Benchmarking in Optimization:
Best Practice and Open Issues

Thomas Bartz-Beielstein[1], Carola Doerr[2], Daan van den Berg[3], Jakob Bossek[6],
Sowmya Chandrasekaran[1], Tome Eftimov[5], Andreas Fischbach[1], Pascal Kerschke[6],
William La Cava[9], Manuel López-Ibáñez[7], Katherine M. Malan[8], Jason H. Moore[9],
Boris Naujoks[1], Patryk Orzechowski[9,10], Vanessa Volz[11], Markus Wagner[4], and
Thomas Weise[12]

[1]*Institute for Data Science, Engineering, and Analytics, TH Köln, Germany*
[2]*Sorbonne Université, CNRS, LIP6, Paris, France*
[3]*Yamasan Science & Education*
[4]*Optimisation and Logistics, School of Computer Science, The University of Adelaide, Adelaide, Australia*
[5]*Computer Systems Department, Jožef Stefan Institute, Ljubljana, Slovenia*
[6]*Statistics and Optimization Group, University of Münster, Münster, Germany*
[7]*School of Computer Science and Engineering, University of Málaga, Málaga, Spain*
[8]*Department of Decision Sciences, University of South Africa, Pretoria, South Africa*
[9]*Institute for Biomedical Informatics, University of Pennsylvania, Philadelphia, PA, USA*
[10]*Department of Automatics, AGH University of Science and Technology, Krakow, Poland*
[11]*modl.ai, Copenhagen, Denmark*
[12]*Institute of Applied Optimization, School of Artificial Intelligence and Big Data, Hefei University, Hefei, China*
benchmarkingbestpractice@gmail.com

December 18, 2020

Version 2

Abstract

This survey compiles ideas and recommendations from more than a dozen researchers with different backgrounds and from different institutes around the world. Promoting best practice in benchmarking is its main goal. The article discusses eight essential topics in benchmarking: clearly stated goals, well-specified problems, suitable algorithms, adequate performance measures, thoughtful analysis, effective and efficient designs, comprehensible presentations, and guaranteed reproducibility. The final goal is to provide well-accepted guidelines (rules) that might be useful for authors and reviewers. As benchmarking in optimization is an active and evolving field of research this manuscript is meant to co-evolve over time by means of periodic updates.

arXiv:2007.03488v2 [cs.NE] 16 Dec 2020

1

Contents

1 Introduction

Introducing a new algorithm without testing it on a set of benchmark functions appears to be very strange to every optimization practitioner, unless there is a strong theoretical motivation justifying the interest in the algorithm. Taking theory-focused papers aside, from the very beginning in the 1960s nearly every publication in Evolutionary Computation (EC) was accompanied by benchmarking studies. One of the key promoters of the EC research domain, Hans-Paul Schwefel [1975], wrote in his PhD thesis:

> The extremely large and constantly increasing number of optimization methods inevitably leads to the question of the best strategy. There does not seem to be a clear answer. Because, if there were an optimal optimization process, all other methods would be superfluous . . . [1]

Famous studies, e.g., from Moré et al. [1981], were performed in this period and established test functions that are today well known among algorithm developers. Some of them can still be found in the portfolio of recent benchmark studies, e.g., Rosenbrock's function [Rosenbrock, 1960]. In the 1960s, experiments could be rerun only a very limited number of times, using different starting points or random seeds. This situation has changed drastically: nowadays, new algorithms can be run a hundred or even a thousand times. This enables very complex and sophisticated benchmark suites such as those available in the Comparing Continuous Optimizers (COCO) [Hansen et al., 2016b] platform or in Nevergrad [Rapin and Teytaud, 2018]. However, the questions to be answered by benchmarking remain basically the same, e.g.,

- how well does a certain algorithm perform on a given problem?

- why does an algorithm succeed/fail on a specific test problem?

Specifying the goal of a benchmark study is as important as the study itself, as it shapes the experimental setup – i.e., the choice of problem instances, of the algorithm instances, the performance criteria, and the statistics. Typical goals that a user or a researcher wishes to answer through a benchmarking study are discussed in Section 2.

But not only computational power has increased significantly in the last decades. Theory made important progress as well. In the 1980s, some researchers claimed that there is an algorithm that is able to outperform all other algorithms on average [Goldberg, 1989]. A set of no free lunch theorems (NFLTs), presented by Wolpert and Macready [1997] changed this situation [Adam et al., 2019]. Statements about the performance of algorithms should be coupled with the problem class or even the problem instances. Brownlee [2007] summarizes NFLT consequences and gives the following recommendations:

> 1) bound claims of algorithm or parameter suitability to the problem instances being tested,
> 2) research into devising problem classes and matching suitable algorithms to classes is a good thing,
> 3) be cautious about generalizing performance to other problem instances, and
> 4) be very cautious about generalizing performance to other problem classes or domains.

Haftka [2016] describes NFLT consequences as follows:

> Improving an algorithm for one class of problem is likely to make it perform more poorly for other problems.

Some authors claim that this statement is too general and should be detailed as follows: improving the performance of an algorithm, e.g., via parameter tuning, for a subset of problems may make it perform worse for a different subset. This does not work so well for classes of problems, unless the classes are finite and small. It also does not work for any two arbitrary subsets, since the subsets may be correlated in precisely the way that leads to better performance of the algorithm. A number of works discuss limitations

---

[1]German original quote: "Die überaus große und ständig steigende Zahl von Optimierungsmethoden führt zwangsläufig zu der Frage nach der besten Strategie. Eine eindeutige Antwort scheint es nicht zu geben. Denn, gäbe es ein optimales Optimierungsverfahren, dann würden sich alle anderen Methoden erübrigen. . ."

for the consequences and the impact of NFLT, such as García-Martínez et al. [2012] and McDermott [2020]. For example, Culberson [1998] stated: "In the context of search problems, the NFL theorem strictly only applies if arbitrary search landscapes are considered, while the instances of basically any search problem of interest have compact descriptions and therefore cannot give rise to arbitrary search landscapes ".

Without doubt, NFLT has changed the way how benchmarking is *considered* in EC. Problems caused by NFLT are still subject of current research, e.g., Liu et al. [2019] discuss paradoxes in numerical comparison of optimization algorithms based on NFLT. Whitley et al. [2002] examine the meaning and significance of benchmarks in light of theoretical results such as NFLT.

Independently of the ongoing NFLT discussion, benchmarking gains a central role in current research, both for theory and practice. Three main aspects that need to be addressed in every benchmark study are the choice of

1. the performance measures,

2. the problem (instances), and

3. the algorithm (instances).

Excellent papers on how to set up a good benchmark test exist for many years. Hooker and Johnson are only two authors that published papers still worth reading today [Hooker, 1994, 1996, Johnson et al., 1989, 1991, Johnson, 2002b]. McGeoch [1986] can be considered as a milestone in the field of experimental algorithmics, which builds the cornerstone for benchmark studies. Gent and Walsh [1994] stated that the empirical study of algorithms is a relatively immature field – and we claim that this situation has unfortunately not significantly changed in the last 25 years. Reasons for this unsatisfactory situation in EC are manifold. For example, EC has not agreed upon general methodology for performing benchmark studies like the fields of statistical Design of Experiments (DOE) or data mining [Chapman et al., 2000, Montgomery, 2017]. These fields provide a general methodology to encourage the practitioner to consider important issues before performing a study. Some journals provide explicit minimal standard requirements.[2]

The question remains: why are minimum standards not considered in every paper submitted to EC conferences and journals? Or, formulated alternatively: why have such best practices not become minimum required standards? One answer might be: setting up a sound benchmark study is very complicated. There are many pitfalls, especially stemming from complex statistical considerations [Črepinšek et al., 2014]. So, to do nothing wrong, practitioners oftentimes report only average values decorated with corresponding standard deviations, p-values, or boxplots. Another answer might be: practical guidelines are missing. Researchers from computer science would apply these guidelines if examples were available. This paper is a joint initiative from several researchers in EC. It presents best-practice examples with references to relevant publications and discusses open issues. This joint initiative was established during the Dagstuhl seminar 19431 on *Theory of Randomized Optimization Heuristics*, which took place in October 2019. Since then, we have been compiling ideas covering a broad range of disciplines, all connected to EC.

We are aware that every version of this paper represents a snapshot, because the field is evolving. New theoretical results such as no-free lunch might come up from theory and new algorithms (quantum computing, heuristics supported by deep learning techniques, etc.) appear on the horizon, and new measures, e.g., based on extensive resampling (Monte Carlo), can be developed in statistics.

We consider this paper as a starting point, as a first trial to support the EC community in improving the quality of benchmark studies. Surely, this paper cannot cover every single aspect related to benchmarking. Although this paper mainly focuses on single-objective, unconstrained optimization problems, its findings can be easily transferred to other domains, e.g, multi-objective or constrained optimization. The objectives in other problem domains may differ slightly and may require different performance measures – but the

---

[2]See https://www.springer.com/cda/content/document/cda_downloaddocument/Journal+of+Heuristic+Policies+on+Heuristic+Search.pdf?SGWID=0-0-45-1483502-p35487524 for guidelines of the *Journal of Heuristics* and https://static.springer.com/sgw/documents/1593723/application/pdf/Additional_submission_instructions.pdf for similar ones of the journal *Swarm Intelligence*.

content of most sections should be applicable. Each of the following sections presents references to best-practice examples and discusses open topics. The following aspects, which are considered relevant to every benchmark study, are covered in the subsequent sections:

1. Goals: what are the reasons for performing benchmark studies (Section 2)?

2. Problems: how to select suitable problems (Section 3)?

3. Algorithms: how to select a portfolio of algorithms to be included in the benchmark study (Section 4)?

4. Performance: how to measure performance (Section 5)?

5. Analysis: how to evaluate results (Section 6)?

6. Design: how to set up a study, e.g., how many runs shall be performed (Section 7)?

7. Presentation: how to describe results (Section 8)?

8. Reproducibility: how to guarantee scientifically sound results and how to guarantee a lasting impact, e.g., in terms of comparability (Section 9)?

The paper closes with a summary and an outlook in Section 10.

Generalization of benchmarking results. As discussed above in the context of the NFLT, we recommend being very precise in the description of the algorithms and the problem instances that were used in the benchmark study. Performance extrapolation or generalization always needs to be flagged as such, and where algorithms are compared to each other, it should be made very clear what the basis for the comparison is. We suggest to very carefully distinguish between algorithms (e.g., "the" Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [Hansen, 2000]) and algorithm instances (e.g., the pycma-es [Hansen et al., 2020] with population size 8, budget 100, restart strategy X, etc.).[3] A similar rule applies to the problems (e.g., "the" sphere function) vs. a concrete problem instance (the five-dimensional sphere function $f : \mathbb{R}^5 \to \mathbb{R}, x \mapsto \alpha \sum_{i=1}^{d} x_i^2 + \beta$ centered at at the origin, multiplicative scaling α, and additive shift β). Going one step further, one may even argue that we only benchmark a certain *implementation* of an algorithm instance, which is subject to a concrete choice of implementation language, compiler and operating system optimizations, and concrete versions of software libraries.

Algorithm and problem instances may be (and in the context of this survey often are) randomized, so that the performance of the algorithm instance on a given problem instance is a series of (typically highly correlated) random variables, one for each step of the algorithm. In practice, replicability is often achieved by fixing the random number generator and storing the random seed, which plays an important role in guaranteeing reproducibility as discussed in Sec. 9.

2   Goals of Benchmarking Activities

The motivations for performing benchmark studies on optimization algorithms are as diverse as the algorithms and the problems that are being used in these studies. Apart from scientifically motivated goals, benchmarking can also be used as a means to popularize an algorithmic approach or a particular problem. In this section, we focus on summarizing the most common scientifically-motivated goals for benchmarking studies. Figure 1 summarizes these goals. The relevance of these goals can differ from study to study, and the proposed categorization is not necessarily unique, but should be understood as an attempt to find something that represents the benchmarking objectives well within the broader scientific community. Most of the goals listed below are, ultimately, aimed at contributing towards a better deployment of the algorithms in practice, typically through a better understanding of the interplay between the algorithmic design choices

---

[3]Algorithm instances are also referred to as "algorithm configurations" in the literature [Birattari et al., 2002].

Common Goals of Benchmarking Studies

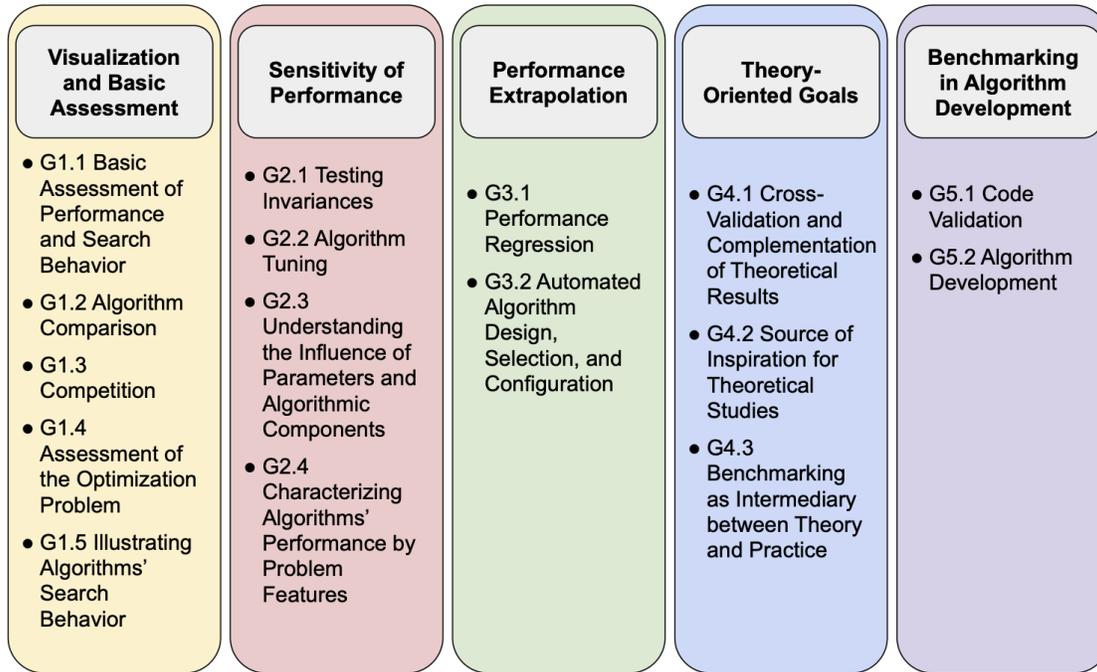| Visualization and Basic Assessment	Sensitivity of Performance	Performance Extrapolation	Theory-Oriented Goals	Benchmarking in Algorithm Development
• G1.1 Basic Assessment of Performance and Search Behavior • G1.2 Algorithm Comparison • G1.3 Competition • G1.4 Assessment of the Optimization Problem • G1.5 Illustrating Algorithms' Search Behavior	• G2.1 Testing Invariances • G2.2 Algorithm Tuning • G2.3 Understanding the Influence of Parameters and Algorithmic Components • G2.4 Characterizing Algorithms' Performance by Problem Features	• G3.1 Performance Regression • G3.2 Automated Algorithm Design, Selection, and Configuration	• G4.1 Cross-Validation and Complementation of Theoretical Results • G4.2 Source of Inspiration for Theoretical Studies • G4.3 Benchmarking as Intermediary between Theory and Practice	• G5.1 Code Validation • G5.2 Algorithm Development

Figure 1: Summary of common goals of benchmark studies.

and the problem instance characteristics. However, benchmarking also plays an important role as intermediary between the scientific community and users of optimization heuristics and as intermediary between theoretically and empirically-guided streams within the research community.

2.1 Visualization and Basic Assessment of Algorithms and Problems

(G1.1) *Basic Assessment of Performance and Search Behavior.*
The arguably most basic research question that one may want to answer with a benchmark study is how well a certain algorithm performs on a given problem. In the absence of mathematical analyses, and in the absence of existing data, the most basic approach to gain insight into the performance is to run one or more instances of the algorithm (ideally several times, if the algorithm or the problem are stochastic) on one or more problem instances, and to observe the behavior of the algorithm. With this data, one can analyze what a typical performance profile looks like on some problem instance, how the solution quality evolves over time, how robust the performance is, etc. The evaluation criteria can be diverse, as we shall discuss in Section 5. But what is inherent to all studies falling into this goal G1.1, is that they are aimed to answer a rather basic question "How well does this particular algorithm perform on this particular problem instance?" or "How does a particular run of this algorithm on the given problem look like?".

(G1.2) *Algorithm Comparison.*
The great majority of benchmark studies do not focus on a single algorithm, but rather *compare* the performance and/or the search behavior of two or more algorithms. The comparison of algorithms

serves, most notably, the purpose of understanding strengths and weaknesses of different algorithmic approaches for different types of problems or problem instances during the different stages of the optimization process. These insights can be leveraged to design or to select, for a given problem class or instance, a most suitable algorithm instance.

(G1.3) *Competition.*
One particular motivation to compare algorithms is to determine a "winner", i.e., an algorithm that performs better than any of its competitors, for a given performance measure and on a given set of problem instances. Benchmarking is of great value in selecting the most adequate algorithm especially in real-world optimization settings [Beiranvand et al., 2017]. The role of competitive studies for benchmarking is discussed quite controversially [Hooker, 1996], as competitive studies may promote algorithms that overstate the importance of the problems that they are tested upon, and thereby create over-fitting. At the same time, however, one cannot neglect that competitions can provide an important incentive to contribute to the development of new algorithmic ideas and better selection guidelines.

(G1.4) *Assessment of the Optimization Problem.* In many real-world problems like scheduling, container packing, chemical plant control, or protein folding, the global optimum is unknown, while in other problems it is necessary to deal with limited knowledge, or lack of explicit formulas. In those situations, computer simulations or even physical experiments are required to evaluate the quality of a given solution candidate. In addition, even if a problem is explicitly modelled by a mathematical formula, it can nevertheless be difficult to grasp its structure or to derive a good intuition for what its fitness landscape looks like. Similarly, when problems consist of several instances, it can be difficult to understand in what respect these different instances are alike and in which aspects they differ. Benchmarking simple optimization heuristics can help to analyze and to visualize the optimization problem and to gain knowledge about its characteristics.

(G1.5) *Illustrating Algorithms' Search Behavior.*
Understanding how an optimization heuristic operates on a problem can be difficult to grasp when only looking at the algorithm and problem description. One of the most basic goals that benchmarking has to offer are numerical and graphical illustrations of the optimization process. With these numbers and visualizations, a first idea about the optimization process can be derived. This also includes an assessment of the stochasticity when considering several runs of a randomized algorithm or an algorithm operating upon a stochastic problem. In the same vein, benchmarking offers a hands-on way of visualizing effects that are difficult to grasp from mathematical descriptions. That is, where mathematical expressions are not easily accessible to everyone, benchmarking can be used to illustrate the effects that the mathematical expressions describe.

