solutions produced by an ideal heuristic. Problems are categorised into optimisation (finding optimal solutions), semi-optimisation problems (finding near-optimal or good solutions), and satisficing (finding qualified solutions).

For continuous test problems, other notable examples include Van Veldhuizen's classification and in-depth analysis of 20th century optimisation test problems [34] to Deb et al.'s early suite of test problems [5] and more recently the review and scalable test problem toolkit presented in [10].

As a result of an extensive literature review, Huband et al. developed a concise yet relatively extensive classification of multi-objective continuous test problems [10]. This classification, although specifically developed for multi-objective continuous problems, provides a firm foundation from which to extend into a classification for all domains. [10] categorise problems by recommendations and features, giving guidance to the construction of new test problems. Recommendations relate to the desirable attributes of any toolkit and features relate to key characteristics of the test functions. However, the classification can also be decomposed into three primary sets of features: parameter-space geometry; objective-space geometry; and function mapping.

[33] and [32] present classifications for the field of dynamic optimisation problems.

Taxonomies and Problem Relations.

Few problem taxonomies have been produced for optimisation problems. More recently, [2] presented an updated numerical taxonomy of single objective optimisation problems with continuous encodings. The study used an experimental approach which collected performance data from optimisation runs on a large set of problems and used distance measures on these performance profiles to draw a cladogram of problems which forms the taxonomy.

Frameworks.

Various optimisation algorithm frameworks have been proposed, such as the well known PISA framework [3]. In contrast few problem specific frameworks have been presented in the literature. An early example is [29], which provides a small framework for problem construction within the wider LANCELOT framework for non-linear optimisation.

More recently, [8] proposed a framework for Real-parameter BBOB which has been the source of workshops and research relating specifically to the task of benchmarking. A framework for selecting and implementing single-objective benchmark problems is provided as part of the larger BBOB framework. A new hyper-heuristic framework was presented in [21], which provided a implementation of an interface between problems, heuristics and hyper-heuristic optimisation methods. Again this framework provided some guidelines for problem design and requirements but as part of a larger framework focused on hyper-heuristic methods.

Summary.

Clearly these frameworks, taxonomies and classifications are useful for their specific domains but none achieve a general ontology which can be used across the set of problem types, including combinatorial, heterogeneous, dynamic and so on. A strong, strictly defined ontology is needed to better understand the relationship between different problems and will also provide a framework with which problem understanding literature can be organised. Furthermore it could provide a well defined, strict nomenclature for use in publications. The confusion between terms like the types of "spaces" (search, parameter, decision, solution, genotypic, phenotypic, fitness, objective, etc.) is just one example of the inconsistency in terms used throughout optimisation literature.

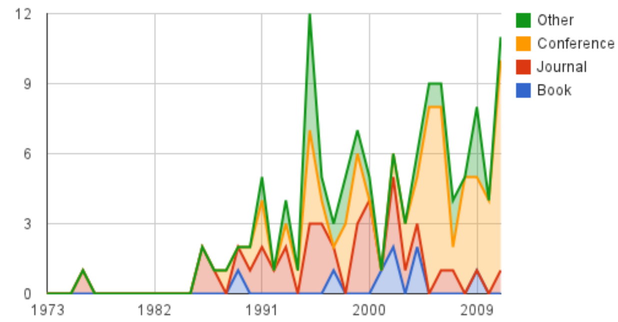

Figure 1: Change in publication type over time. The prevalence of journal publications and books visible at the turn of the century has recently given way to conference publications.

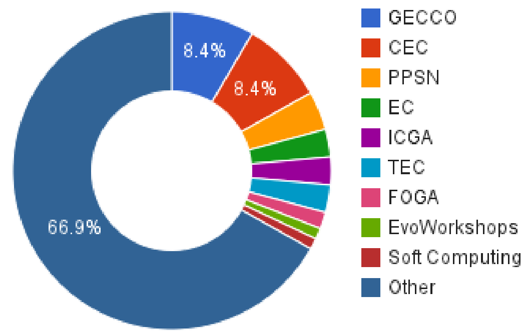

Figure 2: The distribution of publications in various sources. Following the trend for publications in conference proceedings, almost a quarter of the publications surveyed are from evolutionary computing conferences. Two thirds are published as technical reports and theses rather than in more accessible journals and conference proceedings.

