

Towards Large Scale Automated Algorithm Design by Integrating Modular Benchmarking Frameworks

Amine Aziz-Alaoui
ISAE-SUPAERO, Université de
Toulouse, France

Carola Doerr
Sorbonne Université, CNRS, LIP6
Paris, France

Johann Dreo
Thales Research & Technology
Palaiseau, France

ABSTRACT

We present a first proof-of-concept use-case that demonstrates the efficiency of interfacing the algorithm framework ParadisEO with the automated algorithm configuration tool irace and the experimental platform IOHprofiler. By combining these three tools, we obtain a powerful benchmarking environment that allows us to systematically analyze large classes of algorithms on complex benchmark problems. Key advantages of our pipeline are fast evaluation times, the possibility to generate rich data sets to support the analysis of the algorithms, and a standardized interface that can be used to benchmark very broad classes of sampling-based optimization heuristics.

In addition to enabling systematic algorithm configuration studies, our approach paves a way for assessing the contribution of new ideas in interplay with already existing operators—a promising avenue for our research domain, which at present may have a too strong focus on comparing entire algorithm instances.

CCS CONCEPTS

• **Theory of computation** → *Bio-inspired optimization*.

ACM Reference Format:

Amine Aziz-Alaoui, Carola Doerr, and Johann Dreo. 2021. Towards Large Scale Automated Algorithm Design by Integrating Modular Benchmarking Frameworks. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3449726.3463155>

1 INTRODUCTION

When confronted with an optimization problem in practice, one of the major challenges that we face is the selection (and the configuration) of an algorithm that corresponds well to the given problem structure, optimization objective(s), and the available resources (compute, possibility to parallelize computations, accessibility of the problem, etc.). A vast amount of different optimization techniques exist, which renders this *algorithm selection problem* non-trivial.

In practice, algorithm selection is often biased by personal preferences and experiences, as well as by practical aspects such as the availability of ready-to-use implementations. Supporting practitioners in making more systematic choices is one of the key objectives

of our research domain. A key tool for deriving such recommendations is *algorithm benchmarking*, i.e., the analysis of empirical performance data and search trajectories of one or several algorithms on one or several optimization problems [6, 29]. Several important benchmarking tools and software frameworks have been developed by our community to ensure sound and meaningful data extraction. These platforms address different stages of the algorithm selection process. They cover, for example, instance selection and generation [51, 55, 59], feature extraction [35], algorithm configuration [7, 31, 37, 38], experimentation [21, 29, 48, 52], data analysis [12, 23, 25], and performance extrapolation [34]. However, most of these tools are developed in isolation, paying little attention to building compatible interfaces to other benchmarking modules. This significantly hinders their wider adoption.

With this work we demonstrate the benefits of a fully modular benchmarking pipeline design, which keeps the different steps of the benchmarking study in mind. We see our work as a proof of concept for better compatibility between benchmarking software. On the practical side, our pipeline paves a way for assessing the benefits of new algorithmic ideas in the context of and in interplay with other operators and ideas that our community has to offer.

1.1 Our Contribution

Concretely, we propose in this work a benchmarking pipeline that integrates the modular algorithm framework ParadisEO [11, 33] with the algorithm configuration tool irace [38], the experimental platform IOHexperimenter [21], and the data analysis and visualization module IOHAnalyzer [53]. We test our pipeline on tuning a family of genetic algorithms, inspired by [58], on the so-called W-model problem instances suggested in [54].

Quality of the results: We show that irace is capable of finding algorithm instances which outperform all baseline algorithms selected by hand, and this for each of the 19 problem instances that we consider. The relative advantage of the best out of 15 irace suggestions over the best baseline algorithm, measured in terms of volume under the discretized Empirical Attainment Function (see Sec. 3.1), varies between 1% and 30%, with a median gain of 13%.

Scalability: Targeting per-instance algorithm design on synthetic benchmark, our algorithmic framework is capable of generating large set of solvers, up to several millions of unique configurations. We show that it is possible to tackle such spaces thanks to fast computations. For instance, we give irace a budget of 100 000 target runs for each of 19 problems, and it completes the full task in approximately 3 hours on a laptop. In our experience, our C++ pipeline is at least 10 times faster than heavily optimized counterparts in Python, not mentioning that most of the available modular frameworks are not always heavily optimized.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '21 Companion, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8351-6/21/07...\$15.00

<https://doi.org/10.1145/3449726.3463155>

Take-away for instance selection: As a side result, we observe that similar algorithm instances can be suggested by irace for some problems, suggesting that the diversity in performance profiles sought in [54] may be weaker than intended. Our work suggests that an approach like ours may result in a more reliable instance selection, since it will be less biased by a small set of baseline algorithms, but rather be built on a large and diverse set of possible algorithm instances.

Extendability: Our pipeline is ready to perform large benchmark studies, covering large classes of continuous and discrete optimization algorithms. For example, local searches, particle swarm optimization, estimation of distribution algorithms and using numerical or bitstring encodings. Similarly, the pipeline gives direct access to all problems collected in IOHprofiler, which comprises in particular the BBOB functions from the COCO framework [29], the Nevergrad problem suites [48], the W-model instances [55], the PBO suite [22], etc. Additionally, any benchmark or solver which would be plugged into IOHprofiler would be easily used to further extend this study.

1.2 Comparison to Previous Works

Our work is a top-down approach for automatic algorithm design [44], which uses a parametrized algorithmic framework to instantiate many *algorithm instances*. Following [39], we observe that this differs from bottom-up “grammar-based” approaches [16, 17, 47] or Grammatical Evolution [42, 49], which allow for easily designed algorithm space, but complicates algorithm instantiation and optimization. In our case, the width of the design space is already large and we target fast algorithm instantiation. We thus favor the top-down approach. In this first study, we only consider categorical parameters, for the sake of implementation simplicity.