2.2 Sensitivity of Performance in Algorithm Design and Problem Characteristics

(G2.1) *Testing Invariances.*
Several researchers argue that, ideally, the performance of an optimization algorithm should be invariant with respect to certain aspects of the problem embedding, such as the scaling and translation of the function values [Vrielink and van den Berg, 2019], dimensional increase [De Jonge and van den Berg, 2020], or a rotation of the search space (see Hansen [2000] and references therein for a general discussion and Lehre and Witt [2012], Rowe and Vose [2011] for examples formalizing the notion of *unbiased* algorithms).

Whereas certain invariances, such as comparison-baseness, are typically easily inferred from a pseudocode description of the algorithm, other invariances (e.g., invariance with respect to translation or rotation) might be harder to grasp. In such cases, benchmarking can be used to test, empirically, whether the algorithm possesses the desired invariances.

(G2.2) *Algorithm Tuning.*

Most optimization heuristics are configurable, i.e, we are able to adjust their search behavior (and, hence, performance) by modifying their parameters. Typical parameters of algorithms are the number of individuals kept in the memory (its 'population size'), the number of individuals that are evaluated in each iteration, parameters determining the distribution from which new samples are generated (e.g., the mean, variance, and direction of the search), the selection of survivors for the next generation's population, and the stopping criterion. Optimization heuristics applied in practice often comprise tens of parameters that need to be tuned.

Finding the optimal configuration of an algorithm for a given problem instance is referred to as offline parameter tuning [Eiben and Jelasity, 2002, Eiben and Smith, 2015]. Tuning can be done manually or with the help of automated configuration tools [Akiba et al., 2019, Bergstra et al., 2013, Olson and Moore, 2016], Benchmarking is a core ingredient of the parameter tuning process. A proper design of experiment is an essential requirement for tuning studies [Bartz-Beielstein, 2006, Orzechowski et al., 2018, 2020]. Parameter tuning is a necessary step before comparing a viable configuration of a method with others, as we are disregarding those combinations of parameters, which do not yield promising results.

Benchmarking can help to shed light on suitable choices of parameters and algorithmic modules. Selecting a proper parameterization for a given optimization problem is a tedious task [Fialho et al., 2010]. Besides the selection of the algorithm and the problem instance, tuning requires the specification of a *performance measure*, e.g., best solution found after a pre-specified number of function evaluations (to be discussed in Sec. 5) and a *statistic*, i.e., number of repeats, which will be discussed in Sec. 7.

Another important concern with respect to algorithm tuning is the robustness of the performance with respect to these parameters, i.e., how much does the performance deteriorate if the parameters are mildly changed? In this respect, parameter recommendations with a better robustness might be preferable over less robust ones, even if compromising on performance [Paenke et al., 2006].

(G2.3) *Understanding the Influence of Parameters and Algorithmic Components.*

While algorithm tuning focuses on finding the best configuration for a given problem, *understanding* refers to the question: *why* does one algorithm perform better than a competing one? Understanding requires additional statistical tools, e.g., analysis of variance or regression techniques. Questions such as "Does recombination have a significant effect on the performance?" are considered in this approach. Several tools that combine methods from statistics and visualization are integrated in the software package Sequential Parameter Optimization Toolbox (SPOT), which was designed for understanding the behavior of optimization algorithms. SPOT provides a set of tools for model based optimization and tuning of algorithms. It includes surrogate models, optimizers and DOE approaches [Bartz-Beielstein et al., 2017].

(G2.4) *Characterizing Algorithms' Performance by Problem (Instance) Features and Vice Versa.*

Whereas *understanding* as discussed in the previous paragraph tries to get a deep insight into the elements and working principles of algorithms, *characterization* refers to the relationship between algorithms and problems. That is, the goal is to link features of the problem with the performance of the algorithm(s). A classical example for a question answered by the characterization approach is how the performance of an algorithm scales with the number of decision variables.

Problem instance features can be high-level features such as its dimensionality, its search constraints, its search space structure, and other basic properties of the problem. Low-level features of the problem, such as its multi-modality, its separability, or its ruggedness can either be derived from the problem formulation or via an exploratory sampling approach [Kerschke and Trautmann, 2019a,b, Malan and Engelbrecht, 2013, Mersmann et al., 2010, 2011, Muñoz Acosta et al., 2015a,b].

2.3 Benchmarking as Training: Performance Extrapolation

(G3.1) *Performance Regression.*
The probably most classical hope associated with benchmarking is that the generated data can be used to extrapolate the performance of an algorithm for other, not yet tested problem instances. This extrapolation is highly relevant for *selecting* which algorithm to choose and how to *configure* it, as we shall discuss in the next section. Performance extrapolation requires a good understanding of how the performance depends on problem characteristics, the goal described in G2.4.

In the context of machine learning, performance extrapolation is also referred to as *transfer learning* [Pan and Yang, 2010]. It can be done manually or via sophisticated regression techniques. Regardless of the methodology used to extrapolate performance data, an important aspect in this regression task is a proper selection of the instances on which the algorithms/configurations are tested. For performance extrapolation based on supervised learning approaches, a suitable selection of feature extraction methods is another crucial requirement for a good fit between extrapolated and true performance.

(G3.2) *Automated Algorithm Design, Selection, and Configuration.*
When the dependency of algorithms' performance with respect to relevant problem characteristics is known and performance can be reasonably well extrapolated to previously unseen problem instances, the benchmarking results can be used for designing, selecting, or configuring an algorithm for the problem at hand. That is, the goal of the benchmark study is to provide training data from which rules can be derived that help the user choose the best algorithm for her optimization task. These guidelines can be human-interpretable such as proposed in Bartz-Beielstein [2006], Liu et al. [2020] or they can be implicitly derived by AutoML techniques [Hutter et al., 2019, Kerschke and Trautmann, 2019a, Kerschke et al., 2019, Olson and Moore, 2016].

2.4 Theory-Oriented Goals

(G4.1) *Cross-Validation and Complementation of Theoretical Results.*
Theoretical results in the context of optimization are often expressed in terms of asymptotic running time bounds [Auger and Doerr, 2011, Doerr and Neumann, 2020, Neumann and Witt, 2010], so that it is typically not possible to derive concrete performance values from them, e.g., for a concrete dimension, target values, etc. To analyze the behavior in small dimensions and/or to extend the regime for which the theoretical bounds are valid, a benchmarking study can be used to complement existing theoretical results.

(G4.2) *Source of Inspiration for Theoretical Studies.*
Notably, empirical results derived from benchmarking studies are an important source of inspiration for theoretical works. In particular when empirical performance does not match our intuition, or when we observe effects that are not well understood by mathematical means, benchmarking studies can be used to pinpoint these effects, and to make them accessible to theoretical studies, see [Doerr et al., 2019] for an example.

(G4.3) *Benchmarking as Intermediary between Theory and Practice.*
The last two goals, G4.1 and G4.2, together with G1.1 and G1.2 highlight the role of benchmarking as an important intermediary between empirically-oriented and mathematically-oriented sub-communities within the domain of heuristic optimization [Müller-Hannemann and Schirra, 2010]. In this sense, benchmarking plays a similar role for optimization heuristics as *Algorithm Engineering* [Kliemann and Sanders, 2016] does for classical algorithmics.

2.5 Benchmarking in Algorithm Development

(G5.1) *Source Code Validation.*
Another important aspect of benchmarking is that it can be used to verify that a given program performs as it is expected to. To this end, algorithms can be assessed on problem instances with known properties. If the algorithm consistently does not behave as expected, a source code review might be necessary.

(G5.2) *Algorithm Development.*
In addition to *understanding* performances, benchmarking is also used to identify weak spots with the goal to develop better performing algorithms. This also includes first empirical comparisons of new ideas to gain first insights into whether or not it is worth investigating further. This can result in a loop of empirical and theoretical analysis. A good example for this is parameter control: it has been observed early on that a dynamic choice of algorithms' parameters can be beneficial over static ones [Karafotias et al., 2015]. This led to the above mentioned loop of evaluating parameters empirically and stimulated theoretical investigations.

2.6 Open Issues and Challenges

Several of the goals listed above require fine-grained records about the traces of an algorithm, raising the issue of storing, sharing, and re-using the data from the benchmark studies. Several benchmark environments offer a data repository to allow users to re-use previous experimental results. However, compatibility between the data formats of different platforms is rather weak, and a commonly agreed-upon standard would be highly desirable, both for a better comparability and for a resource-aware benchmarking culture. As long as such standards do not exist, tools that can flexibly interpret different data formats can be used. For example, the performance assessment module IOHanalyzer of the IOHprofiler benchmarking environment [Doerr et al., 2018] can deal with various different formats, including those from the two most widely adopted benchmarking environments in EC, Nevergrad [Rapin and Teytaud, 2018] and COCO [Hansen et al., 2016b].

Coming back to a resource-aware benchmarking culture, we repeat a statement already made in the introduction: two of the most important steps of a benchmarking study are the formulation of a clear research question that shall be answered, and the design of an experimental setup that best answers this question through a well-defined set of experiments. It is often surprising to see how many scientific reports do not clearly explain the main research question that the study aims to answer, (n)or how the reported benchmarking data supports the main claims.

Finally, we note that also the goals themselves undergo certain "trends", which are not necessarily stable over time. The above collection of goals should therefore be seen as a snapshot of what we observe today, some of the goals mentioned above may gain or lose in relevance.

3 Problem Instances

A critical element of algorithm benchmarking is the choice of problem instances, because it can heavily influence the results of the benchmarking. Assuming that we (ultimately) aim at solving real-world problems, ideally, the problem set should be representative of the real-world scenario under investigation, otherwise it is not possible to derive general conclusions from the results of the benchmarking. In addition, it is important that problem sets are continually updated to prevent the over-tuning of algorithms to particular problem sets.

This section discusses various aspects related to problem sets used in benchmarking. The four questions we address are:

1. What are the desirable properties of a good problem set?

2. How to evaluate the quality of a problem set?

3. What benchmark problem sets are publicly available?

4. What are the open problems in research related to problem sets for benchmarking?

3.1 Desirable Characteristics of a Problem Set

This section describes some of the general properties that affect the usefulness of suites of problems for benchmarking, see Whitley et al. [1996] and Shir et al. [2018] for position statements.

(B1.1) *Diverse.*
A good benchmark suite should contain problems with a range of difficulties [Olson et al., 2017]. However, what is difficult for one algorithm could be easy for another algorithm and for that reason, it is desirable for the suite to contain a wide variety of problems with different characteristics. In this way, a good problem suite can be used to highlight the strengths and weaknesses of different algorithms. Competition benchmark problems are frequently distinguished based on a few simple characteristics such as modality and separability, but there are many other properties that can affect the difficulty of problems for search [Kerschke and Trautmann, 2019b, Malan and Engelbrecht, 2013, Muñoz Acosta et al., 2015b] and the instances in a problem suite should collectively capture a wide range of characteristics.

(B1.2) *Representative.*
At the end of a benchmarking exercise, claims are usually made regarding algorithm performance. The more representative the benchmarking suite is of the class of problems under investigation, the stronger the claim about algorithm performance will hold. The problem instances should therefore include the difficulties that are typical of real world instances of the problem class under investigation.

(B1.3) *Scalable and tunable.*
Ideally a benchmark set/framework includes the ability to tune the characteristics of the problem instances. For example, it could be useful to be able to set the dimension of the problem, the level of dependence between variables, the number of objectives, and so on.

(B1.4) *Known solutions / best performance.*
If the optimal solution(s) of a benchmark problem are known, then it makes it easier to measure exact performance of algorithms in relation to the known optimal performance. There are, however, simple problems for which optimal solutions are not known even for relatively small dimensions (e.g. the Low Auto-correlation Binary Sequence (LABS) problem [Packebusch and Mertens, 2016]). In these cases it is desirable to have the best known performance published for particular instances.

3.2 Evaluating the Quality of a Problem Set

Although it is trivial to assess whether a problem suite provides information on the optimal solution or is tunable, it is not as obvious to assess whether a problem set is diverse or representative. In this section, we provide a brief overview of existing ways of evaluating the quality of problem sets.

(B2.1) *Feature space.* One of the ways of assessing the diversity of a set of problem instances is to consider how well the instances cover a range of different problem characteristics. When these characteristics are measurable in some way, then we can talk about the instances covering a wide range of feature values. Garden and Engelbrecht [2014] use a self-organizing feature map to cluster and analyse the Black-Box-Optimization-Benchmarking (BBOB) and Congress on Evolutionary Computation (CEC) problem sets based on fitness landscape features (such as ruggedness and the presence of multiple funnels). In a similar vein, Škvorc et al. [2020] use Exploratory Landscape Analysis (ELA) features [Mersmann et al., 2011] combined with clustering and a *t*-distributed stochastic neighbor embedding visualization approach to analyse the distribution of problem instances across feature space.

(B2.2) *Performance space.*

Simple statistics such as mean and best performance aggregate much information without always enabling the discrimination of two or more algorithms. For example, two algorithms can be very similar (and thus perform comparably) or they might be structurally very different but the aggregated scores might still be comparable. From the area of algorithm portfolios, we can employ ranking-based concepts such as the marginal contribution of an individual algorithm to the total portfolio, as well as the Shapley values, which consider all possible portfolio configurations [Fréchette et al., 2016]. Still, for the purpose of benchmarking and better understanding of the effect of design decisions on an algorithm's performance, it might be desirable to focus more on instances that enable the user to tell the algorithms apart in the performance space.

This is where the targeted creation of instances comes into play. Among the first articles that evolved small Traveling Salesperson Problem (TSP) instances that are difficult or easy for a single algorithm is that by Mersmann et al. [2013], which was then followed by a number of articles also in the continuous domain as well as for constrained problems. Recently, this was extended to the explicit discrimination of pairs of algorithms for larger TSP instances [Bossek et al., 2019], which required more disruptive mutation operators.

(B2.3) *Instance space.*

Smith-Miles and colleagues [Smith-Miles and Tan, 2012] introduced a methodology called *instance space analysis*[4] for visualizing problem instances based on features that are correlated with difficulty for particular algorithms. This approach can be seen as combining problem feature and algorithm performance into a single space. Regions of good performance (so called 'footprints') in the instance space indicate the types of problems that specific algorithms can relatively easily solve. Visualizations of instance spaces can also be useful for indicating the spread of a set of problem instances across the space of features and can therefore be used to assess whether a benchmark suite covers a diverse range of instances for the algorithms under study. Example applications of the methodology include analysis of the TSP [Smith-Miles and Tan, 2012] and continuous black-box optimisation problems [Muñoz and Smith-Miles, 2017]. An interesting extension of the approach is to evolve problem instances that fill the gaps in instance space left open by existing problem instances [Smith-Miles and Bowly, 2015], or to directly evolve diverse sets of instances [Neumann et al., 2019].

3.3 Available Benchmark Sets

Over the years, competitions and special sessions at international conferences have provided a wealth of resources for benchmarking of optimization algorithms. Some studies on metaheuristics have also made problems available to be used as benchmarks. This section briefly outlines some of these resources, mostly in alphabetical order of their key characteristic. We have concentrated on benchmark problems that are fundamentally different in nature, and that have documentation and code available online. It is also due to our focus on fundamental differences that we typically do not go into the details of configurable instances and parameterized instance generators.

It is worth mentioning upfront that a number of the benchmark problems mentioned below are available within the optimization software platform Nevergrad [Rapin and Teytaud, 2018].

(B3.1) *Artificial discrete optimization problems.*

Subjectively, this area is among those with the largest number of benchmark sets. Here, many are inspired by problems encountered in the real world, which then have given rise to many fundamental problems in computer science. Noteworthy subareas of discrete optimization are combinatorial optimization, integer and constraint programming—and for many of them large sets of historically grown sets of benchmarks exist. Examples include the Boolean satisfiability and maximum satisfiability

---

[4]https://matilda.unimelb.edu.au/matilda/

competitions[5], travelling salesperson problem library[6] and the mixed integer programming library of problems[7].

In contrast to these instance-driven sets are the more abstract models that define variable interactions at the lowest level (i.e., independent of a particular problem) and then construct an instance based on fundamental characteristics. Noteworthy examples here are (for binary representations) the NK landscapes [Kauffman, 1993] (which has the idea of tunable ruggedness at its core), the W-Model [Weise and Wu, 2018] (with configurable features like length, neutrality, epistasis, multi-objectivity, objective values, and ruggedness), and the Pseudo-Boolean Optimization (PBO) suite of 23 binary benchmark functions by Doerr et al. [2020], which covers a wide range of landscape features and which extends the W-model in various ways (in particular, superposing its transformations to other base problems).

(B3.2) *Artificial real-parameter problems.*
Benchmark suites have been defined for special sessions, workshops and competitions at both the Association for Computing Machinery (ACM) Genetic and Evolutionary Computation Conference (GECCO) and the Institute of Electrical and Electronics Engineers (IEEE) CEC. Documentation and code are available online—for GECCO BBOB[8], and for CEC[9].

(B3.3) *Artificial mixed representation problems.*
Benchmark suites combining discrete and continuous variables include mixed-integer NK landscapes [Li et al., 2006], mixed-binary and real encoded multi-objective problems [McClymont and Keedwell, 2011], mixed-integer problems based on the CEC functions [Liao et al., 2014], and a mixed-integer suite based on the BBOB functions (bbob-mixint) with a bi-objective formulation (bbob-biobj-mixint) [Tušar et al., 2019].

(B3.4) *Black-box optimization problems.*
For all benchmarks listed here, the problem formulation and the instances typically are publicly available, which inevitably leads to a specialization of algorithms to these. The Black-Box Optimization Competition[10] has attempted to address this shortcoming with its single- and multi-objective, continuous optimization problems. Having said this, in 2019, the evaluation code has been made available.

(B3.5) *Constrained real-parameter problems.*
Most real-parameter benchmark problems are unconstrained (except for basic bounds on variables) and there is a general lack of constrained benchmark sets for EC. Exceptions include a set of 18 artificial scalable problems for the CEC 2010 Competition on Constrained Real-Parameter Optimization [11], six constrained real-parameter multi-objective real-world problems presented by Tanabe and Ishibuchi [2020] and a set of 57 real-world constrained problems[12] for both the GECCO and CEC 2020 conferences.

(B3.6) *Dynamic single-objective optimization problems.*
Benchmark problems for analysing Evolutionary Algorithms (EAs) in dynamic environments should ideally allow for the nature of the changes (such as severity and frequency) to be configurable. A useful resource on benchmarks for dynamic environments is the comprehensive review by Nguyen et al. [2012].