2.1.4 Visualisation

Visualisation is a useful approach to problem understanding as it allows us to visually examine and compare problem features. That said, it has received relatively little attention thus far. One case in which visualisation has been used is [1], where Ashlock and Schonfeld use a 2D graph to represent a fitness landscape in which a node is an optima and edge weights indicate the distance between the optima.

3. PUBLICATION ANALYSIS

3.1 Annual Trend

There was an increasing trend of publications relating to problem understanding from the mid 1980's through to 2011 (see Fig 1). At the turn of the century, the majority of research was published in journals or books, following the same maturing pattern as Evolutionary Computation research.

Despite recent publications being mainly conference papers, the number of publications appearing in other media is slowly growing. Indeed, the increased ratio of conference publications to journal articles suggests that research into problem understanding is expanding and that these published works are a prelude to an increase in the number of journal articles devoted to the topic.

3.2 Publication Sources

In recent years, publications relating to problem understanding have been predominantly published as part of conference proceedings at events such as GECCO, CEC and PPSN (see Fig 2). An analysis of the sources of these publications and the other publications produced by the authors indicate that the Evolutionary Computation community has been largely responsible for the literature published to date, although other outlets are occasionally used.

However, although the number of publications is increasing, largely by Evolutionary Computation focused dissemination, 66.9% of the publications were in disparate sources, such as: technical reports and theses; application specific journals like the Journal of Mathematical Chemistry; and smaller conferences and workshops. This wide distribution of publications has lead to a more disjointed research effort in this area with isolated research clusters while also making relevant papers less accessible and difficult to locate.

4. AUTHOR ANALYSIS

Figure 3 shows collaborations within the field of problem understanding, as well as demonstrating how certain authors have been cited. Each individual author included in the study is shown by a node in the graph, and coloured to indicate the number of publications relevant to the field. The figure clearly shows a large corpus of researchers connected by collaborative efforts (black lines) and cross-references to other authors' work (dashed, grey lines). The authors in the large 'research ring' are most commonly established EC researchers. A prominent example is Deb, who as well as being the most frequently published of the surveyed authors has a considerable number of citations in the field. A factor in this is his significant contributions to work on multi-objective test problems (e.g., [5]). There are, however, a significant number of isolated authors and groups of authors whose efforts are not suitably recognised by the wider field nor recognise work produced by the wider community. In many cases this is due to the narrow focus of the work, such as Sartenaer [29]. Some authors, such as Yao, connect research clusters through association.

The lack of cross-referencing and isolation of the smaller groups of authors is partly due to the change in research trends (such as the inclusion of multi-objective and dynamic problems) as well as period over which the survey is conducted where papers become "lost in the pile" and/or authors become inactive. This time period also accounts for the lack of collaboration between individuals and groups, where some authors do not publish in the same time period.

5. A COLLABORATIVE SYSTEM

Within the field of optimisation and its community there are a wide number of proposed frameworks and individual databases of resources and literature. Each of these are managed and curated by a specific individual or group. Whilst this is certainly a significant improvement over individuals reinventing the wheel (or re-implementing a basic EA) there is still a lack of cohesion in their efforts. As we note later, these frameworks are essentially competing to produce the same thing - a framework for optimisation methods.

In this section we propose a wider system, taking inspiration from the biosciences and other similar subjects who recognise and utilise open, accessible and connected databases and frameworks far more effectively. It is argued that such a system would act as an ideal base from which researchers could access resources to conduct experiments and potentially play host to a problem specific framework for optimisation research.

5.1 The Status Quo

There are a wide variety of optimisation frameworks available on the internet. A short and by no means exhaustive search finds the following frameworks: PISA (Any), Opt4J (Java), jMetal (Java), JavaEvA (Java), Watchmaker (Java), ECJ (Java), JCLEC (Java), JGAP (Java), GAA (Java), Paradiseo (C++), Evolving Objects (C++), UOF (C++) and Open Beagle (C++).