A similar approach to ours was suggested in [9, 10, 40] for multi-objective optimization. Those studies also use irace, but the authors implemented their own modular algorithm frameworks that are restricted to multi-objective optimization. Our work significantly scales up this kind of study, by leveraging larger algorithm design spaces, larger sets of benchmarks, with more problems and allowing a more detailed analysis of the results.

Few other studies consider a bi-objective measurement of performance for automated algorithm design. Most notably, [41] introduced the use of the hypervolume for given quality and time budgets. In our case, we use the volume under the curve of the empirical cumulative histogram of quality and time attainments [15]. This should behave as the sum of hypervolumes defined for a set of thresholds covering the whole domain, instead of a single one.

1.3 Structure of the Paper

Sec. 2 briefly introduces the individual modules of our algorithm design pipeline and how they interplay with each other. The use-case on which we apply this pipeline, as well as the experimental setup are summarized in Sec. 3. The results of our empirical analysis are described in Sec. 4. We conclude our paper in Sec. 5 with a discussion on promising avenues for future work.

1.4 Availability of Code and Data

The code and data used for this study have been archived at [4]. Up-to-date versions of the code are available at <https://github.com/jdreio/paradiseo>

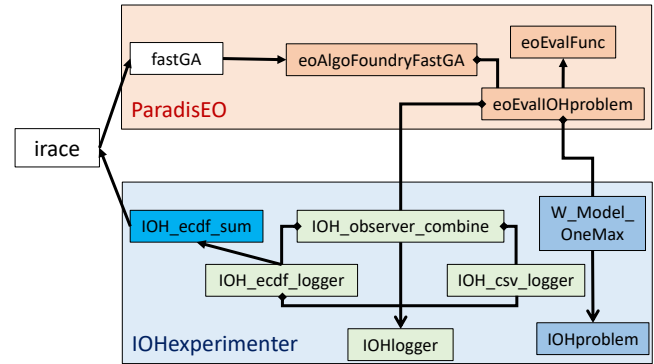


Figure 1: Summary diagram of the FastGA evaluation pipeline involving the Paradiseo (upper part, red colors) and IOHexperimenter (lower part, blue colors) frameworks along with the irace entry point. The execution starts from the irace run command on the left, goes through the Paradiseo modules, which call the IOHexperimenter problem (in blue) and loggers (in green). After the run of the algorithm, a statistic is computed on the ECDF data (in cyan), which is then returned to irace as performance metric (i.e., this is the “fitness value” that the evaluation associates to the configuration under evaluation). Involved classes are represented using the UML convention. For the sake of clarity, the IOHprofiler prefix is written as IOH and the type of the eoAlgoFoundryFastGA slots are indicated as {double} instead of eoOperatorFoundry<double>.

[//github.com/jdreio/paradiseo](https://github.com/jdreio/paradiseo) and <https://github.com/IOHprofiler/IOHexperimenter>.

2 THE MODULAR BENCHMARK PIPELINE

Figure 1 summarizes our automated algorithm design pipeline for the concrete use-case that will be studied in Sec. 3. The pipeline links an algorithm configurator with an algorithm generator and a benchmark platform. The algorithm configurator asks the algorithm generator to instantiate an algorithm, which then solves a problem of the benchmark platform while being observed by a logger. After this run, the logger’s data are summarized as a scalar performance measure, which is sent back to the algorithm configurator. We briefly present in this section the different components of our pipeline, and explain the reasons behind our choices.

Algorithm Framework: Paradiseo. Many evolutionary algorithms share similar design patterns, and are often composed of similar operators. This has given rise to several platforms which aim at supporting their users in designing evolutionary heuristics by compiling a set of readily-available operators within a standardized software environment. Given the substantial work that has been put into these frameworks, we decided to build our pipeline around one of the most powerful toolboxes. To this end, we have ranked 39 frameworks among the ones easily available on the web, based on an adhoc metric combining rapidity, activity, features and license, e.g., [1, 2, 14, 26, 27, 46, 50], to name only a few. Since speed is a major concern for our pipeline, we favor frameworks written in

C++. To select an up-to-date framework and to ensure availability of support in case of technical issues, we also checked the contribution activity in recent years. These two criteria reduced our choices to Paradiseo [11, 33], OpenBeagle [27], and ECF [1]. Among these three, Paradiseo covers the largest portfolio of algorithm families, which are composed in the framework by assembling atomic functions (called *operators*). Paradiseo is also the most actively maintained framework among the three, so that we decided to use it for our work.

The upper part of Figure 1 shows the core classes of Paradiseo involved in our setting.

Algorithm Configuration: irace. Several algorithm configuration tools have been developed in the last decade. Among the most common ones used in our community are irace [38], SMAC [31], SPOT [7], GGA [3], and hyperband [37]. We have chosen irace¹ for this study, for practical considerations (previous experience, availability of documentation, support from development team).

Experimental Environment: IOHexperimenter. The IOHprofiler project [21] is a modular platform for algorithm benchmarking of iterative optimization heuristics (IOH). Within this project, IOHexperimenter provides synthetic benchmarks which are very fast to execute and a standardized way of observing algorithms behavior through so-called *loggers*. We have chosen this platform, because it is fast and its modular design made it particularly easy for us to integrate the algorithm design framework (being written in C++, as Paradiseo). IOHprofiler is also actively maintained, and provides access to broad ranges of different optimization processes.

Compared to Nevergrad [48], we particularly like the detailed logging options, which provide information about the anytime behavior of the algorithms—information that is currently not available in Nevergrad. Compared to the COCO [29] environment, IOHprofiler makes it considerably easier to test algorithms' performance on our own benchmark problems or suites. Finally, the project also supports interactive performance analysis and visualization module, IOHanalyzer, which we used for the interpretation of our data.

The lower part of Figure 1 shows the classes related to the loggers and the problems that are used in our experimental study in Sec. 3.