---

[5]http://www.satcompetition.org/, https://maxsat-evaluations.github.io/

[6]http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/ and http://www.math.uwaterloo.ca/tsp/index.html

[7]https://miplib.zib.de/

[8]https://coco.gforge.inria.fr

[9]https://github.com/P-N-Suganthan/2020-Bound-Constrained-Opt-Benchmark

[10]https://www.ini.rub.de/PEOPLE/glasmtbl/projects/bbcomp/

[11]https://github.com/P-N-Suganthan/CEC2010-Constrained

[12]https://github.com/P-N-Suganthan/2020-RW-Constrained-Optimisation

(B3.7) *Expensive optimization problems.*
The GECCO 2020 Industrial Challenge provides a suite of discrete-valued electrostatic precipitator problems with expensive simulation-based evaluation[13]. An alternative approach to benchmarking expensive optimization (used by CEC competitions) is to limit the number of allowed function evaluations for solving existing benchmark problems.

(B3.8) *Multimodal optimization (niching).*
Benchmark problem sets for niching include the GECCO and CEC competitions on niching methods for multimodal optimization problems[14] and the single-objective multi-niche competition problems[15].

(B3.9) *Noisy.*
The original version of the Nevergrad platform [Rapin and Teytaud, 2018][16] had a strong focus on noisy problems, but the platform now also covers discrete, continuous, mixed-integer problems with and without constraints, with and without noise, explicitly modelled problems and true black-box problems, etc. The electroencephalography (EEG) data optimization problem set of the CEC Optimization of Big Data 2015 Competition[17] also includes noisy versions of the problem.

(B3.10) *Problems with interdependent components.*
While much research tackles combinatorial optimization problems in isolation, many real-world problems are combinations of several sub-problems [Bonyadi et al., 2019]. The Travelling Thief Problem [Bonyadi et al., 2013] has been created as an academic platform to systematically study the effect of the interdependence, and the 9 720 instances [Polyakovskiy et al., 2014][18] vary in four dimensions. A number of single- and multi-objective as well as static and dynamic extensions of the Travelling Thief Problem have been proposed since then [Sachdeva et al., 2020].

(B3.11) *Real-world discrete optimization.*
The GECCO competition on the optimal camera placement problem (OCP) and the unicost set covering problem (USCP) include a set of discrete real-world problem instances[19]. Other real-world problems include the Mazda benchmark problem[20], which is a scalable, multi-objective, discrete-valued, constrained problem based on real-world car structure design, and a benchmark suite of combinatorial logic circuit design problems [de Souza et al., 2020] that cover a range of characteristics influencing the difficulty of the problem.

(B3.12) *Real-world numerical optimization.*
A set of 57 single-objective real-world constrained problems were defined for competitions at a number of conferences[21]. Other benchmarks include electroencephalography (EEG) data optimization problems [Goh et al., 2015], sum-of-squares clustering benchmark problem set [Gallagher, 2016], the Smart Grid Problems Competitions for real-world problems in the energy domain[22], and the Game Benchmark for EAs [Volz et al., 2019] of test functions inspired by game-related problems[23].

3.4 Open Issues

We see a number of opportunities for research on problems sets for benchmarking.

---

[13] https://www.th-koeln.de/informatik-und-ingenieurwissenschaften/gecco-challenge-2020_72989.php
[14] http://epitropakis.co.uk/gecco2020/
[15] https://github.com/P-N-Suganthan/CEC2015-Niching
[16] https://github.com/facebookresearch/nevergrad
[17] http://www.husseinabbass.net/BigOpt.html
[18] https://cs.adelaide.edu.au/~optlog/research/combinatorial.php
[19] http://www.mage.fst.uha.fr/brevilliers/gecco-2020-ocp-uscp-competition/
[20] http://ladse.eng.isas.jaxa.jp/benchmark/
[21] https://github.com/P-N-Suganthan/2020-RW-Constrained-Optimisation
[22] http://www.gecad.isep.ipp.pt/ERM-competitions/home/
[23] http://www.gm.fh-koeln.de/~naujoks/gbea/gamesbench.html

First, the number of real-world benchmark seems to be orders of magnitude smaller than the actual number of real-world optimisation problems that are tackled on a daily basis—this is especially true for continuous optimization. When there are some proper real-world problems available (e.g. data sets for combinatorial problems, or the CEC problems mentioned), they are often single-shot optimizations, i.e., only a single run can be conducted, which then makes it difficult to retrieve generalizable results. Having said this, a recent push towards a collection and characterization has been made with a survey[24] by the Many Criteria Optimization and Decision Analysis (MACODA) working group.

Second, the availability of diverse instances and of source code (of fitness functions, problem generators, but also of algorithms) leaves much to be desired. Ideal are large collections of instances, their features, algorithms, and their performance—the Algorithm Selection Library (ASlib)[25] [Bischl et al., 2016] has such data, although for a different purpose. As a side effect, these (ideally growing) repositories offer a means against the reinvention of the wheel and the benchmarking against so-called "well-established" algorithms that are cited many times—but maybe just cited many times because they can be beaten easily.

Third, and this is more of an educational opportunity: we as the community need to make sure that we watch our claims when benchmarking. This includes that we not only make claims like "my approach is better than your approach", but that we also investigate what we can learn about the problem and about the algorithms (see e.g. the discussion in [Agrawal et al., 2020] in the context of data mining), so that we can inform again the creation of new instances. Or to paraphrase this: we need to clarify what conclusions can we actually attempt to draw, given the performance comparison is always "with respect to the given benchmark suite".

Fourth, it is an advantage of test problem suites that they can provide an objective means of comparing systems. However, there are also problems related to test problem suites: Whitley et al. [2002] discuss the potential disadvantage that systems can become overfitted to work well on benchmarks and therefore that good performance on benchmarks does not generalize to real-world problems. Fischbach and Bartz-Beielstein [2020] list and discuss several drawbacks of these test suites, namely: (i) problem instances are somehow artificial and have no direct link to real-world settings; (ii) since there is a fixed number of test instances, algorithms can be fitted or tuned to this specific and very limited set of test functions; (iii) statistical tools for comparisons of several algorithms on several test problem instances are relatively complex and not easy to analyze.

Last, while for almost all benchmark problems and for a wide range of real-world problems the fitness of a solution is deterministic, there are also many problems out there where the fitness evaluations are conducted under noise. Hence, the adequate handling of noise can be critical so as to allow algorithms to explore and exploit the search space in a robust manner. Branke et al. [2001] discuss strategies for coping with noise, and Jin and Branke [2005] present a good survey. While noise (in computational experiments) is often drawn from relatively simple distributions, real-world noise can be non-normal, time-varying, and even be dependent on system states. To validate experimental outcomes from such noisy environments, mechanisms way beyond "do n many repetitions" are needed, and Bokhari et al. [2020] compare five such approaches.

4 Algorithms

To understand strengths and weaknesses of different algorithmic ideas, it is important to select a suitable set of algorithms that is to be tested within the benchmark study. While the algorithm portfolio is certainly one of the most subjective choices in a benchmarking study, there are nevertheless a few design principles to respect. In this section we summarize the most relevant of these guidelines.

4.1 Algorithm Families

To assess the quality of different algorithmic ideas, it is useful to compare algorithm instances from different families. For example, one may want to add solvers from the families of

---

[24]https://sites.google.com/view/macoda-rwp/home
[25]https://github.com/coseal/aslib_data

- one-shot optimization algorithms (e.g., pure random search, Latin Hypercube Design (LHD) [McKay, Michael D and Beckman, Richard J and Conover, William J, 2000], or quasi-random point constructions),

- greedy local search algorithms (e.g., randomized local search, Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [Shanno, 1970], conjugate gradients [Fletcher, 1976], and Nelder-Mead [Nelder and Mead, 1965])

- non-greedy local search algorithms (e.g., Simulated Annealing (SANN) [Kirkpatrick et al., 1983], Threshold Accepting [Dueck and Scheuer, 1990], and *Tabu Search* [Glover, 1989])

- single-point global search algorithms (e.g., $(1 + \lambda)$ Evolution Strategies [Eiben and Smith, 2015] and Variable Neighborhood Search [Mladenović and Hansen, 1997])

- population-based algorithms (e.g., Particle Swarm Optimization (PSO) [Kennedy and Eberhart, 1995, Shi and Eberhart, 1998], ant colony optimization [Dorigo et al., 2006, Socha and Dorigo, 2008], most EAs [Bäck et al., 1997, Eiben and Smith, 2015], and Estimation of Distribution Algorithms (EDAlgos) [Larrañaga and Lozano, 2002, Mühlenbein and Paaß, 1996] such as the CMA-ES [Hansen et al., 2003])

- surrogate-based algorithms (e.g., Efficient Global Optimization (EGO) algorithm [Jones et al., 1998] and other Bayesian optimization algorithms)

Note that the "classification" above is by no means exhaustive, nor stringent. In fact, classification schemes for optimization heuristics always tend to be fuzzy, as hybridization between one or more algorithmic ideas or components is not unusual, rendering the attribution of algorithms to the different categories subjective; see [Birattari et al., 2003, Boussaïd et al., 2013, Stork et al., 2020] for examples.

4.2 Challenges and Guidelines for the Practitioner

The following list summarizes considerations that should guide the selection of the algorithm portfolio. This list is not meant to recommend certain algorithms to solve a given problem class, but to give an overview of aspects that should be taken into consideration before a benchmark study starts.

(C4.1) *Problem features*

The arguably most decisive criterion for the algorithm portfolio is the type of problems that are to be benchmarked. Where information such as gradients are available, gradient-based search methods should be included in the benchmark study. Where mixed-integer decision spaces are to be explored, different algorithms are relevant than for purely numerical or purely combinatorial problems. Also, other characteristics such as the degree of variable interaction, the (supposed) shape of the objective value landscape etc. should determine the algorithm portfolio.

We recommend gathering and using all available information about the problem, e.g., its landscape features [Kerschke et al., 2019] and algorithms performances on the problem class from the past [Kerschke et al., 2019]. Even if the goal of the benchmark study does not have a competitive character, a deeper look into results of preceding competitions or workshops can give useful hints which algorithms to select. Benchmark data repositories such as those collected in [Wang et al., 2020] are designed to support the user in these tasks.

(C4.2) *Budget and convergence*

Compute power and the availability of resources to interpret the benchmarking results has a strong impact on the *number of algorithms* that can be compared, whereas the *budget* that can be allocated to each algorithm is typically driven by the research question or application of interest. The budget can have a decisive influence on the selection of algorithms. For example, surrogate-assisted algorithms tend to be algorithms of choice for small budgets, whereas evolution strategies tend to be more competitive for mid- and large-sized budgets.

(C4.3) *State of the art*
Results of a benchmark study can easily be biased if only outdated algorithms are added. We clearly recommend familiarizing oneself with the state-of-the-art algorithms for the given problem type, where state of the art may relate to performance on a given problem class or the algorithm family itself. Preliminary experiments may lead to a proper pre-selection (i.e., exclusion) of algorithms. The practitioner should be certain to compare versus the best algorithms. Consequently, always compare to the most current versions and implementations of algorithms. This also counts for the programming platform and its versions. For algorithm implementations on programming platforms and operating systems the practitioner is not familiar with, nowadays there exist methods and technologies to solve this inconvenience, e.g., container (like Docker) or virtualizations. For details about considerations regarding the experimental design, e.g., the number of considered algorithms, the number of repetitions, the number of problem instances, the number of different parameter settings, or sequential designs if the state of the art is unknown, please see Section 7.

(C4.4) *Hyperparameter handling*
All discussed families of algorithms require one or several control parameters. To enable a fair comparison of their performances and to judge their efficiency, it is crucial to avoid bad parameter configurations and to properly tune the algorithms under consideration [Beiranvand et al., 2017, Eiben and Smit, 2011]. Even a well-working parameter configuration for a certain setup, i.e., a fixed budget, may work comparably worse on a significantly different budget. As mentioned in Section 2 under goal (G2.2), the robustness of algorithms with respect to their hyperparameters can be an important characteristic for users, in which case this question should be integrated into (or even be the subject of) the benchmarking study. Furthermore, the practitioner should be certain that the algorithm implementation is properly using the parameter setup. It may occur that some implementations do not warn the user if the parameter setting is out of their bounds.

Several tools developed for automatic parameter configuration are available, e.g., iterated racing (irace) [López-Ibáñez et al., 2016], Iterated Local Search in Parameter Configuration Space (ParamILS) [Hutter et al., 2009], SPOT [Bartz-Beielstein et al., 2005], Sequential Model-based Algorithm Configuration (SMAC) [Hutter et al., 2011], GGA [Ansótegui et al., 2015], and hyperband [Li et al., 2017] to name a few. As manual tuning can be biased, especially for algorithms unknown to the experimenter, automated tuning is state of the art and highly recommended. Giving rise to a large amount of research in the field of automated algorithm configuration and hyperparameter optimization, there exist several related benchmarking platforms, like the algorithm configuration library (ACLib) [Hutter et al., 2014] or the hyperparameter optimization library (HPOlib) [Eggensperger et al., 2013], which deal particularly with this topic.

(C4.5) *Initialization*
A good benchmarking study should ensure that the results achieved do not happen by chance. It can be important to consider that algorithm performances are not erroneously rated due to the (e.g., random) initialization of the algorithm. According to a given problem instance, a random seed-based starting point can be beneficial for algorithms if they are placed near or at one or more local optima. Consequently, the practitioner should be aware that the performance of algorithms can be biased by the initialization of algorithms with respect to, e.g., their random seeds, the starting points, the sampling strategy, combined with the difficulty of the chosen problem instance.

We recommend letting all candidate algorithms use the same starting points, especially when the goal of the benchmarking study is to compare (goals G1.2 and G1.3) or to analyze the algorithms search behavior (G1.1). This recommendation also extents to the comparison with historical data. Additionally, the design of experiment (see Section 7) can reflect the considerations by properly handling the number of problem instances, repetitions, sampling strategies (in terms of the algorithm parametrization), and random seeds. For (random) seed handling and further reproducibility handling, we refer to Section 9.

(C4.6) *Performance assessment*
Not all algorithms support the configuration of the same stopping criteria, which may influence the search [Beiranvand et al., 2017] and which has to be taken into account in the interpretation of the results. For example, implementation of algorithms may not respect the given number of objective function evaluations. If not detected by the practitioner, this can largely bias the evaluation of the benchmark.

4.3 Challenges and Open Issues

The selection of the algorithms to be included in a benchmarking study depends to a great extent on the users experience, the availability of off-the-shelf or easy-to-adjust implementations, the availability of data about the problems and/or algorithms, etc. Identifying relevant algorithms, data sets, research papers often requires a major effort. Even where data and implementations are easily available, formats can greatly differ between different studies, hindering their efficient use. We therefore believe that common data formats, common benchmark interfaces, and a better compatibility between existing software to assist benchmarking of optimization heuristics is greatly needed.

Another major issue in the current benchmarking landscape concerns a lack of detail in the description of the algorithms. Especially for complex, say, surrogate-assisted optimization heuristics, not all parameters and components are explicitly mentioned in the paper. Where code is available in an open access mode, the user can find these details there, but availability of algorithms implementations is still a major bottleneck in our community.

5 How to Measure Performance?

The performance of algorithms can be measured with regard to several objectives, of which *solution quality* and *consumed budget* are the most obvious two (see Figure 2). In fact, when benchmarking algorithms one usually examines them with regard to one of the following two questions:

- "How fast can the algorithms achieve a given solution quality?" (Section 5.1) or
- "What solution quality can the algorithms achieve with a given budget?" (Section 5.2)

These two scenarios correspond to vertical and horizontal cuts in a performance diagram as discussed by Hansen et al. [2012] and Finck et al. [2015], respectively (see Figure 2). The fixed-budget scenario (vertical cut) comes with the benefit that its results are well-defined as any real computation has a limited budget. Whereas fixing the desired solution quality (horizontal cut) allows to draw conclusions that are easier to interpret; statements such as "algorithm instance b is ten times faster than algorithm instance a in solving this problem" are likely more tangible compared to "the solution quality achieved by algorithm instance b is 0.2% better than the one of algorithm instance a." However, as not all algorithm runs may hit the chosen target, users of *fixed-target measures* need to define how they treat such non-successful runs.

Depending on the chosen time budgets or targeted objective values, different algorithms may yield better results or shorter run times, respectively. Therefore, instead of measuring the objectives using a fixed value, algorithms can also be assessed regarding their *anytime behaviour* [Bossek et al., 2020b, Jesus et al., 2020]. In those cases, the performance does not correspond to a singular point, but instead to an entire curve in the time-quality diagrams. Note that all three views have different implications, and each of them has its justification. As a result, it depends on the application at hand, which perspective should be focussed.

In addition, the *robustness of the found solution* – which might be affected by the algorithm's stochasticity, a noisy optimization problem, or the smoothness of the landscape in a solution's vicinity – can also be in a study's focus. However, as outlined in Section 5.3, measuring this objective can be very challenging.

5.1 Measuring Time

It should be noted that *time* can be measured in different ways with clock or CPU time being the most intuitive. In several combinatorial optimization problems like solving TSP [Kerschke et al., 2018b] or Boolean
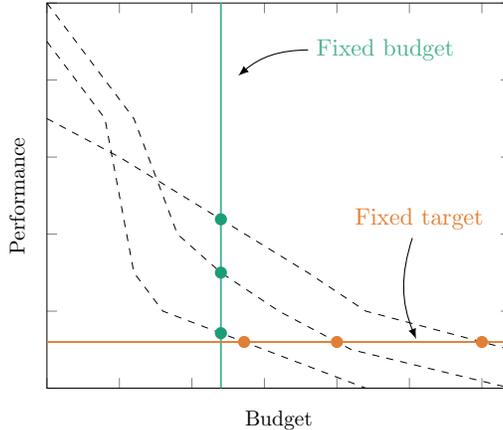
Figure 2: Visualization of a fixed-budget perspective (vertical, green line) and a fixed-target perspective (horizontal, orange line) inspired by Figure 4 in [Hansen et al., 2012]. Dashed lines show three exemplary performance trajectories.

Satisfiability (SAT) problems [Xu et al., 2008]. CPU time is the default. However, as it is highly sensitive to a variety of external factors – such as hardware, programming language, work load of the processors – results from experiments that relied on CPU time are much less replicable and thus hardly comparable. In an attempt of mitigating this issue, Johnson and McGeoch [2002] proposed a normalized time, which is computed by dividing the runtime of an algorithm by the time a standardized implementation of a standardized algorithm requires for the same (or at least a comparable) problem instance.

An alternative way of measuring time are Function Evaluations (FEs), i.e., the number of fully evaluated candidate solutions. In fact, in case of sampling-based optimization, like in classical continuous optimization, this machine-independent metric is *the* most common way of measuring algorithmic performances [Hansen et al., 2016a]. Yet, from the perspective of actual clock time, they risk giving a wrong impression as the FEs of different algorithms might be of different time complexity [Weise et al., 2014]. In such cases, counting algorithm steps in a domain-specific method – e.g., the number of distance evaluations on the TSP [Weise et al., 2014] or bit flips on the Maximum Satisfiability (MAX-SAT) problem [Hains et al., 2013] – may be useful. Nevertheless, within the EC community, counting FEs is clearly the most commonly accepted way to measure the efforts spend by an algorithm to solve a given problem instance.