Each of these projects take a different approach to implementing an optimisation framework that is often a combination of problem, heuristic, experiment and visualisation frameworks. However, none of these frameworks use a standard interface for communicating between each of its components or one another. Furthermore, none of these frameworks are easily machine accessible. An algorithm in one of these frameworks, for example, cannot be identified through labelled meta-data. As a result, web-services and open access source repositories cannot be effectively established and so researchers are unable to easily access, download and combine elements from any of these frameworks.

In addition to the software frameworks discussed above, the community has not effectively established a central knowledge base, like Wikipedia or perhaps WolframAlpha, to connect and coordinate the various efforts. For example, a large number of academics maintain online bibliographies which are manually managed to ensure they are relevant to the specific focus in question. While the concept of a curated list is desirable, they rarely interact with existing bibliographic databases like Google Scholar or CiteseerX and are scattered across the internet residing in many different institution sites and/or academic homepages.

5.2 Achievable Aims

5.2.1 *News Bulletin / Information Portal*

The primary function of any centralised system (assumedly web based) should be to provide a bulletin of important information, events and recent or forthcoming publications. This can be easily achieved through a combination of a site driven news feed as well as the consumption, filtering and reproduction of existing feeds already available elsewhere. In addition, the system should maintain a list of links to other relevant resources, existing bibliographies and software frameworks.

As well as providing news feed services, the site should link in with social media services such as LinkedIn, Twitter and Google+ to encourage users to connect and collaborate. The system should foster these forums for discussion and provide an automatic broadcast service of discussions tagged to the site (e.g., Twitter posts tagged "#probund").

5.2.2 *Connected Bibliography*

Linked.

The system should maintain a general "Problem Understanding" bibliography which can be contributed to by all users. Each reference should be linked, where possible, to other bibliographic sites such as Google Scholar, CiteseerX, DBLP, etc., as well as the originating site, digital copies of the publication and/or associated resources.

Curated.

The system should provide a bibliography service where users can manage a list of bibliography references, similar to existing bibliographies like Carlos A. Coello Coello's extensive bibliography (`http://www.lania.mx/~ccoello/EMOO/`). The references should feed into a centralised list to remove duplicates and