Data Records: fast ECDF Logger. In our use-case, we decide to tune algorithms for good anytime performance, and to use volume under the approximated empirical cumulative density function (ECDF) curve as objective. To this end, we implement within IOHexperimenter an efficient way of computing these values. This "ECDF logger" will be described in Sec. 3.1.

Data Analysis and Visualization with IOHanalyzer. Data analysis and visualization is performed via IOHanalyzer [53], another module of the IOHprofiler project [21].

3 USE-CASE AND EXPERIMENTAL SETUP

Our use-case is the optimization of the anytime performance of a genetic algorithm on selected instances of the W-model problem. Our performance measure (Sec. 3.1), the algorithmic framework (Sec. 3.2), and the problems (Sec. 3.3) are introduced in the first

three subsections. We then summarize the experimental setup of the whole pipeline in Sec. 3.4. Our objective is to find the best algorithm for each instance, which would be the first step of a per-instance, landscape-aware algorithm selection, for instance. We thus do not consider training versus test sets.

3.1 Anytime Performance Measure: AUC

In order to allow for large scale experiments, we implement a fast logger within IOHexperimenter, which essentially stores a histogram of the two-dimensional distribution of the number of runs having reached a quality/time target. The time dimension is given as the number of calls to the objective function, linearly discretized between zero and the allowed budget. The quality dimension is given as the absolute value of the best solution found during the run, linearly discretized between zero and the known V_{\max} bound (see Table 1). This is essentially a discrete version of the Empirical Attainment Function [15], which is related itself to the multivariate Empirical Cumulative Distribution Function [28].

Figure 2 shows two examples of such histograms, arbitrarily chosen. The matrix defines the considered quality/time targets (v, t) . The color of each cell corresponds to the probability that the algorithm has identified, within the first t function evaluations, a solution of quality at least v . The darker a cell, the larger the fraction of runs that could successfully meet the quality/time target. Note here that we assume minimization as objective.

Using the histogram of the performance ECDF instead of its continuous counterpart allows to keep the data in-memory, in compact data structures, without having to rely on slow disk accesses.

The performance of the considered algorithm is computed as a statistic on this histogram. In our study, we use the volume under the curve (3D counterpart of the area under the curve, AUC) of the discretized ECDF, approximated as the sum of the EAF histogram. This allows for a compromise between quality and time, which is easily available because we consider synthetic benchmarks with known bounds.

3.2 The $(\mu \uparrow \lambda)$ "Fast" GA Family

We chose for our use-case a family of $(\mu \uparrow \lambda)$ GAs, which is to a large extent inspired by the study [58]. Algorithm 1 summarizes the framework, called "FastGA" in the implementation.

Essentially, given a parent population of μ points, each of the λ offspring is created by first deciding which variation operator is applied (line 9): with probability p_c the offspring is generated by first recombining two search points from the parent population (lines 11–13) and then randomly deciding (with probability p_m) whether or not to apply a mutation operator to the so-created offspring (lines 15–18). When crossover was not selected in line 10, the offspring is created by mutation (lines 20–21). When all λ offspring have been created, the iteration is completed by a replacement step (line 25).

Implementation of this Family in Paradiseo: We implement this family of GAs through Paradiseo's "foundries", which allow to register a set of operators (e.g., several kind of mutations) within a "slot" (e.g., the step at which mutation is called within the algorithm). Before each call, it is possible to instantiate a specific operator among the registered ones, for each slot, thus assembling one of the algorithm instance among all the possible combinations

¹Version 3.4.1 of <https://cran.r-project.org/web/packages/irace/>, ran with R 3.6.3.

Algorithm 1: A Configurable Family of $(\mu \pm \lambda)$ Genetic Algorithms.

```

1 Input: Budget  $B$ , configuration  $(\mu, \lambda, p_c, p_m)$ , choice of the
  operators and conditional parameters. Note that  $P$  and  $P'$ 
  are multi-sets, i.e., the same point may appear multiple
  times;
2 Initialization:
3    $P \leftarrow \text{InitialSampling}(\mu)$ ;
4   evaluate the  $\mu$  points in  $P$ ;
5    $\text{Evals} \leftarrow \mu$ ;
6 Optimization: while  $\text{Evals} < B$  do
7    $P' \leftarrow \emptyset$ ;
8   for  $i = 1, \dots, \lambda$  do
9     Sample  $r_c \in [0, 1]$  u.a.r.;
10    if  $r_c \leq p_c$  then
11       $(y^{(i,1)}, y^{(i,2)}) \leftarrow \text{SelectC}(P)$ ;
12       $(y'^{(i,1)}, y'^{(i,2)}) \leftarrow \text{Crossover}(y^{(i,1)}, y^{(i,2)})$ ;
13      Sample  $z^{(i,1)} \in \{y'^{(i,1)}, y'^{(i,2)}\}$  u.a.r. ;
14      Sample  $r_m \in [0, 1]$  u.a.r. ;
15      if  $r_m \leq p_m$  then
16         $z^{(i,2)} \leftarrow \text{Mutation}(z^{(i,1)})$ ;
17      else
18         $z^{(i,2)} \leftarrow z^{(i,1)}$ ;
19    else
20       $z^{(i,1)} \leftarrow \text{SelectM}(P)$ ;
21       $z^{(i,2)} \leftarrow \text{Mutation}(z^{(i,1)})$ ;
22    Evaluate  $z^{(i,2)}$ ;
23     $\text{Evals} \leftarrow \text{Evals} + 1$ ;
24     $P' \leftarrow P' \cup \{z^{(i,2)}\}$ ;
25   $P \leftarrow \text{Replace}(P, P', \mu)$ ;

```

of operators. Note that operators can be simple numbers, like a probability. Operators are referenced within slots by their indices.