From a practical point of view, both options have their merits. If the budget is given by means of clock time – e.g., models have to be trained until the next morning, or they need to be adjusted within seconds in case of predictions at the stock market – then results relying on CPU time are more meaningful. On the other hand, in case single FEs are expensive – e.g., in case of physical experiments or cost-intensive numerical simulations – the number of required FEs is a good proxy for clock time, and a more universal measure, as discussed above. To satisfy all perspectives, best practice would be to report both: FEs *and* the CPU time. Moreover, in situations in which single FEs are expensive, CPU time should ideally be separated into the part used by the expensive FE, and the part used by the algorithm at each iteration.

In surrogate-based optimization, algorithms commonly slow down over time due to an ever-increasing complexity of their surrogate models [Bliek et al., 2020, Ueno et al., 2016]. In this case, it is even possible that the algorithm becomes more expensive than the expensive FE itself. By measuring the CPU time used by the algorithm separately from the CPU time used by the FEs, it can be verified that the number of FEs is indeed the limiting factor. At the same time, this reveals more information about the benchmark, namely how expensive it is exactly, and whether all FEs have the same cost. However, if the FEs are expensive not because of the computation time but due to some other costs (deteriorating biological samples, the use of expensive equipment, human interaction, etc.), then just measuring the number of FEs could be sufficient.

Noticeably, many papers in the EC community also use *generations* as machine-independent time mea-

sures. However, it might not be a good idea to *only* report generations, because the exact relationship between FEs and generations is not always clear. This makes results hard to compare with, e.g., local search algorithms, so if generations are reported, FEs should be reported as well.

5.2 Measuring Solution Quality

There are several "natural" quality metrics, e.g., fitness in continuous optimization, the tour length in TSP, the accuracy of a classification algorithm in a machine learning task, or the number of ones in a binary bit string in case of OneMax. However, interpreting these objective values on their own is usually quite difficult and also very specific to the respective problem instance. Instead, one could ideally try to use more intuitive and less problem-dependent alternatives as described below [Johnson, 2002a, Talbi, 2009].

If an instance's optimal solution is known, the (absolute or relative) difference to the optimal target quality could be used. Alternatively, a best-known lower bound for the optimal objective value could be used for normalizing the results. For instance, in case of the TSP, results are often compared to the Held-Karp lower bound [Johnson, 2002a]. As *absolute* differences are very specific to the scaling of the problem's objective values, it is highly recommended to rather look at the *relative* excess over the optimal solution – something that has been common practice in solving TSPs for decades [Christofides, 1976]. Noticeably, despite the varying scales of the objective values across the different problem instances, absolute differences are in the focus of continuous optimization benchmarks like BBOB [Hansen et al., 2016b].

Another alternative is to use results of a clearly specified and ideally simple heuristic for normalization [Johnson, 2002a]. The relative excess over the best-known solution is also often reported. This is typical in the Job Shop Scheduling domain, as shown by many references listed in Weise [2019]. However, this requires an exact knowledge of the related work and may be harder to interpret later in the future. For some problems, reference solutions may be available and the excess over their quality can be reported.

Constraint optimization Under constraint optimization, a solution is either feasible or not, which is decided based on a set of constraints. Here, the absolute violations of each constraint can be summed up as a performance metric [Hellwig and Beyer, 2019, Kumar et al., 2020].

5.3 Measuring Robustness

In terms of robustness analyses, one can differentiate between three reasons for volatility among the results: (i) stochastic search behavior of the considered algorithm (e.g., in randomized search heuristics), (ii) noisy problems, (iii) ruggedness or smoothness of the problem landscape.

From a practical point of view, rugged landscapes can be highly problematic for the outcome of an optimization problem. For instance, when controlling parameters of an airplane or conducting medical surgeries, the global optimum will likely not be targeted, if slight variations (in search space) can have hazardous effects on the objective space and thus the on the whole system [Branke, 1998, Tsutsui et al., 1996]. In such scenarios, one likely is much more interested in finding (local) optima, whose objective values are still very close to the global optimum, but change only slightly when perturbing the underlying solutions.

Another common issue in real-world applications is the existence of noise. In particular, in case of physical experiments or stochastic simulations models, the outcome of an experiment may vary despite using the same candidate solution [Arnold, 2012, Cauwet and Teytaud, 2016].

The third issue related to robustness investigations is the stochastic nature of the algorithms themselves. In fact, many sampling-based optimization algorithms are nowadays randomized search heuristics and as such their performances will vary if the experiment is repeated, i.e., if the algorithm is executed again using the same input. Therefore, it is common to use performance metrics that aggregate the results of several (ideally independent) runs to provide reliable estimates of the algorithm performance.

Location: Measures of Central Behaviors In case of a fixed budget (i.e., vertical cut) approach, solution qualities are usually aggregated using the arithmetic mean. But of course, other location measures

like the median or the geometric mean [Fleming and Wallace, 1986] can be useful alternatives when interested in robust metrics or when aggregating normalized benchmark results, respectively.

In scenarios, in which the primary goal is to achieve a desired target quality (horizontal cut), it might be necessary to aggregate successful and failed runs. In this case, two to three metrics are mostly used for aggregating performances across algorithm runs, and we will list them below.

The gold standard in (single-objective) continuous optimization is the Expected Running Time (ERT) [Auger and Hansen, 2005, Price, 1997], which computes the ratio between the sum of consumed budget across all runs and the number of successful runs [Hansen et al., 2012]. Thereby, it estimates the average running time an algorithm needs to find a solution of the desired target quality (under the assumption of independent restarts every T time units until success).

In other optimization domains, like TSP, SAT, etc., the *Penalized Average Runtime (PAR)* [Bischl et al., 2016] is more common. It penalizes unsuccessful runs with a multiple of the maximum allowed budget – penalty factors ten (PAR10) and two (PAR2) are the most common versions – and afterwards computes the arithmetic mean of the consumed budget across all runs. The *Penalized Quantile Runtime (PQR)* [Bossek et al., 2020a, Kerschke et al., 2018a] works similarly, but instead of using the arithmetic mean for aggregating across the runs, it utilizes quantiles—usually the median – of the (potentially penalized) running times. In consequence, PQR provides a robust alternative to the respective PAR scores.

Spread and Reliability   A common measure of reliability is the *estimated success probability*, i.e., the fraction of runs that achieved a defined goal. Bossek et al. [2020a] use it to take a multi-objective view by combining the probability of success and the average runtime of successful runs. Similarly, Hellwig and Beyer [2019] aggregate the two metrics in a single ratio, which they called *SP*.

As measures of dispersion of a given single performance metric, statistics like standard deviations as well as quantiles are used, whereas the latter are more robust.

For constraint optimization, a *feasibility rate* (FR) [Kumar et al., 2020, Wu et al., 2017] is defined as the fraction of runs discovering at least one feasible solution. The number of constraints violated by the median solution [Kumar et al., 2020] and the mean amount of constraint violation over the best results of all runs [Hellwig and Beyer, 2019] can also be used.

5.4   Open Issues

Although each of the different optimization domains has established its preferable performance metric, research in this field is still facing open issues. For instance, so far performance is mostly measured using fixed values (budget or target). However, depending on the use case, comparing the anytime behavior of algorithms might occasionally be of interest as well.

Aside from facing challenges like measuring quality *and* time simultaneously, we also have to integrate costs for violating constraints (constraint optimization), quantify variation or uncertainty (robust/noisy optimization), measure the spread across the local optima (multimodal optimization), or capture the proximity of the population to the local and/or global optima of the problem.

6   How to Analyze Results?

6.1   Three-Level Approach

Once the performance measure for the algorithm's performance is selected by the user and all data related to it is collected in experiments, the next step is to analyse the data and draw conclusions from it. From the detailed characterization of possible benchmark goals in Section 2, we will focus on goals (G1.2) and (G1.3), i.e., algorithm comparison and competition of several algorithms. Therefore, we will consider:

- single-problem analysis and

- multiple-problem analysis.

In both scenarios, multiple algorithms will be considered, i.e., following the notation introduced in Section 2, there are at least two different algorithm instances, say, a_j and a_k from algorithm A or at least two different algorithm instances $a_j \in A$ and $b_k \in B$, where A and B denote the corresponding algorithms. Single-problem analysis is a scenario where the data consists of multiple runs of the algorithms on a single problem instance $\pi_i \in \Pi$. This is necessary because many optimization algorithms are stochastic in nature, so there is no guarantee that the result will be the same for every run. Additionally, the path leading to the final solution is often different. For this reason, it is not enough to perform just a single algorithm run per problem, but many runs are needed to make a conclusion. In this scenario, the result from the analysis will give us a conclusion which algorithm performs the best on that specific problem.

Otherwise, in the case of multiple-problem analysis, focusing on (G1.2), we are interested in comparing the algorithms on a set of benchmark problems. The best practices of how to select a representative value for multiple-problem analysis will be described in Section 7.

No matter of what we are performing, i.e., single-problem or multiple-problem analysis, the best practices analyzing the results of the experiments suggest that the analysis can be done as a three-level approach, which consists of the following three steps:

1. Exploratory Data Analysis (EDA)

2. Confirmatory Analysis

3. Relevance Analysis

This section focuses on analyzing the empirical results of an experiment using descriptive, graphical, and statistical tools, which can be used for the three-level approach for analysis. More information about various techniques and best practices analyzing the results of experiments can be found in Crowder et al. [1979], Golden et al. [1986], Barr et al. [1995], Bartz-Beielstein et al. [2004], Chiarandini et al. [2007], García et al. [2009], Bartz-Beielstein et al. [2010], Derrac et al. [2011], Eftimov et al. [2017], Beiranvand et al. [2017]. Mersmann et al. [2010], and more recently Kerschke and Trautmann [2019a], present methods based on ELA to answer two basic questions that arise when benchmarking optimization algorithms. The first one is: which algorithm is the 'best' one? and the second one: which algorithm should I use for my real world problem? In the following, we summarize the most accepted and standard practices to evaluate the considered algorithms stringently. These methods, if adhered to, may lead to wide acceptance and applicability of empirically tested algorithms and may be a useful guide in the jungle of statistical tools and methods.

6.2 Exploratory Data Analysis

6.2.1 Motivation

Exploratory Data Analysis (EDA) is an elementary tool that employs descriptive and graphical techniques to better understand and explore empirical results. It must be performed to validate the underlying assumptions about the distribution of the results, e.g., normality or independence, before implementing any statistical technique that will be discussed in Section 6.3.

We recommend starting with EDA to understand basic patterns in the data. It is useful to prepare (statistical) hypotheses, which are the basis of confirmatory analysis. In EDA, visual tools are preferred, whereas confirmatory analysis is based on probabilistic models. EDA provides a flexible way to analyze data without preconceptions. Its tools stem from descriptive statistics and use an inductive approach, because in the beginning, there is no theory that has to be validated. One common saying is "let the data speak", so data suggest interesting questions, e.g., unexpected outliers might indicate a severe bug in the algorithm. EDA is a very flexible way to generate hypotheses, which can be analyzed in the second step (confirmatory analysis). Although EDA might provide deeper understanding of the algorithms, it does not always provide definitive answers. Then, the next step (confirmatory analysis) is necessary. And, there is also the danger of overfitting: focussing on very specific experimental designs and results might cause a far too pessimistic (or optimistic) bias. Finally, it is based on experience, judgement, and artistry. So there is no standard cookbook available, but many recipes.

The following are the key tools available in EDA. It can provide valid conclusions that are graphically presented, without requiring further statistical analysis. For further reading about EDA, the reader is referred to [Tukey, 1977].

6.2.2   The Glorious Seven

Descriptive statistics include the mean, median, best and worst (minimum and maximum, respectively), first and third quartile, and standard deviation of the performance measures of the algorithms. These seven so-called summary statistics measure the central tendency and the variability of the results. Note, they might be sensitive to outliers, missing, or biased data. Most importantly, they do not provide a complete analysis of the performance, because they are based on a very specific data sample. For example, mean and standard deviation are affected by outliers, which might exist because of an algorithms' poor runs and variability. Both can be caused by an inadequate experimental design, e.g., selection of improper starting points for the algorithm or too few function evaluations. The median is more robust statistic than the mean if sufficiently many data points are available. The best and the worst value of the performance measure provide insights about the performance data, but they consider only one out of n data points, i.e., they are determined by one single data and therefore not very robust compared to the mean or median that are considering all data points. The quantiles are cut points which split a probability distribution into continuous intervals with equal probabilities. Similar to the median, they require a certain amount of data points and are probably meaningless for small data. Bartz-Beielstein [2006] presents a detailed discussion of these basic statistics.

6.2.3   Graphical Tools

Visualising final results.   Graphical tools can provide more insight into the results and their distributions. The first set of graphical tools can be used to analyse the final results of the optimization runs. *Histograms* and *boxplots* are simple but effective tools and provide more information for further analysis of the results. Box plots visualize the distribution of the results. They illustrate the statistics introduced in Section 6.2.2 in a very compact and comprehensive manner and provide means to detect outliers. Histograms provide information about the shape of the distribution. Because the shape of histograms is highly affected by the size of the boxes, we strongly recommend combining histograms with density plots.

Visualising run-time behaviour.   The second set of tools can be used to analyse the algorithm performance over time, i.e., information about the performance for every lth iteration is required. Suitable for the analysis of the performances of the optimization algorithms are *convergence plots* in which the performance of the algorithm can be evaluated against the number of function evaluations. This helps us to understand the dynamics of multiple algorithms in a single plot.

Histograms and box plots are also used in the graphical multiple problem analysis. Besides these common tools, specific tools for the multiple problem analysis were developed, e.g., *performance profiles* proposed in Dolan and Moré [2002]. They have emerged as an important tool to compare the performances of optimization algorithms based on the cumulative distribution function of a performance metric (CPU time, achieved optimum). It is the ratio of a performance metric obtained by each algorithm versus the best value the performance metric among all algorithms that is being compared. Such plots help to visualize the advantages (or disadvantages) of each competing algorithm graphically. Performance profiles are not applicable, if the (true or theoretical) optimum is unknown. However, there are solutions for this problem, e.g., using the best known solution so far or a guessed (most likely) optimum based on the user's experience; however, the latter is likely to be error-prone.

As performance profiles are not evaluated against the number of function evaluations, they cannot be used to infer the percentage of the test problems that can be solved with some specific number of function evaluations. To attain this feature, the *data profiles* were designed for fixed-budget derivative free optimization algorithms [Moré and Wild, 2009]. It is appropriate to compare the best possible solutions obtained from various algorithms within a fixed budget.

6.3 Confirmatory Analysis

6.3.1 Motivation

The second step in the three-level approach is referred to as *confirmatory analysis*, which is based in inferential statistics, because it implements a deductive approach: a given assumption (statistical hypothesis) is tested using the experimental data. Since the assumptions are formulated as statistical hypotheses, confirmatory analysis heavily relies on probability models. Its final goal is to provide definite answers to specific questions, i.e., questions for a specific experimental design. Because it uses probability models, its emphasis is on complex numerical calculations. Its main ingredients are hypothesis tests and confidence intervals. Confirmatory analysis usually generates more precise results for a specific context than EDA. But, if the context is not suitable, e.g., statistical assumptions are not fulfilled, a misleading impression of precision might occur.

Often, EDA tools are not sufficient to clearly analyze the differences in the performances of algorithms, mainly when the differences are of smaller magnitude. The need to perform statistical analysis and various procedures involved in making decisions about selecting the best algorithm are widely discussed in [Amini and Barr, 1993, Barr et al., 1995, Carrano et al., 2011, Chiarandini et al., 2007, Eftimov et al., 2017, García et al., 2009, Golden et al., 1986, McGeoch, 1996]. The basic idea of statistical analysis is based on hypothesis testing. Before analysing the performance data, we should define two hypotheses i) the null hypothesis H_0 and ii) the alternative hypothesis H_1. The null hypothesis states that there is no significant statistical difference between the two algorithms' performances, while the alternative hypothesis directly contradicts the null hypothesis by indicating the statistical significance between the algorithms' performances. Hypothesis testing can be two-sided or one-sided. We will consider the one-sided case in the following, because it allows us to ask if algorithm instance a is better than algorithm instance b. Let $p(a)$ denote the performance of algorithm a. In the context of minimization, smaller performance values will be better, because we will compare the best solutions or the run times. The statement "a outperforms b" is equivalent to "$p(a) < p(b)$", which can be formulated as the statistical hypothesis $H_1 : p(b) - p(a) > 0$. It is a common agreement in hypotheses testing that this hypothesis H_1 will be tested against the null hypothesis $H_0 : p(b) - p(a) \leq 0$, which states that a is not better than b.

After the hypotheses are defined, we should select an appropriate statistical test, say T, for the analysis. The test statistic T which is a function of a random sample that allows researchers to determine the likelihood of obtaining the outcomes if the null hypothesis is true. The mean of the best found values from n repeated runs of an algorithm is a typical example of a test statistic. Additionally, a significance level α should be selected. Usually, a significance level of 95% is used. However, the selection of this value depends on the experimental design and the scientific question to be answered.

6.3.2 Assumptions for the Safe Use of the Parametric Tests

There are *parametric* and *non-parametric* statistical tests. To select between them, there are assumptions for the safe use of the parametric tests. Common assumptions include independence, normality, and homoscedasticity of variances. The independence assumption is directly met as the results of independent runs of the algorithm with randomly generated initial seeds are being compared. To check the normality assumption several tests can be performed including *Kolmogorov-Smirnov test* [Sheskin, 2003], *Shapiro-Wilk test* [Shapiro and Wilk, 1965], and *Anderson Darling test* [Anderson and Darling, 1952]. The normality assumption can be also checked by using graphical representation of the data using *histograms, empirical distribution functions* and *quantile-quantile* plots (Q-Q plots) [Devore, 2011]. The *Levene's test* [Levene, 1961] and *Bartlett's test* [Bartlett, 1937] can be performed to check if the assumption of equality of variances is violated. We should also mention that there are transformation approaches that may help to attain the normality, but this should be done with a great care, since we are changing the decision space. If the required assumptions are satisfied then we are selecting a parametric test since it has higher power than a non-parametric one, otherwise we should select a non-parametric one.