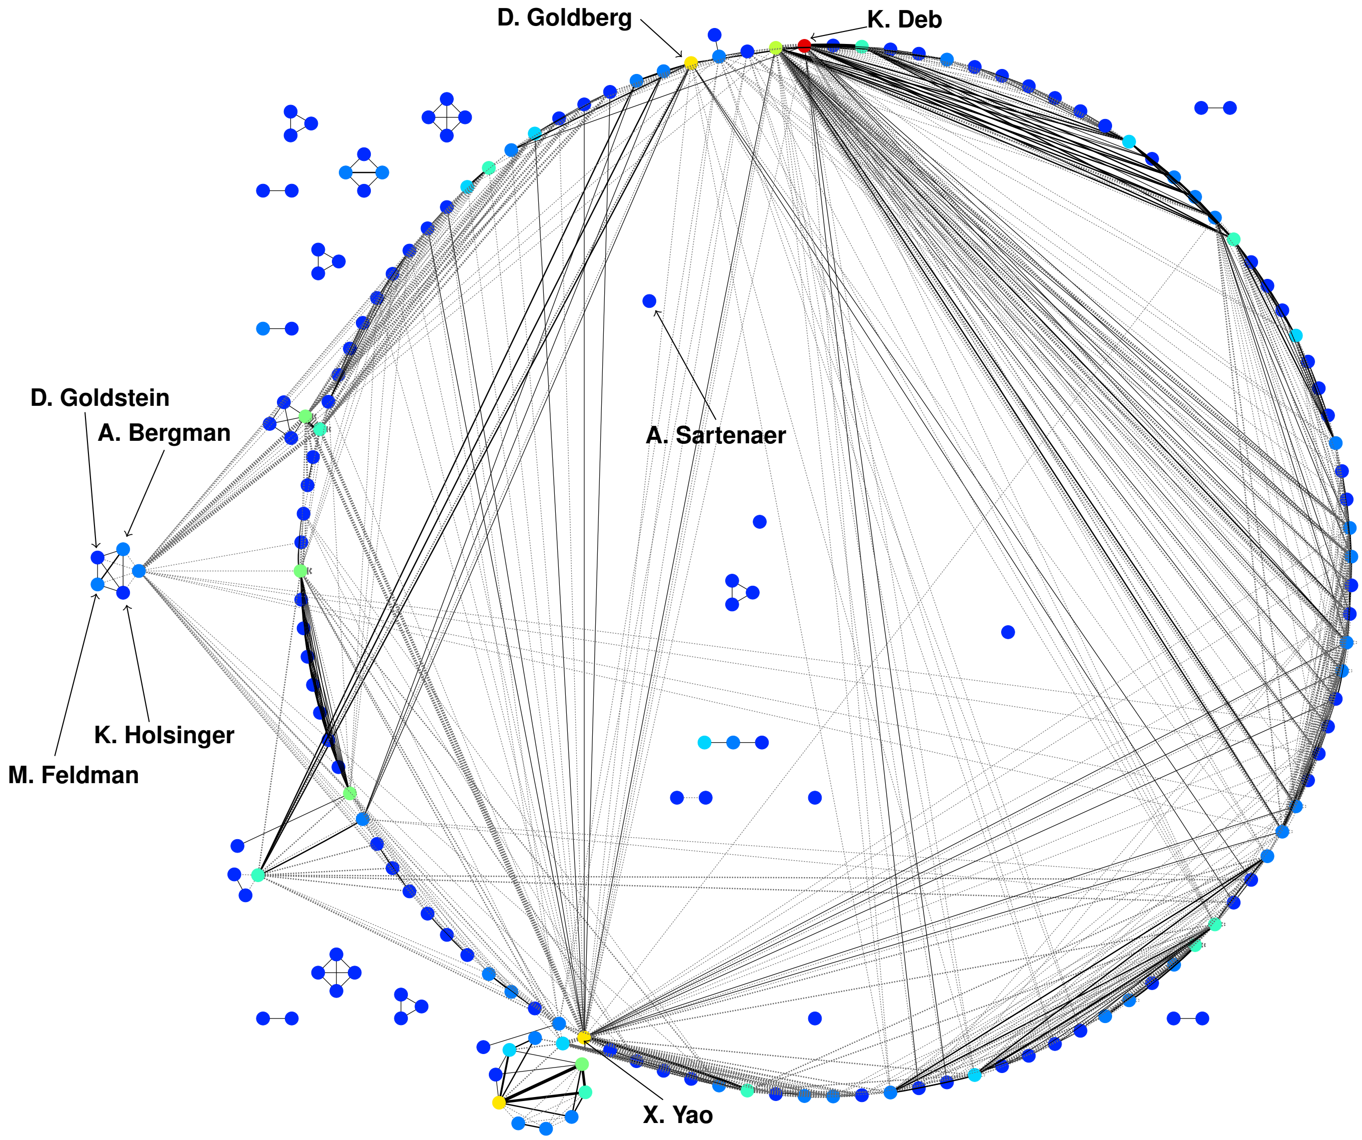

Figure 3: A sociometric graph of authors working on aspects of problem understanding. Each author of the 133 surveyed papers is represented by a node, coloured according to how many publications they contributed to. A solid black edge between authors implies that the two connected authors have collaborated, while a grey dashed line indicates that one author cited another. Clearly, although there is some cooperation, some of the research is still disjoint and isolated.

allow existing references to be easily imported into other curated lists. Furthermore, we suggest that each sub-bibliography be organised by topic or focus (such as Landscape Analysis) rather than by the user(s) curating the list. This will allow for groups to manage the list so that it is more sustainable and persistent.