Most of the operators we use were already available in *ParadisEO*, to the exception of mutations operators with indices 1–5 (see below), which we implemented for this study. We also implemented the algorithm 1 as the *eoFastGA* class², in which to plug the operators.

We consider the following operators and parametrizations, which result in a (large) total number of 1 630 475 different configurations of Algorithm 1. Numbers in brackets indicate the indices of the corresponding operators within its slot.

InitialSampling(μ): Initialization of the Algorithm (1 option). We only consider independent uniform sampling, *i.e.*, the μ points are i.i.d. uniform samples. The corresponding *ParadisEO* operator is *eoInitFixedLength*.

²All our code is contributed to the *ParadisEO* project.




Figure 2: Example of two ECDF histograms. Abscissa shows the time dimension and the ordinate shows the target quality. The colormap shows the probability for the considered algorithm to reach a quality/time target falling in each bucket of the histogram. The figures show ECDF histograms for the 1ptGA algorithm (see Sec. 3.2) after 50 independent runs on problem 2 (left) and of problem 5 (right), respectively, using 20 buckets on each axis.

Crossover rate p_c (6 options). We consider $p_c \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8\}$. Being only able to use the *integer* and *categorical* interface for *irace*.

SelectC(P): Selection of two points for the crossover operation (7 options). Note that in the implementation, the selection operator (line 11) is called twice to select the two candidate points.

- [0] *eoRandomSelect()*: Uniformly select a point from P (without removing the first selected individual from the set P used by the second selection). (1 option).
- [1] *eoStochTournamentSelect(k)*: Select a point from P with tournament selection, *i.e.*, we select uniformly at random k different points in P and the best one of these is selected. k denotes the tournament size as percentage of population (*i.e.*, $k \in [0, 1]$). (1 option, $k = 0.5$).
- [2] *eoSequentialSelect()*: Select the best point from P (with respect to the objective function value). This operator is sometimes referred to as *elitist selection* or *truncation selection*. When called twice, it selects the two distinct best points from P . (1 option).
- [3] *eoProportionalSelect()*: Select a point from P with so-called fitness-proportional selection, *i.e.*, point $x \in P$ is chosen with probability $f(x)/\sum_{y \in P} f(y)$. (1 option).
- [4–6] *eoDetTournamentSelect(k)*: Like *eoDetTournamentSelect*, but k is deterministic. (3 different options, each one for $k \in [2, 6, 10]$).

Crossover(x, y): Bivariate Variation Operators (11 options).

- [0–4] *eoUBitXover(b_c)*: Uniform crossover with bias (or “preference” in *ParadisEO*) b_c , setting (independently for each position $i \in [1..n]$) $z_i = x_i$ with probability b_c and setting $z_i = y_i$ otherwise. z denotes the offspring element coming from the crossover of x and y . (5 different options, $b_c \in [0.1, 0.3, 0.5, 0.7, 0.9]$).
- [5–9] *eoNptsBitXover(k)*: k -point crossover, which selects i_1, \dots, i_k uniformly at random and without replacement from $[1..n]$ and sets $z_i = x_i$ for $i \in [1..i_1] \cup [i_2 + 1..i_3] \cup \dots$ and

sets $z_i = y_i$ for $i \in [i_1 + 1..i_2] \cup [i_3 + 1..i_4] \cup \dots$ (5 different options, $k \in [1, 3, 5, 7, 9]$).

[10] `eo1PtBitXover()`: Classic 1-point crossover. (1 option). We mistakenly added this option even if it is the same as the previous crossover with $k = 1$.

Mutation probability p_m (6 options): We consider $p_m \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8\}$.

Mutation(x): Univariate Variation Operator (11 options). All mutation operators are unary unbiased in the sense proposed in [36]. For a compact representation, we follow the characterization suggested in [19] and define the mutation operators via the distributions that they define over the possible mutation strengths $k \in [0..n]$. After sampling k from the operator-specific distribution, the k -bit flip operator, $\text{flip}_k(\cdot)$, is applied; it flips the entries in k uniformly chosen, pairwise different bits (i.e., the k bits are chosen u.a.r. without replacement).

[0] `eoUniformBitMutation()`: The “uniform” mutation operator, which samples k uniformly at random in the set $[0..n]$. (1 option).

[1] `eoStandardBitMutation($p = 1/n$)`: This is the standard bit mutation with mutation rate p . It chooses k from the binomial distribution $\mathcal{B}(n, p)$. (1 option).

[2] `eoConditionalBitMutation($p = 1/n$)`: A conditional standard bit mutation operator with mutation rate p . It chooses k' from $\mathcal{B}(n - 1, p)$ and applies the $\text{flip}_k(\cdot)$ operator with $k = k' + 1$. (1 option).

[3] `eoShiftedBitMutation($p = 1/n$)`: The “shifted” standard bit mutation with mutation rate p , suggested in [13]. It samples k' from the binomial distribution $\mathcal{B}(n, p)$. When $k' = 0$, it uses $k = 1$ and it uses $k = k'$ otherwise. (1 option).

[4] `eoNormalBitMutation(p, σ^2)`: The “normal” mutation operator suggested in [57]. It samples k from the normal distribution $\mathcal{N}(pn, \sigma^2)$. When $k > n$, k is replaced by a value chosen uniformly at random in the set $[0..n]$. (1 option, $p = 1/n$ and $\sigma^2 = 1.5$).

[5] `eoFastBitMutation(β)`: The “fast” mutation operator suggested in [20]. It samples k' from the power-law distribution $\mathcal{P}[L = k] = (C_{n/2}^\beta)^{-1} k^{-\beta}$ with $C_{n/2}^\beta = \sum_{i=1}^{n/2} i^{-\beta}$. When k' is larger than n , it samples a uniform value k in $[0..n]$, and it uses $k = k'$ otherwise. (1 option, $\beta = 1.5$).