Additionally to the assumptions for the safe use of the parametric tests, before selecting an appropriate statistical test, we should take care if the performance data is paired or unpaired. Paired data is data in which

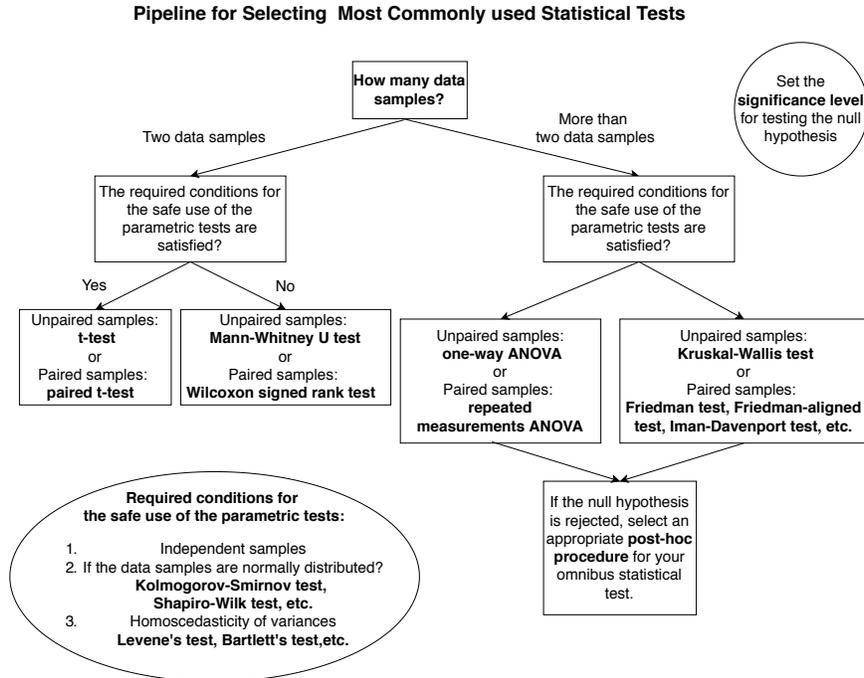**Pipeline for Selecting Most Commonly used Statistical Tests**

Figure 3: A pipeline for selecting an appropriate statistical test [Eftimov et al., 2020].

natural or matched couplings occur. This means that each data value in one sample is uniquely paired to a data value in the other sample. The choice between paired and unpaired samples depends on experimental design, and researchers need to be aware of this when designing their experiment. Using Common Random Numbers (CRN) is a well-known technique for generating paired samples. If the same seeds are used during the optimization, CRNs might reduce the variances and lead to more reliable statistical conclusions [Kleijnen, 1988, Nazzal et al., 2012].

6.3.3 A Pipeline for Selecting an Appropriate Statistical Test

A pipeline for selecting an appropriate statistical test for benchmarking optimization algorithms is presented in Figure 3. Further, we are going to explain some of them depending upon the benchmarking scenario (i.e., single-problem or multiple-problem analysis).

Single-problem analysis. As we previously mentioned, in this case, the performance measure data is obtained using multiple runs of k algorithm instances $a_1, \ldots, a_k$ on one selected problem instance π_j.

The comparison of samples in pairs is called a *pairwise comparison*. Note, that a pairwise comparison of algorithms does not necessarily mean that the corresponding samples are paired. In fact, most pairwise comparisons use unpaired samples, because the setup for pairwise sampling is demanding, e.g., implementing random number streams etc. If more than two samples are compared at the same time, a *multiple comparison* is performed.

For pairwise comparison, the t test [Sheskin, 2003] is the appropriate parametric one, while its non-parametric version is the *Mann-Whitney U test* (i.e., *Wilcoxon-rank sum test*) [Hart, 2001]. In the case when more than two algorithms are involved, the parametric version is the one-way *ANOVA* [Lindman, 1974, Montgomery, 2017], while its appropriate non-parametric test is the *Kruskal-Wallis rank sum test* [Kruskal and Wallis, 1952]. Here, if the null hypothesis is rejected, then we should continue with a post-hoc procedure to define the pairs of algorithms that contribute to the statistical significance.

Multiple-problem analysis. The mean of the performance measure from multiple runs can be used as representative value of each algorithm on each problem. However, as stated above, averaging is sensitive to outliers, which needs to be considered especially because optimization algorithms could have poor runs. For this reason, the median of the performance measure from the multiple runs can also be used as more robust statistic.

Both mean and median are sensitive to errors inside some ϵ-neighborhood (i.e., small difference between their values that is not recognized by the ranking schemes of the non-parametric tests), which can additionally affect the statistical result. For these reasons, Deep Statistical Comparison (DSC) for comparing evolutionary algorithms was proposed [Eftimov et al., 2017]. Its main contribution is its ranking scheme, which is based on the whole distribution, instead of using only one statistic to describe the distribution, such as mean or median.

The impact of the selection of the three above-presented transformations, which can be used to find a representative value for each algorithm on each problem, to the final result of the statistical analysis in the multiple-problem analysis is presented in [Eftimov and Korošec, 2018].

Statistical tests. No matter which transformation is used, once the data for analysis is available, the next step is to select an appropriate statistical test. For pairwise comparison, the *t test* is the appropriate parametric one [Sheskin, 2003], while its relevant non-parametric version is the *Wilcoxon signed rank test* [Wilcoxon, 1945]. In the case when more than two algorithms are involved, the parametric version is the repeated measurements *ANOVA* [Lindman, 1974, Montgomery, 2017], while its appropriate non-parametric tests are the *Friedman rank-based test* [Friedman, 1937], *Friedman-aligned test* [García et al., 2009], and *Iman-Davenport test* [García et al., 2009]. Additionally, if the null hypothesis is rejected, same as the single-problem analysis, we should continue with a post-hoc procedure to define the pairs of algorithms that contribute to the statistical significance.

Other non-parametric tests are the non-parametric rank based tests, which are suitable when the distribution assumptions are questionable [Sheskin, 2003]. Using them, the data is ranked and then the p-value is calculated for the ranks and not the actual data. This ranking helps to eliminate the problem of skewness and in handling extreme values. The *permutation test* [Pesarin, 2001] estimates the permutation distribution by shuffling the data without replacement and identifying almost all possible values of the test statistic. The *Page's trend test* [Derrac et al., 2014] is also a non-parametric test that can be used to analyse convergence performance of evolutionary algorithms.

Post-hoc procedures. When we have more than two algorithms that are involved in the comparison, the appropriate statistical test can find statistical significance between the algorithms' performances, but it is not able to define the pairs of algorithms that contribute to this statistical significance. For this reason, if the null hypothesis is rejected, we should continue with a *post-hoc* test.

The *post-hoc* testing can be done in two scenarios: i) all pairwise comparisons and ii) multiple comparisons with a control algorithm. Let us assume that we have k algorithms involved in the comparison, so in the first scenario we should perform $k(k-1)/2$ comparisons, and in the second one $k-1$.

In the case of all pairwise comparisons, the post-hoc test statistic should be calculated. It depends on the appropriate statistical test that is used to compare all algorithms together, which rejected the null hypothesis. After that, the obtained p-values are corrected with some post-hoc procedure. For example, if the null hypothesis in the *Friedman test*, *Friedman aligned-ranks test*, or *Iman–Davenport test*, is rejected, we can use the *Nemenyi*, *Holm*, *Shaffer*, and *Bergmann* correction to adapt the p-values and handle multiple-testing issues.

The multiple comparisons with a control algorithm is the scenario when our newly developed is the control algorithm, and we are comparing it with state-of-the-art algorithms. Same as the previous scenario, the post-hoc statistic depends on the appropriate statistical test that is used to compare all algorithms together, which rejected the null hypothesis, and the obtained p-values are corrected with some post-hoc procedure. In the case of *Friedman test*, *Friedman aligned-ranks test* , or *Iman–Davenport test*, appropriate post-hoc procedures are: *Bonferroni*, *Holm*, *Hochberg*, *Hommel*, *Holland*, *Rom*, *Finner*, and *Li*.

Another way to the multiple comparisons with a control algorithms is to perform all comparisons between the control algorithm and each other algorithm using some pairwise test. In this case, we should be careful when making a conclusion, since we are losing the control on the Family-Wise Error Rate (FWER) when performing multiple pairwise comparisons. All obtained p-values will come from independent pairwise comparisons. The calculation of the true statistical significance for combining pairwise comparisons is presented in [Eftimov et al., 2017, García et al., 2009].

More information about different post-hoc procedures and their application in benchmarking theory in evolutionary computation is presented in [García et al., 2009].

6.4   Relevance Analysis

6.4.1   Motivation

The third step of the recommended approach is related to the practical relevance of our statistical findings: are the differences really meaningful in practice or are they only statistical "artifacts" caused by an inadequate experimental design? A typical example for these artifacts is a difference in performance, say δ, which is statistically significant but of no practical relevance, because a value as small as δ cannot be measured in real-world scenarios. So, there is still a gap when transferring the learned knowledge from theory to practice. This happens because the statistical significance that exists is not scientifically meaningful in a practical sense.

Example 6.1 (Assembly line). *Let us assume that two optimization algorithms that should minimize the average time of a production process, e.g., an assembly line, are compared. The mean difference in performance is $\delta = 10^{-14}$, which is statistically significant. However, this difference has no meaning in reality, because it is far below the precision of the assembly line timer.*

For this reason, when we are performing a statistical analysis, we should also try to find the relevance of the statistical significance to real world applications. We should also mention that the practical significance depends on the specific problem being solved. Additionally, this is also true in benchmarking performed for scientific publications, where the comparisons of the performance measures can be affected by several factors such as computer accuracy (i.e., floating points), variable types (4-byte float, 8-byte float, 10-byte float), or even the stopping criteria that is the error threshold when the algorithms are stopped. All these factors can result in different values, which does not represent the actual performance of the algorithms even if statistical significance is found.

6.4.2   Severity: Relevance of Parametric Test Results

In order to probe the meaningfulness of the statistically significant result, it is suggested to perform a post-data analysis. One such post data analysis is the severity measure, a meta statistical principle [Bartz-Beielstein et al., 2010, Mayo and Spanos, 2006]. Severity describes the degree of support to decisions made using classical hypothesis testing. Severity takes into account the data and performs a post-data evaluation to scrutinize the decisions made by analyzing how well the data fits the testing framework. The severity is the actual power attained in the post data analysis and can be described separately for the decision of either rejecting or not rejecting the null hypothesis.

The conclusions obtained from the hypothesis testing is dependent on sample size and can suffer from the *problem of large n*. Severity deals with this problem directly [Mayo and Spanos, 2006].

6.4.3   Multiple-Problem Analysis

We present two approaches that investigate the scientific meaningfulness of statistically significant results in the multiple-problem setting. One approach is the Chess Rating Systems for Evolutionary Algorithms (CRS4EA), which is an empirical algorithm for comparing and ranking evolutionary algorithms [Veček et al., 2014]. It makes a chess tournament where the optimization algorithms are considered as chess players and

a comparison between the performance measures of two optimization algorithms as the outcome of a single game. A draw limit that defines when two performance measure values are equal should be specified by the user and it is a problem specific. At the end, each algorithm has its own rating which is a result from the tournament and the statistical analysis is performed using confidence intervals calculated using the algorithms rating.

The second approach is the practical Deep Statistical Comparison (pDSC), which is a modification of the DSC approach used for testing for statistical significance [Eftimov and Korošec, 2019]. The basic idea is that the data on each problem should be pre-processed with some practical level specified by a user, and after that involved with DSC to find relevant difference. Two pre-processing steps are proposed: i) sequential pre-processing, which pre-processes the performance measures from multiple runs in a sequential order, and ii) a Monte-Carlo approach, which pre-processes the performance measure values by using a Monte-Carlo approach to avoid the dependence of the practical significance on the order of the independent runs. A comparison between the CRS4EA and pDSC is presented in [Eftimov and Korošec, 2019]. Using these two approaches, the analysis is made for a multiple-problem scenario. Additionally, the rankings from pDSC obtained on a single-problem level can be used for single-problem analysis.

6.5 Open Issues

An important aspect not addressed in this iteration of the document is the analysis of the benchmark problems themselves rather than the performance of the algorithms operating thereon. That is, which means exist for investigating structural characteristics of the benchmarking problem at hand? How can one (automatically) extract its most relevant information? How should this information be interpreted? There exists a variety of approaches for this, and each of them helps to improve the understanding of the respective problem, and in consequence may facilitate the design, selection and/or configuration of a suitable algorithm.

Linked to the above is a discussion of methods for visualizing problem landscapes. Visualizing the landscape of a continuous problem, or plotting approximate tours for a given TSP instance, usually improves our understanding of its inherent challenges and reveals landscape characteristics such as multimodality. Moreover, such visualizations also help to study the search behavior of the algorithms under investigation. Unfortunately, the vast majority of works treats the issue of visualizing problems very poorly, so we will make sure to address this particular issue in the continuation of this document.

7 Experimental Design

7.1 Design of Experiments (DoE)

Unfortunately, many empirical evaluations of optimization algorithms are performed and reported without addressing basic experimental design considerations [Brownlee, 2007]. An important step to make this procedure more transparent and more objective is to use DOE and related techniques. They provide an algorithmic procedure to make comparisons in benchmarking more transparent. Experimental design provides an excellent way of deciding which and how many algorithm runs should be performed so that the desired information can be obtained with the least number of runs.

DOE is planning, conducting, analyzing, and interpreting controlled tests to evaluate the influence of the varied factors on the outcome of the experiments. The importance and the benefits of a well designed planned experiment have been summarized by Hooker [1996]. Johnson [2002b] suggests to report not only the run time of an algorithm, but also explain the corresponding adjustment process (preparation or tuning before the algorithm is run) in detail, and therefore to include the time for the adjustment in all reported running times to avoid a serious underestimate.

The various key implications involved in the DOE are clearly explained in Kleijnen [2001]. A comprehensive list of the recent publications on design techniques can be found in Kleijnen [2017]. The various design strategies in the Design and Analysis of Computer Experiments (DACE) are discussed by Santner

et al. [2003]. Wagner [2010] discusses important experimental design topics, e.g., "How many replications of each design should be performed?" or "How many algorithm runs should be evaluated?"

This section discusses various important practical aspects of formulating the design of experiments for a stochastic optimization problem. The key principles are outlined. For a detailed reading of the DOE, the readers are referred to Montgomery [2017] and Kleijnen [2015].

7.2 Design Decisions

Design decisions can be based on geometric or on statistical criteria [Pukelsheim, 1993, Santner et al., 2003]. Regarding geometric criteria, two different design techniques can be distinguished: The samples can be placed either (1) on the boundaries, or (2) in the interior of the design space. The former technique is used in DOE whereas DACE uses the latter approach. An experiment is called *sequential* if the experimental conduct at any stage depends on the results obtained so far. Sequential approaches exist for both variants.

We recommend using factorial designs or space-filling designs instead of the commonly used One-factor-at-a-time (OFAT) designs. When several factors are involved in an experiment, the OFAT design strategy is inefficient as it suffers from various limitations including huge number of experimental runs and inability to identify the interactions among the factors involved. It is highly recommended using a multi-factorial design [Montgomery, 2017]. The Factorial designs are robust and faster when compared with OFAT. For a complete insight into Fully-Factorial and Fractional Factorial designs the readers are redirected to Montgomery [2017]. The Taguchi design [Roy, 2001] is a variation of the fractional factorial design strategy, which provides robust designs at better costs with fewer evaluations. The Plackett and Burman design [Plackett and Burman, 1946] are recommended for screening. The modern space-filling designs are sometimes more efficient and require fewer evaluations than the fractional designs, especially in case of non-linearity. Further information about space-filling designs can be found in Santner et al. [2003].

However, it is still an open question which design characteristics are important: "... extensive empirical studies would be useful for better understanding what sorts of designs perform well and for which models" [Santner et al., 2003, p. 161].

7.3 Designs for Benchmark Studies

In the context of DOE and DACE, runs of an optimization algorithm instance will be treated as experiments. There are many degrees of freedom when an optimization algorithm instance is run. In many cases optimization algorithms require the determination of parameters (e.g., the population size in Evolution Strategys (ESs)) before the optimization run is performed. From the viewpoint of an experimenter, design variables (factors) are the parameters that can be changed during an experiment. Generally, there are two different types of factors that influence the behavior of an optimization algorithm:

1. problem-specific factors, i.e., the objective function,

2. algorithm-specific factors, i.e., the population size of an ES and other parameters which need to be set to derive an executable algorithm instance.

We will consider experimental designs that comprise problem-specific factors and algorithm-specific factors. Algorithm-specific factors will be considered first. *Implicit parameters* can be distinguished from *explicit parameters* (synonymously referred to as *endogeneous* and *exogeneous* in [Beyer and Schwefel, 2002]). The latter are explicitly exposed to the user, whereas the former are often hidden, i.e., either made inaccessible to the user (e.g., when the algorithm code is not made available) or simply "hidden" in the implementation and not easily identifiable as a parameter that can be optimized.

An *algorithm design* is a set of parameters, each representing one specific setting of the design variables of an algorithm and defining an algorithm instance. A design can be specified by defining ranges of values for the design parameters. Note that a design can contain none, one, several, or even infinitely many design points, each point representing an algorithm instance. Consider the set of explicit strategy parameters for PSO algorithms with the following values: swarm size $s = 10$, cognitive parameter $c_1 \in [1.5, 2]$, social parameter

$c_2 = 2$, starting value of the inertia weight $w_{max} = 0.9$, final value of the inertia weight $w_{scale} = 0$, percentage of iterations for which w_{max} is reduced $w_{iterscale} = 1$, and maximum value of the step size $v_{max} = 100$. This algorithm design contains infinitely many design points, because c_1 is not fixed.

Problem designs provide information related to the optimization problem, such as the available resources (number of function evaluations) or the problem's dimension.

An *experimental design* consists of a problem design and an algorithm design. Benchmark studies require complex experimental designs, because they are combinations of several problem and algorithm designs. Furthermore, as discussed in Section 5, one or several performance measures must be specified.

7.4   How to Select a Design for Benchmarking

The following points have to be considered when designing an benchmark study[26]:

- What are the main goals of the experiment? (see Section 2)

- What is/are the test problem(s) and which (type of) instances do we select? (see Section 3)

- How many algorithms are to be tested? (see Section 4)

- How many test problems/test classes are relevant for the study? (see Section 3)

- How tuning of algorithms has to be performed? (see Section 4)

- What validation procedures are considered to evaluate the results of the experiment? (see Section 5)

- How will the results be analyzed? (see Section 6)

- How will the results be presented? (see Section 8)

- How are randomization and replicability of the experiment achieved? (see Section 9)

7.5   Tuning Before Benchmarking

Brownlee [2007] discusses the importance of tuning an algorithm before benchmarking. Bartz-Beielstein and Preuss [2010] state that comparisons of tuned versus untuned algorithms are not fair and should be avoided. During a benchmark study the employed parameter settings are extremely important as they largely define the obtained performance. Depending on the availability of code for the algorithms under scope and time for parameter searches, there are different possibilities to make a fair comparison:

- In the best case, the code for all methods is available. It is then possible to perform a parameter search for each problem and each algorithm via a tuning method. Taking the best parameter sets for each method for each problem ensures comparing the algorithms at their peak performance.

- If algorithm runs on the chosen problems take too long for a full tuning process, one may however perform a simple space-filling design on the parameter space, e.g., a LHD or a low-discrepancy point set [Matoušek, 2009] with only a few design points and repeats. This prevents misconfigurations of algorithms as one probably easily gets into the "ball park" [De Jong, 2007] of relatively good parameter settings. Most likely, neither algorithm works at its peak performance level, but the comparison is still fair.