Annotated.
The system should provide the facility for annotating references with abstracts, summaries and related (but not referenced) publications of interest in addition to other publications it references. Where possible, references should be tagged by theme, topic or keyword and link-back to Wiki entries in which it is mentioned.

5.2.3 Wiki Taxonomy

In addition to providing curated bibliographies, the system should utilise Wiki style content management to drive a community created Taxonomy of research, referencing publications where appropriate as well as providing additional content to aid learning and dissemination of new publications. One key aspect of such a Wiki could be the identification of open research questions and key "hot topics" in recent publications which will help to shape the direction of research in the field and facilitate a more coordinated effort.

5.2.4 Source Repository

A foundation element of the system is a centralised, open source, versioned and easily accessible source code repository. The repository should provide a means of collaboratively building a shared software system which is (1) linked to publications where appropriate and (2) linked to the wiki taxonomy. The source repository should be flexible and be easily rearranged to facilitate an continually updating framework and taxonomy which should reflect recent research developments.

5.2.5 *Data and Results Repository*

In addition to hosting software source code for the problem framework, the system should make data and results repositories available for users to upload and publish results for reference in publications in a uniquely identifiable and persistent manner. The data should be machine readable and organised using defined structures (e.g., XML with DTDs) and accessible through web services for wider use and analysis. In principle, all results data should be made available with results providing a "summary" of the data and analysis in related publications.

5.2.6 *Working Groups*

A final aspect of such a system should be the formation and facilitation of "working groups" tasked with specific responsibilities. These could include, ontological review and updating the system taxonomy, development and incorporation of optimisation algorithms and/or problems, the validation of publication references and other maintenance issues, validation of results, identification and flagging of novel experiments or potential avenues of research. These working groups would act as either maintenance and management or review and analysis, and could lead to an interesting and regular source of publications, such as annual literature reviews, publications promoting new ontologies and so forth.

5.3 Blue Sky and the Cloud

5.3.1 *Cloud Computing*

Assuming a reasonably developed source code repository and with the increasingly easy access to elastic cloud computing services, the system could provide the ideal framework with which to launch an open-source, open-access distributed system on which to conduct experiments and store, analyse and publish results. Such an approach would focus the development of experiments within the system itself and distribute the execution from the central system using well established distributed processing methods which will be common to all experiments. A good example of this type of setup can be found in the Ms PacMan vs. Ghosts competition (`http://www.pacman-vs-ghosts.net/`) which provides a similar service, where the experimental runs are conducted on the server rather than by individual users. This approach would create a stronger bond between the experiments and source repository and encourage the use of the open access data repository to enable persistent experimental results which can be referenced by publications.

5.3.2 *Visualisation*

In addition to providing access to data and running experiments through the system using distributed cloud computing, the system would be well placed to provide a visualisation toolkit specialised for the field. The toolkit could incorporate many basic visualisation methods easily, whilst also adapting and including recent and specialised visualisation methods, such as those proposed in the VizGEC workshops at GECCO. Furthermore, these visualisations of results could be automated and included in experimental runs allowing researchers to easily analyse results without having to download the underlying data sets generated by large runs.

6. CONCLUSION

This paper conducted an initial survey of literature relating to problem understanding. An analysis of the publication trends and sources are conducted which identified a steady growth in the number of publications in the field as well as recognising the importance of the Evolutionary Computation community and their contributions to the field.

However, the survey also highlighted the current disparate and disjointed nature of both the publications and researchers involved in the field. This is partly due to the change in research trends (such as the inclusion of multi-objective and dynamic problems) as well as period over which the survey is conducted where papers become "lost in the pile" and/or authors become inactive.

The literature is discussed by theme and broadly categorised into problem analysis, problem construction, problem ontologies and frameworks, and visualisation. Problem Analysis is comprised of: Landscape Analysis; Theory, Feature Definition, Feature Identification and Neighbourhood Structures; Problem Difficulty; Problem/Heuristic Interaction Dynamics. Problem Construction constitutes research into: Test Problems and Problem Generators; Dynamic Problems; Encodings; Real-world Problems. Problem Ontologies and Frameworks cover: Classifications; Taxonomies and Problem Relations; Frameworks. Finally, visualisation is not broken down further due to limited publications found in the search.