[6–10] `eoDetSingleBitFlip(k)`: Deterministically applies $\text{flip}_k(\cdot)$. (5 different options, $k \in [1, 3, 5, 7, 9]$).

SelectM(P): Selection of one point for the mutation operation if crossover was not chosen (7 options). We essentially have the same selection operators as for crossover. The only difference is that we select only one point instead of two.

Replace(P, P', μ): Replacement of population (11 options).

[0] `eoPlusReplacement()`: The best μ points of the multiset $P \cup P'$ are chosen. (1 option).

[1] `eoCommaReplacement()`: The best μ points of the offspring multiset P' are chosen. (1 option).

[2] `eoSSGAWorseReplacement()`: The $\min(\lambda, \mu)$ points of the offspring multiset P' replace the worst points in P . (1 option).

[3–5] `eoSSGAStochTournamentReplacement(k)`: Like `eoSSGADetTournamentReplacement`, k being the tournament size as percentage of population. (3 different options, $k \in [0.51, 0.71, 0.91]$).

[6–10] `eoSSGADetTournamentReplacement(k)`: The μ points are selected through tournament selection. Each tournament involves k uniformly chosen points in $P \cup P'$ and the best ones of these k points is selected. This procedure is repeated μ times, each time removing an already selected point from the multi-set $P \cup P'$. (5 different options, $k \in [2, 4, 6, 8, 10]$).

This concludes our description of the high level operators of our family of $(\mu \nmid \lambda)$ GAs. The set of all combinations generates the algorithm design space on which we let irace search for the configuration(s) that best solve a given problem instance.

Baseline Algorithms. We consider four baseline algorithms, against which we compare the results of the automated design. They were manually chosen without particular justification. (1) $(\lambda + \lambda)$ **EA**: no crossover, plus replacement, standard bit mutation, random selector for mutations. (2) $(\lambda + \lambda)$ **fEA**: no crossover, plus replacement, fast bit mutation, random selector for mutations. (3) $(\lambda + \lambda)$ **xGA**: sequential selections, uniform crossover, standard bit mutation, plus replacement, $p_c = 0.4$, $b_c = 0.4$. (4) $(\lambda + \lambda)$ **1ptGA**: sequential selections, 1-point crossover, standard bit mutation, plus replacement, $p_c = 0.4$, $b_c = 0.4$.

3.3 The W-Model Problems

We evaluate our automated algorithm design pipeline on the W-model functions originally suggested in [55]. In a nutshell, the W-model is a benchmark problem generator, which allows to tune different characteristics of the problems, see below for a description. We selected from this family of benchmark problems the 19 instances suggested in [54], which are summarized in Table 1. Note here that the description differs from that given in [54], since we used the implementation within IOHexperimenter, which was made available in the context of the work [22]. The problem instances listed in Table 1 are identical to those suggested in [54], it is only the representations that differ.

It was suggested in [22] to superpose the W-model transformations to different optimization problems. The instances selected in [54], however, were only selected from transformations applied to the ONEMAX problem $\text{OM} : \{0, 1\} \rightarrow [0..n], x \rightarrow \sum_{i=1}^n x_i$. The ONEMAX problem has a very smooth and non-deceptive fitness landscape. Due to the well-known coupon collector effect [24], it is relatively easy to make progress when the function values are small, and the probability to obtain an improving move decreases considerably with increasing function values. The complexity of the ONEMAX problem can be considerably increased through the following W-model transformations.

(1) **Neutrality** $W(\cdot, \mu_W, \cdot, \cdot)$: The bit string (x_1, \dots, x_n) is reduced to a string (y_1, \dots, y_m) with $m := n/\mu_W$, where μ_W is a parameter of the transformation (we use the subscript W to indicate that these are parameters of the W-model problem generator). For each $i \in [m]$ the value of y_i is the majority of the bit values in the size- μ substring $(x_{(i-1)\mu_W+1}, \dots, x_{i\mu_W})$ of x . That is, $y_i = 1$ if and only if

Table 1: Test problems on which the pipeline is evaluated, taken from [54]. In column “best” we list the baseline algorithm with largest average AUC value, reported in column AUC_b . AUC_i is the average of the 15 mean AUC of elite configurations, as suggested by 15 independent runs of irace with default hyperparameters. AUC-values are w.r.t. to at least 50 validation runs and “rel.” indicates the relative gain $(AUC_i - AUC_b)/AUC_b$.

FID	dim	μ_W	ν_W	γ_W	V_{\max}	best	AUC_b	AUC_i	rel.
1	20	2	6	10	10	xGA	8378	8740	4%
2	20	2	6	18	10	fEA	8402	8754	4%
3	16	1	5	72	16	fEA	8352	8397	1%
4	48	3	9	72	16	EA	8299	8914	7%
5	25	1	23	90	25	fEA	8003	8510	6%
6	32	1	2	397	32	1pt	7055	7311	4%
7	128	4	11	0	32	1pt	6833	8183	20%
8	128	4	14	0	32	EA	6885	8499	23%
9	128	4	8	128	32	xGA	8154	8786	8%
10	50	1	36	245	50	fEA	7216	8122	13%
11	100	2	21	256	50	EA	8314	9139	10%
12	150	3	16	613	50	EA	8034	8730	9%
13	128	2	32	256	64	fEA	8076	9345	16%
14	192	3	21	16	64	fEA	6173	7677	24%
15	192	3	21	256	64	fEA	6797	8292	22%
16	192	3	21	403	64	fEA	7273	8592	18%
17	256	4	52	2	64	xGA	6935	9028	30%
18	75	1	60	16	75	EA	5958	7089	19%
19	150	2	32	4	75	EA	7399	8717	18%

there are at least $\mu_W/2$ ones in this “block”. When $n/\mu_W \notin \mathbb{N}$, the last bits of x are copied to y .