- If no code other than for one's own algorithm is available, one has to resort to comparing with default parameter values. For a new algorithm, these could be determined by a common tuning process over the whole problem set. Note however, that such comparison deliberately abstains from setting good parameters for specific problems, even if this would be attempted for any real-world application.

---

[26]At the moment, this is only a list, which will be extended in forthcoming versions of this survey.

7.6 Open Issues

(O7.1) *Best Designs.*
Same authors consider LHDs as the default choice, even if for numerous applications a superiority of other space-filling or low-discrepancy designs has been demonstrated Santner et al. [2003]. The question when to prefer i.i.d. uniform sampling, LHDs, low-discrepancy point sets, other space-filling designs, or sets minimizing some other diversity criterion is largely open.

(O7.2) *Multiple Objectives.*
Sometimes properties of the objective function are used to determine the quality of a design. Therefore, it remains unclear how to measure the quality in settings where the objective function is unknown. Furthermore, problems occur if wrong assumptions about the objective function, e.g., linearity, are made. And, last but not least, in Multi-Objective Optimization (MOO), where no single objective can be specified, finding the optimal design can be very difficult [Santner et al., 2003].

8 How to Present Results?

8.1 General Recommendations

Several papers have been published in the last years, which give recommendations on how to report results. As Gent and Walsh [1994] already stated in 1994, after having generated some good results in your benchmark study, there are still many mistakes to make. They give the following recommendations:

1. present statistics, i.e., statements such as "algorithm a outperforms b" should be accompanied with suitable test results as described in Section 6,

2. do not push deadlines, i.e., do not reduce the quality of the report, because the deadline is approaching soon. Invest some time in planning: number of experiments, algorithm portfolio, hardness of the problem instances as discussed in Section 7,

3. and report negative results , i.e., present and discuss problem instances on which the algorithms fail (this is a key component of a good scientific report as discussed in this section).

Barr et al. [1995] in their classical work on reporting empirical results of heuristics specify a loose experimental setup methodology with the following steps:

1. define the goals of the experiment,

2. select measure of performance and factors to explore,

3. design and execute the experiment,

4. analyze the data and draw conclusions, and finally

5. report the experimental results.

They then suggest eight guidelines for reporting results, in summary they are; reproducibility, specify all influential factors (code, computing environment, etc.), be precise regarding measures, specify parameters, use statistical experimental design, compare with other methods, reduce variability of results, ensure results are comprehensive. They then go on to clarify these points with examples.

8.2 Reporting Methodologies

Besides recommendations, that provide valuable hints on how to report results, there exist also methodologies, which employ a scientific methodology, e.g., based on hypothesis testing [Popper, 1959, 1975]. Such a methodology was proposed by Bartz-Beielstein and Preuss [2010]. They propose organizing the presentation of experiments into seven parts, as follows:

(R.1) *Research question*
Briefly names the matter dealt with, the (possibly very general) objective, preferably in one sentence. This is used as the report's "headline" and related to the primary model.

(R.2) *Pre-experimental planning*
Summarizes the first—possibly explorative—program runs, leading to task and setup (R-3 and R-4). Decisions on employed benchmark problems or performance measures should be taken according to the data collected in preliminary runs. The report on pre-experimental planning should also include negative results, e.g., modifications to an algorithm that did not work or a test problem that turned out to be too hard, if they provide new insight.

(R.3) *Task*
Concretizes the question in focus and states scientific claims and derived statistical hypotheses to test. Note that one scientific claim may require several, sometimes hundreds, of statistical hypotheses. In case of a purely explorative study, as with the first test of a new algorithm, statistical tests may not be applicable. Still, the task should be formulated as precisely as possible. This step is related to the experimental model.

(R.4) *Setup*
Specifies problem design and algorithm design, including the investigated algorithm, the controllable and the fixed parameters, and the chosen performance measuring. It also includes information about the computational environment (hard- and software specification, e.g., the packages or libraries used). The information provided in this part should be sufficient to replicate an experiment.

(R.5) *Results/Visualization*
Gives raw or produced (filtered) data on the experimental outcome and additionally provides basic visualizations where meaningful. This is related to the data model.

(R.6) *Observations*
Describes exceptions from the expected, or unusual patterns noticed, without subjective assessment or explanation. As an example, it may be worthwhile to look at parameter interactions. Additional visualizations may help to clarify what happens.

(R.7) *Discussion*
Decides about the hypotheses specified in R-3, and provides necessarily subjective interpretations of the recorded observations. Also places the results in a wider context. The leading question here is: What did we learn?

This methodology was extended and refined in Preuss [2015]. It is important to divide parts R-6 and R-7, to facilitate different conclusions drawn by others, based on the same results/observations. This distinction into parts of increasing subjectiveness is similar to the suggestions of Barr et al. [1995], who distinguish between results, their analysis, and the conclusions drawn by the experimenter.

Note that all of these parts are already included in current good experimental reports. However, they are usually not separated but wildly mixed. Thus, we only suggest inserting labels into the text to make the structure more obvious.

We also recommend keeping a journal of experiments with single reports according to the above scheme to enable referring to previous experiments later on. This is useful even if single experiments do not find their way into a publication, as it improves the overview of subsequent experiments and helps to avoid repeated tests.

8.3 Open Issues

Reporting negative results has many benefits, e.g., to demonstrate what has been done and does not work, so someone else will not do the same in the future. And they are valuable tools for illustrating the limitations of new approaches. The presentation of negative results discussed above in 3 is not adequately accepted in the research community (cf. Gent and Walsh [1994]). Whereas a paper improving existing experimental results or outperforming another algorithm regularly gets accepted for publication, papers presenting negative results regularly will not.

9 How to Guarantee Reproducibility?

Reproducibility has been a topic of interest in the experimental analysis of algorithms for many decades. Classical works [Johnson, 2002a] advise *ensuring reproducibility*, but caution that the classical understanding of reproducibility in computer science, i.e., running exactly the same code on the same machine returns exactly the same measurements, differs substantially from the understanding in other experimental sciences, i.e., a different implementation of the experiment under similar conditions returns measurements that lead to the same conclusions.

For example, the "Reproducibility guidelines for AI research"[27] intended to be adopted by the Association for the Advancement of Artificial Intelligence (AAAI) are clearly focused on the concept of reproducibility in computer science.

Trying to clearly define various reproducibility concepts, the ACM distinguishes among:[28]

Repeatability (Same team, same experimental setup) The measurement can be obtained with stated precision by the same team using the same measurement procedure, the same measuring system, under the same operating conditions, in the same location on multiple trials. For computational experiments, this means that a researcher can reliably repeat her own computation.

Reproducibility (Different team, same experimental setup) The measurement can be obtained with stated precision by a different team using the same measurement procedure, the same measuring system, under the same operating conditions, in the same or a different location on multiple trials. For computational experiments, this means that an independent group can obtain the same result using the author's own artifacts.

Replicability (Different team, different experimental setup) The measurement can be obtained with stated precision by a different team, a different measuring system, in a different location on multiple trials. For computational experiments, this means that an independent group can obtain the same result using artifacts which they develop completely independently.

The above classification helps to identify various levels of reproducibility, reserving the term "Replicability" to the most scientifically useful, yet hardest to achieve. There are many practical guidelines and software systems available to achieve repeatibility and reproducibility [Gent et al., 1997, Johnson, 2002a], including code versioning tools (Subversion and Git), data repositories (Zenodo), reproducible documents (Rmarkdown and Jupyter notebooks), and reproducible software environments (OSF[29], CodeOcean and Docker).

Unfortunately, it is not so clear how to successfully achieve Replicability. For achieving replicability, one must give up on exactly reproducing the results and provide statistical guidelines that are commonly accepted by the field to provide sufficient evidence for a conclusion, even under different, but similar, experimental conditions. What constitutes *similar* experimental conditions depends on the experiment and there is no simple answer when benchmarking algorithms. One step towards better replicability is to pre-register experimental designs [Nosek et al., 2018] to fix the hypothesis and design of experiments. Preregistration

---

[27]http://folk.idi.ntnu.no/odderik/reproducibility_guidelines.pdf
[28]Quoting from:
https://www.acm.org/publications/policies/artifact-review-and-badging-current
[29]https://osf.io/

reduces the risk of spurious results due to adaptations to data analysis. However, it is much harder to systematically control for adaptive computational experiments because, unlike randomized controlled trials, they are much easier to run and re-run prior to registration.

10 Summary and Outlook

This survey compiles ideas and recommendations from more than a dozen researchers with different backgrounds and from different institutions around the world. Its main goal is the promotion of best practice in benchmarking. This version is the result of long and fruitful discussions among the authors. The authors agreed on eight essential topics, that should be considered in every benchmark study: goals, problems, algorithms, performance, analysis, design, presentation, and reproducibility. These topics defined the section structure of this article.

While it is definitely not a textbook that explains every single approaches in detail, we hope it is a good starting point for setting up benchmark studies. It is basically a guide (similar to the famous hitch-hiker's guide to EC [Heitkötter and Beasley, 1994]) and has a long list of references, which covers classical papers as well as the most recent ones. Every section presents recommendations, best practice examples, and open issues.

As mentioned above, this survey is only the beginning of a wonderful journey. It can serve as a starting point for many activities that improve the quality of benchmark studies and enhance the quality of research in EC and related fields. Next steps can be as follows:

1. offering tutorials and organizing workshops,

2. compiling videos, which explain how to set up the experiments, analyze results, and report important findings,

3. providing software tools,

4. developing a comprehensible check-list, especially for beginners in benchmarking,

5. including a discussion section in every section, which describes controversial topics and ideas.

Our final goal is to provide well-accepted guidelines (rules) that might be useful for authors, reviewers, and others. Consider the following (rudimentary and incomplete) checklist, that can serve as a guideline for authors and reviewers:

1. goals: did the authors clearly state the reasons for this study?

2. problems: is the selection of problem instances well motivated and justified?

3. algorithms: do comparisons include relevant competitors?

4. performance: is the choice of the performance measure adequate?

5. analysis: are standards from statistics considered?

6. design: does the experimental setup enable efficient and fair experimentation? What measures are taken to avoid "cherry-picking results"?

7. presentation: are the results well organized and explained?

8. reproducibility: data and code availability?

Transparent, well accepted standards will improve the review process in EC and related fields significantly. These common standards might also accelerate the review process, because they improve the quality of submissions and helps reviewers to write objective evaluations. Most importantly, it is not our intention to dictate specific test statistics, experimental designs, or performance measures. Instead, we claim that publications in EC would improve, if authors explain, *why* they have chosen this specific measure, tool, or design. And, last but not least, authors should describe the goal of their study.

Although we tried to include the most relevant contributions, we are aware that important contributions are missing. Because the acceptance of the proposed recommendations is crucial, we would like to invite more researchers to share their knowledge with us. Moreover, as the field of benchmarking is constantly changing, this article will be regularly updated and published on arXiv [Bartz-Beielstein et al., 2020]. To get in touch, interested readers can use the associated e-mail address for this project: benchmarkingbest-practice@gmail.com.

There are several other initiatives that are trying to improve benchmarking standards in query-based optimization fields, e.g., the Benchmarking Network[30], an initiative built to consolidate and to stimulate activities on benchmarking iterative optimization heuristics [Weinand et al., 2020].

In our opinion, starting and maintaining this public discussion is very important. Maybe, this survey poses more questions than answers, which is fine. Therefore, we conclude this article with a famous saying that is attributed to Richard Feynman[31]:

> I would rather have questions that can't be answered than answers that can't be questioned.

Acknowledgments

---

[30]https://sites.google.com/view/benchmarking-network/

[31]https://en.wikipedia.org/w/index.php?title=Talk:Richard_Feynman&oldid=2681873#%22I_would_rather_have_questions_that_cannot_be_answered%22

[32]https://www.dagstuhl.de/19431

Glossary

AAAI Association for the Advancement of Artificial Intelligence. 34

ACM Association for Computing Machinery. 14, 34

ASlib Algorithm Selection Library. 16

BBOB Black-Box-Optimization-Benchmarking. 12, 14

BFGS Broyden-Fletcher-Goldfarb-Shanno. 17

CEC Congress on Evolutionary Computation. 12, 14–16

CMA-ES Covariance Matrix Adaptation Evolution Strategy. 6, 17

COCO Comparing Continuous Optimizers. 4, 11

CRN Common Random Numbers. 26

CRS4EA Chess Rating Systems for Evolutionary Algorithms. 28, 29

DACE Design and Analysis of Computer Experiments. 29, 30

DOE Design of Experiments. 5, 9, 29, 30

EA Evolutionary Algorithm. 14, 15, 17

EC Evolutionary Computation. 4, 5, 11, 14, 20, 35, 36

EDA Exploratory Data Analysis. 23–25

EDAlgo Estimation of Distribution Algorithm. 17

EGO Efficient Global Optimization. 17

ELA Exploratory Landscape Analysis. 12, 23

ERT Expected Running Time. 22

ES Evolution Strategy. 30

FE Function Evaluation. 20, 21

GECCO Genetic and Evolutionary Computation Conference. 14, 15

IEEE Institute of Electrical and Electronics Engineers. 14

irace iterated racing. 18

LABS Low Auto-correlation Binary Sequence. 12

LHD Latin Hypercube Design. 17, 31, 32

MACODA Many Criteria Optimization and Decision Analysis. 16

MAX-SAT Maximum Satisfiability. 20

References

Stavros P Adam, Stamatios-Aggelos N Alexandropoulos, Panos M Pardalos, and Michael N Vrahatis. No Free Lunch Theorem: A Review. In *Approximation and Optimization*, pages 57 – 82. Springer, 2019.

Amritanshu Agrawal, Tim Menzies, Leandro L. Minku, Markus Wagner, and Zhe Yu. Better software analytics via "duo": Data mining algorithms using/used-by optimizers. *Empirical Software Engineering*, 25(3):2099–2136, 2020. doi:10.1007/s10664-020-09808-9. URL https://doi.org/10.1007/s10664-020-09808-9.

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631, 2019.

Mohammad M Amini and Richard S Barr. Network reoptimization algorithms: A statistically designed comparison. *ORSA Journal on Computing*, 5(4):395–409, 1993.

Theodore W Anderson and Donald A Darling. Asymptotic theory of certain" goodness of fit" criteria based on stochastic processes. *The annals of mathematical statistics*, pages 193–212, 1952.

Carlos Ansótegui, Yuri Malitsky, Horst Samulowitz, Meinolf Sellmann, and Kevin Tierney. Model-based genetic algorithms for algorithm configuration. In *Proc. of International Conf. on Artificial Intelligence (IJCAI'15)*, pages 733–739. AAAI, 2015.

Dirk V. Arnold. *Noisy Optimization with Evolution Strategies*, volume 8. Springer, 2012.

Anne Auger and Benjamin Doerr. *Theory of Randomized Search Heuristics*. World Scientific, 2011.

Anne Auger and Nikolaus Hansen. Performance evaluation of an advanced local search evolutionary algorithm. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1777–1784. IEEE, 2005.

Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz. *Handbook of Evolutionary Computation*. IOP Publishing Ltd., GBR, 1st edition, 1997.

Richard S Barr, Bruce L Golden, James P Kelly, Mauricio GC Resende, and William R Stewart. Designing and reporting on computational experiments with heuristic methods. *Journal of Heuristics*, 1(1):9–32, 1995.

Maurice Stevenson Bartlett. Properties of sufficiency and statistical tests. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, 160(901):268–282, 1937.

Thomas Bartz-Beielstein. *Experimental Research in Evolutionary Computation—The New Experimentalism*. Natural Computing Series. Springer, 2006.

Thomas Bartz-Beielstein and Mike Preuss. The Future of Experimental Research. In Thomas Bartz-Beielstein, Marco Chiarandini, Luis Paquete, and Mike Preuss, editors, *Experimental Methods for the Analysis of Optimization Algorithms*, pages 17–46. Springer, Berlin, Heidelberg, New York, 2010.

Thomas Bartz-Beielstein, Konstantinos E Parsopoulos, and Michael N Vrahatis. Design and analysis of optimization algorithms using computational statistics. *Applied Numerical Analysis & Computational Mathematics*, 1(2):413–433, 2004.

Thomas Bartz-Beielstein, Christian WG Lasarczyk, and Mike Preuss. Sequential Parameter Optimization. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, volume 1, pages 773 – 780. IEEE, 2005.

Thomas Bartz-Beielstein, Marco Chiarandini, Luís Paquete, and Mike Preuss. *Experimental methods for the analysis of optimization algorithms*. Springer, 2010.

Thomas Bartz-Beielstein, Lorenzo Gentile, and Martin Zaefferer. In a nutshell: Sequential parameter optimization. Technical report, TH Köln, 2017.

Thomas Bartz-Beielstein, Carola Doerr, Jakob Bossek, Sowmya Chandrasekaran, Tome Eftimov, Andreas Fischbach, Pascal Kerschke, Manuel Lopez-Ibanez, Katherine M. Malan, Jason H. Moore, Boris Naujoks, Patryk Orzechowski, Vanessa Volz, Markus Wagner, and Thomas Weise. Benchmarking in Optimization: Best Practice and Open Issues. *arXiv e-prints*, art. arXiv:2007.03488, July 2020.

Vahid Beiranvand, Warren Hare, and Yves Lucet. Best Practices for Comparing Optimization Algorithms. *Optimization and Engineering*, 18(4):815 – 848, 2017.

James Bergstra, Dan Yamins, and David D Cox. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in science conference*, volume 13, page 20. Citeseer, 2013.

Hans-Georg Beyer and Hans-Paul Schwefel. Evolution strategies: A comprehensive introduction. *Natural Computing*, 1:3–52, 2002.

Mauro Birattari, Thomas Stützle, Luís Paquete, and Klaus Varrentrapp. A racing algorithm for configuring metaheuristics. In W. B. Langdon et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2002*, pages 11–18. Morgan Kaufmann Publishers, San Francisco, CA, 2002.

Mauro Birattari, Luis Paquete, and Thomas Stützle. Classification of metaheuristics and design of experiments for the analysis of components. https://www.researchgate.net/publication/2557723_Classification_of_Metaheuristics_and_Design_of_Experiments_for_the_Analysis_of_Components, 2003. technical report.

Bernd Bischl, Pascal Kerschke, Lars Kotthoff, Thomas Marius Lindauer, Yuri Malitsky, Alexandre Fréchette, Holger H. Hoos, Frank Hutter, Kevin Leyton-Brown, Kevin Tierney, and Joaquin Vanschoren. ASlib: A Benchmark Library for Algorithm Selection. *Artificial Intelligence (AIJ)*, 237:41 – 58, 2016.

Laurens Bliek, Sicco Verwer, and Mathijs de Weerdt. Black-box mixed-variable optimisation using a surrogate model that satisfies integer constraints. *arXiv preprint arXiv:2006.04508*, 2020.