A number of potentially fruitful avenues of research are also identified from the search.

- A detailed survey of this landscape analysis literature would be well placed to identify key methods, features and measurements.
- Devise a more general framework for landscape analysis which provides a more holistic approach to landscape analysis.
- A more detailed analysis of proposed problem features and a study to unify these features and generalise across the whole family of optimisation problems.
- Research into generalised problem difficulty measures independent of optimisation method.
- A detailed survey of problem/heuristic interaction dynamics research.
- Develop a unified test problem framework for single-, multi-objective problems with various encodings.
- Explore the effects of encoding types on optimisation problem difficulty, features and relations.
- Develop a unified ontology of optimisation problem understanding.
- Develop a unified framework for optimisation problem source code.
- Explore new techniques for visualising problems and problem information.

Finally, the paper proposes a new collaborative system which is designed to act as: (1) a managed and linked bibliography of literature; (2) wiki/taxonomy for problem understanding research; (3) a source repository for problems, heuristics and problem understanding methods; (4) data and experimental results repository; and (5) news bulletin for field relevant events and publications. It is argued that such a system would act as an ideal base from which researchers could access resources to conduct experiments and potentially play host to a problem specific framework for optimisation research. Furthermore, with a centralised source repository and the increasingly easy access to elastic cloud computing services, the system would provide the ideal framework with which to launch an open-source, open-access distributed system on which to conduct experiments and store, analyse and publish results.

7. REFERENCES

[1] D. Ashlock and J. Schonfeld. Nonlinear projection for the display of high dimensional distance data. In *IEEE CEC 2005*, pages 2776–2783, 2005.

[2] D. A. Ashlock, K. M. Bryden, S. M. Corns, and Justin Schonfeld. An updated taxonomy of evolutionary computation problems using graph-based evolutionary algorithms. In *IEEE WCCI 2006*, pages 403–410, 2006.

[3] Stefan Bleuler, Marco Laumanns, Lothar Thiele, and Eckart Zitzler. PISA — a platform and programming language independent interface for search algorithms. In *EMO 2003*, pages 494 – 508, 2003.

[4] C.R.Reeves. Landscapes, operators and heuristic search. *Annals of Operational Research*, 86:473–490, 1999.

[5] K. Deb, L. Thiele, M. Laumanns, and E Zitzler. Scalable test problems for evolutionary multi-objective optimization. Technical Report 112, ETH Zurich, Switzerland, 2001.

[6] Kalyanmoy Deb, Ankur Sinha, and Saku Kukkonen. Multi-objective test problems, linkages, and evolutionary methodologies. In *Proceedings of GECCO 2006*, pages 1141–1148, 2006.

[7] M. Farina, K. Deb, and P. Amato. Dynamic multiobjective optimization problems: Test cases, approximation, and applications. In *Evolutionary Multi-Criterion Optimization*, pages 71–71. 2003.

[8] 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.

[9] Mark Hauschild and Martin Pelikan. Advanced neighborhoods and problem difficulty measures. In *Proceedings of GECCO 2011*, 2011.

[10] S. Huband, P. Hingston, L. Barone, and L While. A review of multiobjective test problems and a scalable test problem toolkit. *IEEE TEC*, 10(5):477–506, 2006.

[11] H. Li and Q. Zhang. Multiobjective optimization problems with complicated pareto sets, moea/d and nsga-ii. *IEEE TEC*, 12(2):284–302, 2009.

[12] Guanzhou Lu, Jinlong Li, and Xin Yao. Fitness-probability cloud and a measure of problem difficulty for evolutionary algorithms. In *Proceedings of EVOCOP 2011*, pages 108–117, 2011.

[13] M. Lunacek, D. Whitley, and A. Sutton. The impact of global structure on search. In *PPSNX 2008*, pages 498–507, 2008.

[14] C. MacNish. Benchmarking evolutionary and hybrid algorithms using randomized self-similar landscapes. In *Proceedings of SEAL'06*, pages 361–368, 2006.