(2) **Epistasis** $W(.,., \nu_W, .)$: Epistasis introduces local perturbations to the bit strings. It first “cuts” the input string (x_1, \dots, x_n) into subsequent blocks of size ν_W . Using a permutation $e_{\nu_W} : \{0, 1\}^{\nu_W} \rightarrow \{0, 1\}^{\nu_W}$, each substring $(x_{(i-1)\nu_W+1}, x_{(i-1)\nu_W+2}, \dots, x_{i\nu_W})$ is mapped to another string $(y_{(i-1)\nu_W+1}, y_{(i-1)\nu_W+2}, \dots, y_{i\nu_W}) = e_{\nu_W}((x_{(i-1)\nu_W+1}, x_{(i-1)\nu_W+2}, \dots, x_{i\nu_W}))$. The permutation e_{ν_W} is chosen in a way that Hamming-1 neighbors are mapped to strings of Hamming distance at least $\nu_W - 1$, see [55] for examples.

(3) **Ruggedness and Deceptiveness** $W(.,., \gamma_W)$: This layer perturbs the fitness values, by applying a permutation $\sigma(\gamma_W)$ to the possible fitness values $[0..n]$. The parameter γ_W can be thought of as a parameter which controls the distance of the permutation to the identity. The permutations $\sigma(\gamma_W)$ are chosen in a way such that the “hardness” of the instances monotonically increases with increasing γ_W , see [55] for details.

We convert these functions into a minimization problem by multiplying all values by -1 .

3.4 Experimental Setup

Our test bed is the automated design of Algorithm 1 with the options specified in Sec. 3.2 and with the objective to maximize the AUC as defined in Sec. 3.1, and this for each of the 19 problems listed in Table 1. These instances of the W-model problem were suggested

in [54] based on an empirical study using clustering of algorithm performance data, with the goal to select a diverse collection of benchmark problems. Note here that we tune the algorithms for each problem individually. That is, we apply our algorithm design pipeline 19 independent times.

For the sake of simplicity, we fix the population sizes to $\lambda = \mu = 5$, for the search performed by irace and for our baseline algorithms.

For each use-case, we set the budget of the algorithms to $5n$ function evaluations (FEs). To compute the AUC, we evaluate the performance at 100 linearly distributed budgets $b^1, \dots, b^{100} \in [1, 5n]$ and at 100 linearly distributed target values $v^1, \dots, v^{100} \in [0, V_{\max}]$. Linearization computes the bucket index $i = \lfloor (x - x_{\min}) / (x_{\max} - x_{\min}) \cdot 100 \rfloor$ for both budgets and targets.

To find the best algorithm design, we allow irace a budget of 100 000 target runs. We ensure that irace performed at least 50 independent runs for the first ranked elite configuration, adding additional runs if needed, and keeping all runs if irace conducted more than 50 runs. We do not consider the other elite configurations ranked by irace, even in draw cases. We run irace 15 independent times, to check the robustness of its selection. We compare performance to the four baseline algorithms, which we run 50 independent times each on each of the 19 test problems.

In total, these experiments took around 15×3 hours on a computer with four Intel CPU cores i5-7300HQ at 2.50GHz and Crucial P1 solid-state disks.

4 EXPERIMENTAL RESULTS

Comparison of AUC Values by Function. Table 1 compares the AUC values of the best out of the four baseline algorithms against that of the elite configuration suggested by irace. We observe that, for each of the 19 functions, the elite configurations suggested by irace perform better than the best baseline algorithm. We report in Table 1 the average values, but the differences between the individual irace runs are very small, less than 2.1% difference in AUC value between the best and the worst elite configuration for all 19 problems, and less than 1% performance difference for 9 out of the 19 functions. The relative advantage of the irace recommendations over the best baseline algorithms varies between 1% and 30%. When looking at each of the 15 elite configurations suggested per function, the best relative advantage is 31% for F17, whereas two of the irace elites performed worse than the best of the four baseline algorithms on function F3. For all other functions, all 15 irace elites have a better AUC value than the best of the four baseline algorithms. However, although we see a clear advantage of the irace configurations, we should keep in mind that the irace configurations are specifically tuned for each function, whereas the configurations of the four baseline algorithms are identical for all 19 W-model functions.

Unfortunately, our pipeline does not yet allow to tune a single best solver, i.e., a single configuration that maximizes the AUC under the aggregated ECDF curve. Adding this functionality is a straightforward extension of our framework, which we plan to address in future work. The key challenge here is that Paradiseo does not have the feature to easily reset on the fly the states of solvers between two runs on different problems.

Table 2: Configuration of the best out of the elite recommendation suggested by 15 independent runs of irace, for each of the 19 benchmark problems specified in Table 1, and compared against the configuration of the four baseline algorithms. The “op.” column gives the number of options per operator. All other integer values correspond to the indices with which the different options are listed in Sec. 3.2 and “-” indicates a non-applicable element (e.g., no crossover operator is used when $p_c = 0$).

Operator	op.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	EA	fEA	xGA	1ptGA
p_c	5	1	4	1	2	4	0	3	0	2	4	3	2	3	1	2	2	3	4	4	0	0	2	2
SelectC	7	2	5	3	1	2	-	0	-	2	2	2	2	6	5	5	2	2	2	2	-	-	2	2
Crossover	11	1	2	8	1	2	-	3	-	2	2	10	5	2	9	2	10	2	2	2	-	-	2	5
p_m	5	2	3	2	2	4	-	4	-	4	4	4	4	4	4	4	4	4	4	4	-	-	2	2
SelectM	7	2	4	6	6	3	2	3	2	5	5	2	3	1	2	6	6	5	1	6	0	0	2	2
Mutation	11	8	9	3	9	7	6	10	10	10	9	10	9	10	8	8	10	10	8	9	1	5	1	1
Replace	11	8	9	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 3: Distribution of operators variants recommended by 15 runs of irace, for problems 5 (left), problem 17 (center) and all problems (right). The most selected indices are highlighted in bold and the darker the background color, the more often the operator instance is selected. Empty cells indicates that irace never selected the operator instance, cells with a “-” entry marks indices which are not defined for this operator.