Mahmoud A. Bokhari, Brad Alexander, and Markus Wagner. Towards rigorous validation of energy optimisation experiments. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, GECCO '20, page 1232–1240, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450371285. doi:10.1145/3377930.3390245. URL https://doi.org/10.1145/3377930.3390245.

Mohammad Reza Bonyadi, Zbigniew Michalewicz, and Luigi Barone. The Travelling Thief Problem: The First Step in the Transition from Theoretical Problems to Realistic Problems. In *2013 IEEE Congress on Evolutionary Computation*, pages 1037 – 1044. IEEE, 2013.

Mohammad Reza Bonyadi, Zbigniew Michalewicz, Markus Wagner, and Frank Neumann. Evolutionary Computation for Multicomponent Problems: Opportunities and Future Directions. In *Optimization in Industry: Present Practices and Future Scopes*, pages 13 – 30. Springer, 2019. doi:10.1007/978-3-030-01641-8_2.

Jakob Bossek, Pascal Kerschke, Aneta Neumann, Markus Wagner, Frank Neumann, and Heike Trautmann. Evolving Diverse TSP Instances by Means of Novel and Creative Mutation Operators. In *Proc. of the 15th ACM/SIGEVO Conference on Foundations of Genetic Algorithms*, pages 58 – 71. ACM, 2019.

Jakob Bossek, Pascal Kerschke, and Heike Trautmann. A Multi-Objective Perspective on Performance Assessment and Automated Selection of Single-Objective Optimization Algorithms. *Applied Soft Computing Journal (ASOC)*, 88:105901, March 2020a.

Jakob Bossek, Pascal Kerschke, and Heike Trautmann. Anytime Behavior of Inexact TSP Solvers and Perspectives for Automated Algorithm Selection. In *Proc. of the IEEE Congress on Evolutionary Computation*. IEEE, 2020b. A preprint of this manuscript can be found at https://arxiv.org/abs/2005.13289.

Ilhem Boussaïd, Julien Lepagnot, and Patrick Siarry. A survey on optimization metaheuristics. *Information Sciences*, 237:82–117, 2013. doi:10.1016/j.ins.2013.02.041. URL https://doi.org/10.1016/j.ins.2013.02.041.

Jürgen Branke. Creating Robust Solutions by Means of Evolutionary Algorithms. In *International Conference on Parallel Problem Solving from Nature*, pages 119 – 128. Springer, 1998.

Jürgen Branke, Christian Schmidt, and Hartmut Schmeck. Efficient Fitness Estimation in Noisy Environments. In *Genetic and Evolutionary Computation Conference (GECCO'01)*, pages 243 – 250. Morgan Kaufmann, 2001.

Jason Brownlee. A Note on Research Methodology and Benchmarking Optimization Algorithms. Technical report, Complex Intelligent Systems Laboratory (CIS), Centre for Information Technology Research (CITR), Faculty of Information and Communication Technologies (ICT), Swinburne University of Technology, Victoria, Australia, Technical Report ID 70125, 2007.

Eduardo G Carrano, Elizabeth F Wanner, and Ricardo HC Takahashi. A multicriteria statistical based comparison methodology for evaluating evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 15(6): 848–870, 2011.

Marie-Liesse Cauwet and Olivier Teytaud. Noisy Optimization: Fast Convergence Rates with Comparison-based Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 1101 – 1106, 2016.

Pete Chapman, Julian Clinton, Randy Kerber, Thomas Khabaza, Thomas Reinartz, Colin Shearer, and Rüdiger Wirth. CRISP-DM 1.0: Step-by-Step Data Mining Guide. Technical report, SPSS Inc., 2000.

Marco Chiarandini, Luis Paquete, Mike Preuss, and Enda Ridge. Experiments on Metaheuristics: Methodological Overview and Open Issues. Technical report, Institut for Matematik og Datalogi Syddansk Universitet, 2007.

Nicos Christofides. The Vehicle Routing Problem. *Revue française d'automatique, d'informatique et de recherche opérationnelle (RAIRO). Recherche opérationnelle*, 10(1):55 – 70, 1976. URL http://www.numdam.org/item?id=RO_1976__10_1_55_0.

Matej Črepinšek, Shih-Hsi Liu, and Marjan Mernik. Replication and Comparison of Computational Experiments in Applied Evolutionary Computing: Common Pitfalls and Guidelines to Avoid Them. *Applied Soft Computing*, 19: 161 – 170, June 2014.

Harlan Crowder, Ron S Dembo, and John M Mulvey. On reporting computational experiments with mathematical software. *ACM Transactions on Mathematical Software (TOMS)*, 5(2):193–203, 1979.

Joseph C. Culberson. On the futility of blind search: An algorithmic view of "no free lunch". *Evolutionary Computation*, 6(2):109–127, 1998. doi:10.1162/evco.1998.6.2.109.

Kenneth De Jong. Parameter Setting in EAs: a 30 Year Perspective. In *Parameter Setting in Evolutionary Algorithms*, pages 1 – 18. Springer, 2007.

Marleen De Jonge and Daan van den Berg. Parameter sensitivity patterns in the plant propagation algorithm. In *IJCCI*, page 92–99, 2020.

Lucas Augusto Müller de Souza, José Eduardo Henriques da Silva, Luciano Jerez Chaves, and Heder Soares Bernardino. A benchmark suite for designing combinational logic circuits via metaheuristics. *Applied Soft Computing*, 91:106246, June 2020. doi:10.1016/j.asoc.2020.106246. URL https://doi.org/10.1016/j.asoc.2020.106246.

Joaquín Derrac, Salvador García, Daniel Molina, and Francisco Herrera. A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms. *Swarm and Evolutionary Computation*, 1(1):3–18, 2011.

Joaquín Derrac, Salvador García, Sheldon Hui, Ponnuthurai Nagaratnam Suganthan, and Francisco Herrera. Analyzing convergence performance of evolutionary algorithms: A statistical approach. *Information Sciences*, 289:41–58, 2014.

Jay L Devore. *Probability and Statistics for Engineering and the Sciences*. Cengage learning, 2011.

Benjamin Doerr and Frank Neumann. *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*. Springer, 2020.

Benjamin Doerr, Carola Doerr, and Johannes Lengler. Self-adjusting mutation rates with provably optimal success rules. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1479 – 1487. ACM, 2019.

Carola Doerr, Hao Wang, Furong Ye, Sander van Rijn, and Thomas Bäck. IOHprofiler: A Benchmarking and Profiling Tool for Iterative Optimization Heuristics. *arXiv e-prints*, art. arXiv:1810.05281, Oct 2018. Wiki page of IOHprofiler is available at https://iohprofiler.github.io/.

Carola Doerr, Furong Ye, Naama Horesh, Hao Wang, Ofer M. Shir, and Thomas Bäck. Benchmarking Discrete Optimization Heuristics with IOHprofiler. *Applied Soft Computing*, 88:106027, 2020.

Elizabeth D Dolan and Jorge J Moré. Benchmarking optimization software with performance profiles. *Mathematical programming*, 91(2):201–213, 2002.

Marco Dorigo, Mauro Birattari, and Thomas Stützle. Ant colony optimization: Artificial ants as a computational intelligence technique. *IEEE Computational Intelligence Magazine*, 1(4):28–39, 2006.

Gunter Dueck and Tobias Scheuer. Threshold accepting: a general purpose optimization algorithm appearing superior to simulated annealing. *J. Comput. Phys.*, 90:161–175, 1990.

Tome Eftimov and Peter Korošec. The impact of statistics for benchmarking in evolutionary computation research. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1329–1336, 2018.

Tome Eftimov and Peter Korošec. Identifying practical significance through statistical comparison of meta-heuristic stochastic optimization algorithms. *Applied Soft Computing*, 85:105862, 2019.

Tome Eftimov, Peter Korošec, and Barbara Koroušić Seljak. A Novel Approach to Statistical Comparison of Meta-Heuristic Stochastic Optimization Algorithms Using Deep Statistics. *Information Sciences*, 417:186 – 215, 2017.

Tome Eftimov, Gašper Petelin, and Peter Korošec. Dsctool: A web-service-based framework for statistical comparison of stochastic optimization algorithms. *Applied Soft Computing*, 87:105977, 2020.

Katharina Eggensperger, Matthias Feurer, Frank Hutter, James Bergstra, Jasper Snoek, Holger H. Hoos, and Kevin Leyton-Brown. Towards an Empirical Foundation for Assessing Bayesian Optimization of Hyperparameters. In *NIPS Workshop on Bayesian Optimization in Theory and Practice*, volume 10, December 2013.

Ágoston Endre Eiben and Márk Jelasity. A Critical Note on Experimental Research Methodology in EC. In *Proceedings of the 2002 IEEE Congress on Evolutionary Computation*, volume 1, pages 582 – 587. IEEE, 2002.

Ágoston Endre Eiben and Selmar K Smit. Evolutionary Algorithm Parameters and Methods to Tune Them. In *Autonomous search*, pages 15 – 36. Springer, 2011.

Ágoston Endre Eiben and James E Smith. *Introduction to Evolutionary Computing*. Natural Computing. Springer, 2 edition, 2015.

Álvaro Fialho, Luís Da Costa, Marc Schoenauer, and Michèle Sebag. Analyzing bandit-based adaptive operator selection mechanisms. *Annals of Mathematics and Artificial Intelligence*, 60:25 – 64, 2010.

Steffen Finck, Nikolaus Hansen, Raymond Ros, and Anne Auger. COCO Documentation, Release 15.03, November 2015. URL http://coco.lri.fr/COCOdoc/COCO.pdf.

Andreas Fischbach and Thomas Bartz-Beielstein. Improving the reliability of test functions generators. *Applied Soft Computing*, 92:106315, 2020. ISSN 1568-4946. doi:https://doi.org/10.1016/j.asoc.2020.106315. URL http://www.sciencedirect.com/science/article/pii/S1568494620302556.

Philip J. Fleming and John J. Wallace. How not to lie with statistics: The correct way to summarize benchmark results. *Commun. ACM*, 29(3):218–221, March 1986.

Roger Fletcher. Conjugate gradient methods for indefinite systems. In *Numerical analysis*, pages 73–89. Springer, 1976.

Alexandre Fréchette, Lars Kotthoff, Tomasz Michalak, Talal Rahwan, Holger H. Hoos, and Kevin Leyton-Brown. Using the shapley value to analyze algorithm portfolios. In *Proc. of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3397 —- 3403. AAAI Press, 2016.

Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association*, 32(200):675–701, 1937.

Marcus Gallagher. Towards improved benchmarking of black-box optimization algorithms using clustering problems. *Soft Computing*, 20(10):3835–3849, March 2016. doi:10.1007/s00500-016-2094-1. URL https://doi.org/10.1007/s00500-016-2094-1.

Salvador García, Daniel Molina, Manuel Lozano, and Francisco Herrera. A study on the use of non-parametric tests for analyzing the evolutionary algorithms' behaviour: a case study on the cec'2005 special session on real parameter optimization. *Journal of Heuristics*, 15(6):617, 2009.

Carlos García-Martínez, Francisco J. Rodríguez, and Manuel Lozano. Arbitrary function optimisation with meta-heuristics: No free lunch and real-world problems. *Soft Computing*, 16(12):2115–2133, 2012. doi:10.1007/s00500-012-0881-x.

Robert W. Garden and Andries P. Engelbrecht. Analysis and classification of optimisation benchmark functions and benchmark suites. In *2014 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, July 2014. doi:10.1109/cec.2014.6900240. URL https://doi.org/10.1109/cec.2014.6900240.

Ian P. Gent and Toby Walsh. How not to do it. In *AAAI Workshop on Experimental Evaluation of Reasoning and Search Methods*, 1994.

Ian P. Gent, Stuart A. Grant, Ewen MacIntyre, Patrick Prosser, Paul Shaw, Barbara M. Smith, and Toby Walsh. How not to do it. Technical Report 97.27, School of Computer Studies, University of Leeds, May 1997.

Fred Glover. Tabu search—part i. *ORSA Journal on computing*, 1(3):190–206, 1989.

Sim Kuan Goh, Kay Chen Tan, Abdullah Al-Mamun, and Hussein A. Abbass. Evolutionary big optimization (BigOpt) of signals. In *2015 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, May 2015. doi:10.1109/cec.2015.7257307. URL https://doi.org/10.1109/cec.2015.7257307.

David E Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading MA, 1989.

Bruce L Golden, Arjang A Assad, Edward A Wasil, and Edward Baker. Experimentation in optimization. *European Journal of Operational Research*, 27(1):1–16, 1986.

Raphael T. Haftka. Requirements for papers focusing on new or improved global optimization algorithms. *Structural and Multidisciplinary Optimization*, 54(1):1–1, 2016.

Doug Hains, L. Darrell Whitley, Adele E. Howe, and Wenxiang Chen. Hyperplane Initialized Local Search for MAXSAT. In *Proc. of the Genetic and Evolutionary Computation Conference*, pages 805 – 812. ACM, 2013.

Nikolaus Hansen. Invariance, self-adaptation and correlated mutations in evolution strategies. In *Proc. of International Conference on Parallel Problem Solving from Nature*, pages 355–364. Springer, 2000. ISBN 978-3-540-45356-7.

Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). *Evolutionary computation*, 11(1):1–18, 2003.

Nikolaus Hansen, Anne Auger, Steffen Finck, and Raymond Ros. Real-parameter black-box optimization benchmarking: Experimental setup. Technical report, Université Paris Sud, INRIA Futurs, Équipe TAO, Orsay, France, March 24, 2012. URL http://coco.lri.fr/BBOB-downloads/download11.05/bbobdocexperiment.pdf.

Nikolaus Hansen, Anne Auger, Dimo Brockhoff, Dejan Tušar, and Tea Tušar. COCO: performance assessment. *CoRR*, abs/1605.03560, 2016a. URL http://arxiv.org/abs/1605.03560.

Nikolaus Hansen, Anne Auger, Olaf Mersmann, Tea Tušar, and Dimo Brockhoff. COCO: A Platform for Comparing Continuous Optimizers in a Black-Box Setting. *arXiv preprint*, abs/1603.08785v3, August 2016b. URL http://arxiv.org/abs/1603.08785v3.

Nikolaus Hansen, Youhei Akimoto, yoshihikoueno, Dimo Brockhoff, and Matthew Chan. CMA-ES/pycma: r3.0.3, 2020. URL https://doi.org/10.5281/zenodo.3764210.

Anna Hart. Mann-whitney test is not just a test of medians: differences in spread can be important. *Bmj*, 323(7309):391–393, 2001.

Jörg Heitkötter and David Beasley. The hitch-hiker's guide to evolutionary computation, 1994.

Michael Hellwig and Hans-Georg Beyer. Benchmarking Evolutionary Algorithms For Single-Objective Real-valued Constrained Optimization – A Critical Review. *Swarm and Evolutionary Computation*, 44:927–944, 2019.

John N Hooker. Needed: An Empirical Science of Algorithms. *Operations research*, 42(2):201 – 212, 1994.

John N. Hooker. Testing heuristics: We have it all wrong. *Journal of Heuristics*, 1(1):33–42, 1996. doi:10.1007/BF02430364.

Frank Hutter, Holger H. Hoos, Kevin Leyton-Brown, and Thomas Stützle. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36:267–306, October 2009.

Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Sequential Model-Based Optimization for General Algorithm Configuration. In *International Conference on Learning and Intelligent Optimization*, pages 507 – 523. Springer, 2011.

Frank Hutter, Manuel López-Ibáñez, Chris Fawcett, Marius Thomas Lindauer, Holger H. Hoos, Kevin Leyton-Brown, and Thomas Stützle. AClib: a benchmark library for algorithm configuration. In Panos M. Pardalos, Mauricio G. C. Resende, Chrysafis Vogiatzis, and Jose L. Walteros, editors, *Learning and Intelligent Optimization, 8th International Conference, LION 8*, volume 8426 of *Lecture Notes in Computer Science*, pages 36–40. Springer, Heidelberg, Germany, 2014. doi:10.1007/978-3-319-09584-4_4.

Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. *Automated Machine Learning: Methods, Systems, Challenges*. Springer, 2019.

Alexandre D. Jesus, Arnaud Liefooghe, Bilel Derbel, and Luís Paquete. Algorithm Selection of Anytime Algorithms. In *Proc. of the 2020 Genetic and Evolutionary Computation Conference*, pages 850 – 858. ACM, 2020.

Yaochu Jin and Jürgen Branke. Evolutionary Optimization in Uncertain Environments – A Survey. *IEEE Transactions on Evolutionary Computation*, 9(3):303 – 317, 2005.

David S Johnson, Cecilia R Aragon, Lyle A McGeoch, and Catherine Schevon. Optimization by Simulated Annealing: An Experimental Evaluation. Part I, Graph Partitioning. *Operations Research*, 37(6):865 – 892, 1989.

David S Johnson, Cecilia R Aragon, Lyle A McGeoch, and Catherine Schevon. Optimization by Simulated Annealing: An Experimental Evaluation. Part II, Graph Coloring and Number Partitioning. *Operations Research*, 39(3):378 – 406, 1991.

David Stifler Johnson. A Theoretician's Guide to the Experimental Analysis of Algorithms. In *Proc. of a DIMACS Workshop on Data Structures, Near Neighbor Searches, and Methodology: Fifth and Sixth DIMACS Implementation Challenges*, volume 59 of *DIMACS – Series in Discrete Mathematics and Theoretical Computer Science*, pages 215–250, 2002a.

David Stifler Johnson. A Theoretician's Guide to the Experimental Analysis of Algorithms. *Data Structures, Near Neighbor Searches, and Methodology: Fifth and Sixth DIMACS Implementation Challenges*, 59:215 – 250, 2002b.

David Stifler Johnson and Lyle A. McGeoch. Experimental Analysis of Heuristics for the STSP. In *The Traveling Salesman Problem and its Variations*, volume 12 of *Combinatorial Optimization*, chapter 9, pages 369 – 443. Kluwer Academic Publishers, 2002.

Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.

Giorgos Karafotias, Mark Hoogendoorn, and Ágoston E. Eiben. Parameter control in evolutionary algorithms: Trends and challenges. *IEEE Transactions on Evolutionary Computation*, 19(2):167–187, April 2015.

Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, USA, 1993.

James Kennedy and Russell Eberhart. Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948. IEEE, 1995.

Pascal Kerschke and Heike Trautmann. Automated Algorithm Selection on Continuous Black-Box Problems by Combining Exploratory Landscape Analysis and Machine Learning. *Evolutionary Computation (ECJ)*, 27(1): 99 – 127, 2019a.

Pascal Kerschke and Heike Trautmann. Comprehensive Feature-Based Landscape Analysis of Continuous and Constrained Optimization Problems Using the R-package flacco. In *Applications in Statistical Computing*, pages 93 – 123. Springer, 2019b.