[15] Kent McClymont and Ed Keedwell. Benchmark multi-objective optimisation test problems with mixed encodings. In *IEEE CEC 2011*, pages 2131–2138.

[16] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. Exploratory landscape analysis. In *Proceedings of GECCO 2011*, 2011.

[17] P. Merz. Advanced fitness landscape analysis and the performance of memetic algorithms. *Evol. Comput*, 12(3):303–325, 2004.

[18] Z. Michalewicz, K. Deb, and T. Stidsen. Test-case generator for nonlinear continuous parameter optimization techniques. *IEEE TEC*, 4(3), 2000.

[19] Christian L. Müller and Ivo F. Sbalzarini. Global characterization of the cec2005 fitness landscapes using fitness-distance analysis. In *Proceedings of the 2011 international conference on Applications of evolutionary computation*, 2011.

[20] C.L. Muller and I.F. Sbalzarini. A tunable real-world multi-funnel benchmark problem for evolutionary optimization - and why parallel island models might remedy the failure of cma-es on it. In *Proceedings of IJCCI 2009*, pages 248–253, 2009.

[21] G. Ochoa, M. Hyde, T. Curtois, J.A. Vazquez-Rodriguez, J. Walker, M. Gendreau, G. Kendall, B. McCollum, A.J. Parkes, S. Petrovic, and E.K. Burke. HyFlex: A Benchmark Framework for Cross-domain Heuristic Search. pages 136–147, 2012.

[22] Tatsuya Okabe. Evolutionary multiobjective optimisation with a hybrid representation. In *IEEE CEC 2003*, pages 2262–2269, 2003.

[23] J. Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, 1983.

[24] Martin Pelikan. Nk landscapes, problem difficulty, and hybrid evolutionary algorithms. In *Proceedings of GECCO 2010*, 2010.

[25] Martin Pelikan, K. Sastry, M. V. Butz, and D. E. Goldberg. Performance of evolutionary algorithms on random decomposable problems. In *PPSN 2006*, pages 788–797, 2006.

[26] Tinos R and Yang S. Continuous dynamic problem generators for evolutionary algorithms. In *IEEE CEC 2007*, pages 236–243, 2007.

[27] Philipp Rohlfshagen and Xin Yao. Dynamic combinatorial optimisation problems: an analysis of the subset sum problem. *Soft Computing - A Fusion of Foundations, Methodologies and Applications*, 15:1723–1734, 2011.

[28] Morrison RW and DeJong KA. A test problem generator for non-stationary environments. In *IEEE CEC 1999*, pages 2047–2053, 1999.

[29] A. Sartenaer. Large-scale nonlinear optimization and the lancelot package. Technical Report 95/9, October 1995.

[30] A. M. Sutton, L. D. Whitley, and A. E. Howe. A polynomial time computation of the exact correlation structure of k-satisfiability landscapes. In *Proceedings of GECCO-2009*, pages 365–372, 2009.

[31] K. Tang, X. Li, N. Suganthan, P, Z. Yang, and T. Weise. Benchmark functions for the cec'2010 special session and competition on large-scale global optimization. Technical report, University of Science and Technology of China, 2009.

[32] Alexandru-Adrian Tantar, Emilia Tantar, and Pascal Bouvry. A classification of dynamic multi-objective optimization problems. In *GECCO 2011*, 2011.

[33] Emilia Tantar, Alexandru-Adrian Tantar, and Pascal Bouvry. On dynamic multi-objective optimization – classification and performance measures. In *IEEE CEC 2011*, 2011.

[34] D. A Van Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Air Force Institute of Technology, Wright-Patterson AFB, Ohio, 1999.

[35] Sebastien Verel, Arnaud Liefooghe, and Clarisse Dhaenens. Set-based multiobjective fitness landscapes: a preliminary study. In *Proceedings of GECCO 2011*, 2011.

[36] Abdunnaser Younes, Paul Calamai, and Otman Basir. Generalized benchmark generation for dynamic combinatorial problems. In *Proceedings of GECCO 2005*, pages 25–31, 2005.