	Problem 5											Problem 17											All problems										
Op. index	0	1	2	3	4	5	6	7	8	9	10	0	1	2	3	4	5	6	7	8	9	10	0	1	2	3	4	5	6	7	8	9	10
P_c	5	6	1	3		-	-	-	-	-	-				1	14	-	-	-	-	-	-	18	74	56	57	65	-	-	-	-	-	-
SelectC	1		5	2		4	3	-	-	-			15						-	-	-	-	18	16	141	19	18	30	28	-	-	-	-
Crossover	1	5	2	2	1	1				1	2		15									11	29	78	24	6	8	11	19	18	23	43	
P_m			2	3	10	-	-	-	-	-	-					15	-	-	-	-	-	-	4	9	25	49	183	-	-	-	-	-	-
SelectM			14					1	-	-	-	6	2	1	2	2		2	-	-	-	-	26	17	61	36	27	41	62	-	-	-	-
Mutation	2	5		1	1	2	4												1	9	4	1	2	7		3	1	2	6	10	62	81	96
Replace	15											15											245	2	3	6	2	1	2	5	1	3	

Comparison of the Configurations. Table 2 summarizes the best of the 15 elite configurations that were suggested by irace and compares them against the four baseline algorithms. Table 3 shows the distribution of operators chosen by the 15 irace runs. For the latter, we have chosen problem 17 as an example because we observed here the largest relative gain (see Table 1). We have added problem 5 for comparison, because the distribution of operators suggested for it is very distinct from that of problem 17.

It is worth noting that each operator is selected at least once in the 19×15 elite configurations suggested by irace (Table 3, right), which seems to confirm that i) different operators work well on different problems and ii) that irace searches the full design space, giving some indication that it is not too large or too complex for automated tuning approaches.

We can see that, among all the best configurations proposed by irace across 15 runs, none are similar to one of the baseline algorithms. The probability of mutation p_m is most frequently set to higher values and the most often chosen mutation is deterministic bit flip with larger number of bits (index 10 in the *Mutation* slot, which is the flip₉ mutation operator). This indicates that larger mutation strengths could have been worth investigating, a result that has surprised us, since in most benchmark studies we see small mutation rates as defaults (albeit we do not consider various population sizes in this study). The results confirm the superiority of the *plus* replacement (id. 0 in the *Replace* slot) and support the use of an elitist selection for the crossover candidates (id. 2 in *SelectC*). We can also see that the uniform crossover with $b_c = 0.5$

(id. 2 in *Crossover*) is more often chosen, like a small probability of performing a crossover (id. 1 in p_c).

For some problems, irace almost always suggest a similar algorithm. On problem 17, for example, it often selects a GA with a large probability of applying uniform crossover in combination with deterministic bit flip mutations. For some other problems, a larger variance on the selected operators can be observed. For instance on problem 5, irace selects a high mutation probability along with an elitist mutation selection, but does not show a clear preference for the other slots.

These results support the idea that there is not always a single best solver (i.e., “No Free Lunch”), even when considering limited design and benchmarking spaces. We also see that some problems seem to require certain design choices, whereas others can be solved well by a broad range of configurations. A more detailed analysis of how these preferences correlate with the characteristics of the problems should offer plenty of interesting insights, but is left for future work.

Fixed-budget solution qualities. Figure 3 shows two examples of convergence plots, where we plot the values of the best solutions found so far against the number of objective function evaluations performed, for each baseline algorithm and for the best elite configuration selected by irace. Problems 5 and 17 are chosen to allow for comparison with Table 2.

We observe that the elite configuration on problem 17 is largely more efficient than any of the baseline algorithms. However, on

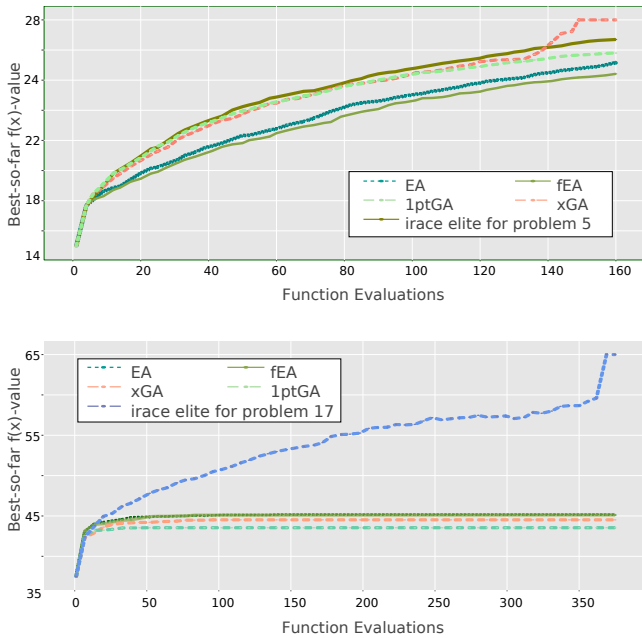


Figure 3: Convergence plots for the baseline algorithms and the elite configurations suggested by irace, on problem 5 (top) and problem 17 (bottom).

problem 5, the elite configuration is only the most efficient until 140 evaluations. It is selected nonetheless, because we consider the AUC of the 2D ECDF, which takes into account the average performance (across all budgets and targets) rather than the terminal budget of the best target. We believe that, whatever the performance metric we choose, there will always exist such artifacts, where some algorithm would be the best, had we chosen another metric. It is clear, however, that even in this plot the elite configuration performs better *most of the time*.