Pascal Kerschke, Jakob Bossek, and Heike Trautmann. Parameterization of State-of-the-Art Performance Indicators: A Robustness Study based on Inexact TSP Solvers. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1737 – 1744. ACM, 2018a.

Pascal Kerschke, Lars Kotthoff, Jakob Bossek, Holger H. Hoos, and Heike Trautmann. Leveraging TSP Solver Complementarity through Machine Learning. *Evolutionary Computation (ECJ)*, 26(4):597 – 620, December 2018b.

Pascal Kerschke, Holger H. Hoos, Frank Neumann, and Heike Trautmann. Automated Algorithm Selection: Survey and Perspectives. *Evolutionary Computation (ECJ)*, 27:3 – 45, 2019.

Scott Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

Jack P. C. Kleijnen. Analyzing Simulation Experiments with Common Random Numbers. *Management Science*, 34 (1):65 – 74, 1988.

Jack P. C. Kleijnen. Experimental Design for Sensitivity Analysis of Simulation Models. Workingpaper, Operations Research, 2001.

Jack P. C. Kleijnen. Design and Analysis of Simulation Experiments. In *International Workshop on Simulation*, pages 3–22. Springer, 2015.

Jack P. C. Kleijnen. Regression and Kriging Metamodels with their Experimental Designs in Simulation: A Review. *European Journal of Operational Research*, 256(1):1–16, 2017.

Lasse Kliemann and Peter Sanders. *Algorithm Engineering: Selected Results and Surveys*, volume 9220. Springer, 2016.

William H Kruskal and W Allen Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

Abhishek Kumar, Guohua Wu, Mostafa Z. Ali, Rammohan Mallipeddi, Ponnuthurai Nagaratnam Suganthan, and Swagatam Das. Guidelines for real-world single-objective constrained optimisation competition. Technical report, 2020.

Pedro Larrañaga and José A. Lozano. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*. Kluwer Academic Publishers, 2002.

Per Kristian Lehre and Carsten Witt. Black-box search by unbiased variation. *Algorithmica*, 64:623–642, 2012.

Howard Levene. Robust tests for equality of variances. *Contributions to probability and statistics. Essays in honor of Harold Hotelling*, pages 279–292, 1961.

Lisha Li, Kevin G. Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research*, 18:185:1–185:52, 2017.

Rui Li, Michael T. M. Emmerich, Jeroen Eggermont, Ernst G. P. Bovenkamp, Thomas Bäck, Jouke Dijkstra, and Johan H. C. Reiber. Mixed-Integer NK Landscapes. In *Proc. of Parallel Problem Solving from Nature*, pages 42–51. Springer, 2006.

Tianjun Liao, Krzysztof Socha, Marco A. Montes de Oca, Thomas Stützle, and Marco Dorigo. Ant colony optimization for mixed-variable optimization problems. *IEEE Transactions on Evolutionary Computation*, 18(4):503–518, 2014.

Harold R Lindman. *Analysis of variance in complex experimental designs.* WH Freeman & Co, 1974.

Jialin Liu, Antoine Moreau, Mike Preuss, Baptiste Rozière, Jérémy Rapin, Fabien Teytaud, and Olivier Teytaud. Versatile black-box optimization. *CoRR*, abs/2004.14014, 2020. URL https://arxiv.org/abs/2004.14014.

Qunfeng Liu, William V. Gehrlein, Ling Wang, Yuan Yan, Yingying Cao, Wei Chen, and Yun Li. Paradoxes in Numerical Comparison of Optimization Algorithms. *IEEE Transactions on Evolutionary Computation*, pages 1–15, 2019.

Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Thomas Stützle, and Mauro Birattari. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58, 2016. doi:10.1016/j.orp.2016.09.002.

Katherine Mary Malan and Andries Petrus Engelbrecht. A Survey of Techniques for Characterising Fitness Landscapes and Some Possible Ways Forward. *Information Sciences (JIS)*, 241:148 – 163, 2013.

Jiří Matoušek. *Geometric Discrepancy*. Springer, Berlin, 2 edition, 2009.

Deborah G Mayo and Aris Spanos. Severe testing as a basic concept in a neyman–pearson philosophy of induction. *The British Journal for the Philosophy of Science*, 57(2):323–357, 2006.

Kent McClymont and Ed Keedwell. Benchmark Multi-Objective Optimisation Test Problems with Mixed Encodings. In *Proceedings of the 2011 IEEE Congress on Evolutionary Computation*, pages 2131 – 2138. IEEE, 2011.

James McDermott. When and why metaheuristics researchers can ignore "no free lunch" theorems. *SN Computer Science*, 1(60):1–18, 2020. doi:10.1007/s42979-020-0063-3.

Catherine C McGeoch. *Experimental Analysis of Algorithms*. PhD thesis, Carnegie Mellon University, Pittsburgh PA, 1986.

Catherine C McGeoch. Toward an experimental method for algorithm simulation. *INFORMS Journal on Computing*, 8(1):1–15, 1996.

McKay, Michael D and Beckman, Richard J and Conover, William J. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics*, 42(1):55 – 61, 2000.

Olaf Mersmann, Mike Preuss, and Heike Trautmann. Benchmarking Evolutionary Algorithms: Towards Exploratory Landscape Analysis. In *Proc. of International Conference on Parallel Problem Solving from Nature*, pages 73 – 82. Springer, 2010.

Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. Exploratory Landscape Analysis. In *Proc. of the 13th Annual Conference on Genetic and Evolutionary Computation*, pages 829 – 836. ACM, 2011.

Olaf Mersmann, Bernd Bischl, Heike Trautmann, Markus Wagner, Jakob Bossek, and Frank Neumann. A Novel Feature-Based Approach to Characterize Algorithm Performance for the Traveling Salesperson Problem. *Annals of Mathematics and Artificial Intelligence*, 69(2):151 – 182, 2013.

Nenad Mladenović and Pierre Hansen. Variable neighborhood search. *Comput. Oper. Res.*, 24(11):1097–1100, 1997. ISSN 0305-0548. doi:10.1016/S0305-0548(97)00031-2. URL https://doi.org/10.1016/S0305-0548(97)00031-2.

Douglas C Montgomery. *Design and Analysis of Experiments*. John Wiley & Sons, 9 edition, 2017.

Jorge J Moré and Stefan M Wild. Benchmarking derivative-free optimization algorithms. *SIAM Journal on Optimization*, 20(1):172–191, 2009.

Jorge J Moré, Burton S Garbow, and Kenneth E Hillstrom. Testing Unconstrained Optimization Software. *ACM Transactions on Mathematical Software (TOMS)*, 7(1):17 – 41, 1981.

Mario Andrés Muñoz Acosta, Michael Kirley, and Saman K. Halgamuge. Exploratory Landscape Analysis of Continuous Space Optimization Problems Using Information Content. *IEEE Transactions on Evolutionary Computation (TEVC)*, 19(1):74 – 87, 2015a.

Mario Andrés Muñoz Acosta, Yuan Sun, Michael Kirley, and Saman K. Halgamuge. Algorithm Selection for Black-Box Continuous Optimization Problems: A Survey on Methods and Challenges. *Information Sciences (JIS)*, 317: 224 – 245, 2015b.

H. Mühlenbein and G. Paaß. From recombination of genes to the estimation of distributions i. binary parameters. In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, *Proc. of Parallel Problem Solving from Nature (PPSN'96)*, pages 178–187. Springer, 1996. ISBN 978-3-540-70668-7.

Matthias Müller-Hannemann and Stefan Schirra. *Algorithm Engineering: Bridging the Gap Between Algorithm Theory and Practice.* Springer, 2010.

Mario A. Muñoz and Kate A. Smith-Miles. Performance Analysis of Continuous Black-Box Optimization Algorithms via Footprints in Instance Space. *Evolutionary Computation*, 25(4):529–554, December 2017.

Dima Nazzal, Mansooreh Mollaghasemi, H Hedlund, and A Bozorgi. Using Genetic Algorithms and an Indifference-Zone Ranking and Selection Procedure Under Common Random Numbers for Simulation Optimisation. *Journal of Simulation*, 6(1):56 – 66, 2012.

John A Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

Aneta Neumann, Wanru Gao, Markus Wagner, and Frank Neumann. Evolutionary diversity optimization using multi-objective indicators. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '19, page 837–845, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361118. doi:10.1145/3321707.3321796. URL https://doi.org/10.1145/3321707.3321796.

Frank Neumann and Carsten Witt. *Bioinspired Computation in Combinatorial Optimization – Algorithms and Their Computational Complexity.* Springer, 2010.

Trung Thanh Nguyen, Shengxiang Yang, and Juergen Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and Evolutionary Computation*, 6:1–24, October 2012. doi:10.1016/j.swevo.2012.05.001. URL https://doi.org/10.1016/j.swevo.2012.05.001.

Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11):2600–2606, March 2018. ISSN 0027-8424, 1091-6490. doi:10.1073/pnas.1708274114.

Randal S Olson and Jason H Moore. TPOT: A tree-based pipeline optimization tool for automating machine learning. In *Workshop on automatic machine learning*, pages 66–74, 2016.

Randal S Olson, William La Cava, Patryk Orzechowski, Ryan J Urbanowicz, and Jason H Moore. Pmlb: a large benchmark suite for machine learning evaluation and comparison. *BioData mining*, 10(1):1–13, 2017.

Patryk Orzechowski, William La Cava, and Jason H Moore. Where are we now? a large benchmark study of recent symbolic regression methods. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1183–1190, 2018.

Patryk Orzechowski, Franciszek Magiera, and Jason H Moore. Benchmarking manifold learning methods on a large collection of datasets. In *European Conference on Genetic Programming (Part of EvoStar)*, pages 135–150. Springer, 2020.

Tom Packebusch and Stephan Mertens. Low autocorrelation binary sequences. *Journal of Physics A: Mathematical and Theoretical*, 49(16):165001, 2016. doi:10.1088/1751-8113/49/16/165001. URL https://doi.org/10.1088%2F1751-8113%2F49%2F16%2F165001.

Ingo Paenke, Jürgen Branke, and Yaochu Jin. Efficient search for robust solutions by means of evolutionary algorithms and fitness approximation. *IEEE Trans. Evolutionary Computation*, 10(4):405–420, 2006.

Sinno Jialin Pan and Qiang Yang. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, Oct 2010.

Fortunato Pesarin. *Multivariate permutation tests: with applications in biostatistics*, volume 240. Wiley Chichester, 2001.

Robin L Plackett and J Peter Burman. The design of optimum multifactorial experiments. *Biometrika*, 33(4):305–325, 1946.

Sergey Polyakovskiy, Mohammad Reza Bonyadi, Markus Wagner, Zbigniew Michalewicz, and Frank Neumann. A Comprehensive Benchmark Set and Heuristics for the Traveling Thief Problem. In *Proc. of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 477 —- 484. ACM, 2014. ISBN 9781450326629.

Karl Raimund Popper. *The Logic of Scientific Discovery*. Hutchinson & Co, 2 edition, 1959.

Karl Raimund Popper. *Objective Knowledge: An Evolutionary Approach*. Oxford University Press, 1975.

Mike Preuss. *Experimentation in Evolutionary Computation*, pages 27–54. Springer, 2015.

Kenneth V. Price. Differential Evolution vs. The Functions of the 2^{nd} ICEO. In *Proc. of the IEEE International Conference on Evolutionary Computation*, pages 153 – 157. IEEE, 1997.

F. Pukelsheim. *Optimal Design of Experiments*. Wiley, New York NY, 1993.

Jeremy Rapin and Olivier Teytaud. Nevergrad - A gradient-free optimization platform. https://GitHub.com/FacebookResearch/Nevergrad, 2018.

Howard Harry Rosenbrock. An Automatic Method for Finding the Greatest or Least Value of a Function. *The Computer Journal*, 3(3):175 – 184, 1960.

Jonathan Rowe and Michael Vose. Unbiased black box search algorithms. In *Proc. of Genetic and Evolutionary Computation Conference*, pages 2035 – 2042. ACM, 2011.

Ranjit K Roy. *Design of experiments using the Taguchi approach: 16 steps to product and process improvement*. John Wiley & Sons, 2001.

Ragav Sachdeva, Frank Neumann, and Markus Wagner. The dynamic travelling thief problem: Benchmarks and performance of evolutionary algorithms, 2020.

Thomas J Santner, Brian J Williams, William I Notz, and Brain J Williams. *The design and analysis of computer experiments*, volume 1. Springer, 2003.

Hans-Paul Schwefel. *Evolutionsstrategie und numerische Optimierung*. PhD thesis, Technische Universität Berlin, Fachbereich Verfahrenstechnik, Berlin, Germany, 1975.

David F Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24 (111):647–656, 1970.

Samuel Sanford Shapiro and Martin B Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, 1965.

David J Sheskin. *Handbook of parametric and nonparametric statistical procedures*. crc Press, 2003.

Yuhui Shi and Russell Eberhart. A Modified Particle Swarm Optimizer. In *Proc. of the 1998 IEEE International Conference on Evolutionary Computation, within the IEEE World Congress on Computational Intelligence*, pages 69–73. IEEE, 1998.

Ofer M. Shir, Carola Doerr, and Thomas Bäck. Compiling a Benchmarking Test-Suite for Combinatorial Black-Box Optimization: A Position Paper. In *Proc. of Genetic and Evolutionary Computation Conference*, pages 1753 – 1760. ACM, 2018.

Urban Škvorc, Tome Eftimov, and Peter Korošec. Understanding the Problem Space in Single-Objective Numerical Optimization Using Exploratory Landscape Analysis. *Applieed Soft Computing (ASOC)*, 90:106138, 2020.

Kate Smith-Miles and Simon Bowly. Generating new test instances by evolving in instance space. *Computers & Operations Research*, 63:102–113, November 2015. doi:10.1016/j.cor.2015.04.022. URL https://doi.org/10.1016/j.cor.2015.04.022.

Kate Smith-Miles and Thomas T. Tan. Measuring Algorithm Footprints in Instance Space. In *2012 IEEE Congress on Evolutionary Computation*. IEEE, 2012.

Krzysztof Socha and Marco Dorigo. Ant colony optimization for continuous domains. *European journal of operational research*, 185(3):1155–1173, 2008.

Jörg Stork, A. E. Eiben, and Thomas Bartz-Beielstein. A new taxonomy of global optimization algorithms, Nov 2020. ISSN 1572-9796. URL https://doi.org/10.1007/s11047-020-09820-4.

El-Ghazali Talbi. *Metaheuristics: From Design to Implementation*. John Wiley & Sons Inc., July 2009. ISBN 978-0-470-27858-1.

Ryoji Tanabe and Hisao Ishibuchi. An Easy-to-Use Real-World Multi-Objective Optimization Problem Suite. *Applied Soft Computing (ASOC)*, 89:106078, 2020.

Shigeyoshi Tsutsui, Ashish Ghosh, and Yoshiji Fujimoto. A Robust Solution Searching Scheme in Genetic Search. In *International Conference on Parallel Problem Solving from Nature*, pages 543 – 552. Springer, 1996.

John Wilder Tukey. *Exploratory Data Analysis*, volume 2. Reading, MA, 1977.

Tea Tušar, Dimo Brockhoff, and Nikolaus Hansen. Mixed-integer benchmark problems for single- and bi-objective optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 718 —- 726. ACM, 2019.

Tsuyoshi Ueno, Trevor David Rhone, Zhufeng Hou, Teruyasu Mizoguchi, and Koji Tsuda. COMBO: an efficient Bayesian optimization library for materials science. *Materials discovery*, 4:18–21, 2016.

Niki Veček, Marjan Mernik, and Matej Črepinšek. A chess rating system for evolutionary algorithms: A new method for the comparison and ranking of evolutionary algorithms. *Information Sciences*, 277:656–679, 2014.

Vanessa Volz, Boris Naujoks, Pascal Kerschke, and Tea Tušar. Single- and Multi-Objective Game-Benchmark for Evolutionary Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 647 – 655. ACM, 2019.

Wouter Vrielink and Daan van den Berg. Fireworks algorithm versus plant propagation algorithm. In *IJCCI*, pages 101–112, 2019.

Tobias Wagner. A subjective review of the state of the art in model-based parameter tuning. In Thomas Bartz-Beielstein, Marco Chiarandini, Luis Paquete, and Mike Preuss, editors, *Workshop on Experimental Methods for the Assessment of Computational Systems (WEMACS 2010)*, Algorithm Engineering Report, pages 1–13. TU Dortmund, Faculty of Computer Science, Algorithm Engineering (Ls11), 2010.

Hao Wang, Diederick Vermetten, Furong Ye, Carola Doerr, and Thomas Bäck. Iohanalyzer: Performance analysis for iterative optimization heuristic. *CoRR*, abs/2007.03953, 2020. URL https://arxiv.org/abs/2007.03953. A benchmark data repository is available through the web-based GUI at iohprofiler.liacs.nl/.

Jann Michael Weinand, Kenneth Sörensen, Pablo San Segundo, Max Kleinebrahm, and Russell McKenna. Research trends in combinatorial optimisation. *arXiv e-prints*, art. arXiv:2012.01294, December 2020.

Thomas Weise. jsspinstancesandresults: Results, data, and instances of the job shop scheduling problem, 2019. URL http://github.com/thomasWeise/jsspInstancesAndResults/. A meta-study of 145 algorithm setups from literature on the JSSP.

Thomas Weise and Zijun Wu. Difficult features of combinatorial optimization problems and the tunable w-model benchmark problem for simulating them. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, GECCO '18, pages 1769 — 1776. ACM, 2018. ISBN 9781450357647.

Thomas Weise, Raymond Chiong, Ke Tang, Jörg Lässig, Shigeyoshi Tsutsui, Wenxiang Chen, Zbigniew Michalewicz, and Xin Yao. Benchmarking optimization algorithms: An open source framework for the traveling salesman problem. *IEEE Computational Intelligence Magazine (CIM)*, 9:40–52, August 2014.

L. Darrell Whitley, Soraya B. Rana, John Dzubera, and Keith E. Mathias. Evaluating Evolutionary Algorithms. *Artificial Intelligence (AIJ)*, 85(1-2):245 – 276, 1996.

L. Darrell Whitley, Jean-Paul Watson, Adele Howe, and Laura Barbulescu. Testing, Evaluation and Performance of Optimization and Learning Systems. In Ian C. Parmee, editor, *Adaptive Computing in Design and Manufacture V*, pages 27 – 39, London, 2002. Springer.

Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83, 1945.

David H. Wolpert and William G. Macready. No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, April 1997.

Guohua Wu, Rammohan Mallipeddi, and Ponnuthurai Nagaratnam Suganthan. Problem definitions and evaluation criteria for the CEC 2017 competition on constrained real-parameter optimization. Technical report, National University of Defense Technology, Changsha, Hunan, PR China and Kyungpook National University, Daegu, South Korea and Nanyang Technological University, Singapore, September 2017.

Lin Xu, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. SATzilla: portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research*, 32:565–606, June 2008.