5 CONCLUSIONS AND FUTURE WORK

By interfacing the three state-of-the-art benchmarking modules from the evolutionary computation literature, irace [38], Paradiseo [33], and IOHprofiler [21], we have introduced in this work a powerful pipeline for the automated design of sampling-based optimization algorithms. We have demonstrated its efficiency on the use-case of tuning a family of genetic algorithms on instances of the W-model [55] suggested in [54].

Our results support the idea that automated algorithm design can lead to increased performances and that there are efficient designs which wait to be studied more thoroughly. We believe that efficient pipelines like the one we introduce has the potential to help raising the level of abstraction at which researchers are working. Using automated algorithm design, it becomes possible to check if a newly designed operator can actually be useful in some algorithms/problems coupling [56]. We also believe that such studies can help deriving generic rules about algorithm design and could probably help theoretical researchers by suggesting *where* to look for interesting structures.

The modular design of the pipeline and its components makes our approach very broadly applicable. It is not restricted to particular types of problems nor to specific algorithms. In particular, extensions to continuous or mixed-integer problems are rather straightforward. Indeed, the Paradiseo framework is designed to separate operators which are independent of the encoding (selection, replacement, etc.) from operator which depends on it (mutation, crossover, etc.), allowing for easy reuse of components and extensions to other algorithmic paradigms (estimation of distribution, local search, multi-objective, etc.). Additionally, the IOHExperimenter provides loggers for vectorial encodings and benchmarks for both numerical and bitstring encodings.

Our work is partially motivated by an industrial application that requires an automated configuration of hardware products. However, we believe that our pipeline is not only interesting for such practical purposes. For researchers, our pipeline offers an elegant way of assessing new algorithm operators and their interplay with already existing ones.

In terms of further development, we plan to add the necessary features which would i) allow for running the same algorithm on multiple problems, while using a single logger that aggregates the results and would ii) support irace's interface for numerical parameters (additionally to categorical and integer ones).

We then plan to test the approach on different algorithms families, with a possible extension to generic “bottom-up” hybridization grammars [43] and studies on the most efficient algorithms design (e.g., on the correlations between elite algorithms' operators).

We also plan to extend the framework by integrating feature extraction methods that use algorithm trajectory data [5, 18] and/or samples specifically made for exploratory landscape analysis [35, 45] to couple the algorithm design to such information, similar to the per-instance configuration approaches made in [8, 30].

Our long-term vision is a pipeline for the automated design of algorithms which adjust their behavior during the optimization process, by taking into account information accumulated so far, similar to the dynamic algorithm configurations studied under the notion of *parameter control* [32]. In contrast to the static designs considered in this work, the automated design of dynamic algorithms requires to select suitable update rules (e.g. based on time, on progress, on self-adaption, etc.).

Finally, we also consider interesting the idea to provide a user-friendly front-end which allows users to assemble a benchmark study by selecting (e.g. through a graphical user interface) one or more algorithms and problems, the budget, etc. and then passing on this study to an automated interface which tunes (if desired) and runs the algorithm(s) and then automatically directs its users to the data summary and visualization platform IOHanalyzer, where the results of the empirical study can be analyzed. We believe that such a pipeline would greatly improve the deployment of evolutionary methods in practice.

Acknowledgments. We thank the GECCO ECADA workshop reviewers for very constructive feedback and pointers to [41]. This work has been supported by the Paris Ile-de-France region. This work was partially done during the M2 master internship of Amine Aziz-Alaoui at École Polytechnique, Institut Polytechnique de Paris, CNRS, LIX, Palaiseau, France.

REFERENCES

- [1] [n.d.]. Evolutionary Computation Framework (ECF), howpublished = <http://ecf.zemris.fer.hr/>, note = Last visited: 2021-02-04.
- [2] [n.d.]. Jnetics, howpublished = <https://jnetics.io/>, note = Last visited: 2021-02-04.
- [3] Carlos Ansótegui, Yuri Malitsky, Horst Samulowitz, Meinolf Sellmann, and Kevin Tierney. 2015. Model-based Genetic Algorithms for Algorithm Configuration. In *Proc. of International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 733–739.
- [4] Amine Aziz-Alaoui, Carola Doerr, and Johann Dreö. 2021. *Towards Large Scale Automated Algorithm Design by Integrating Modular Benchmarking Frameworks – data and analysis scripts*. <https://doi.org/10.5281/zenodo.4729568>
- [5] Lukás Bajer, Zbynek Pitra, Jakub Repický, and Martin Holena. 2019. Gaussian Process Surrogate Models for the CMA Evolution Strategy. *Evolutionary Computation* 27, 4 (2019), 665–697. https://doi.org/10.1162/evco_a_00244
- [6] Thomas Bartz-Beielstein, Carola Doerr, Jakob Bossek, Sowmya Chandrasekaran, Tome Eftimov, Andreas Fischbach, Pascal Kerschke, Manuel López-Ibáñez, Katherine M. Malan, Jason H. Moore, Boris Naujoks, Patryk Orzechowski, Vanessa Volz, Markus Wagner, and Thomas Weise. 2020. Benchmarking in Optimization: Best Practice and Open Issues. *CoRR abs/2007.03488* (2020). [arXiv:2007.03488](https://arxiv.org/abs/2007.03488) <https://arxiv.org/abs/2007.03488>
- [7] Thomas Bartz-Beielstein, Oliver Flasch, Patrick Koch, and Wolfgang Konen. 2010. SPOT: A toolbox for interactive and automatic tuning in the R environment. In *Proc. of the 20. Workshop Computational Intelligence*. Universitätsverlag Karlsruhe, 264–273.
- [8] Nacim Belkhir, Johann Dreö, Pierre Savéant, and Marc Schoenauer. 2017. Per instance algorithm configuration of CMA-ES with limited budget. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'17)*. ACM, 681–688. <https://doi.org/10.1145/3071178.3071343>
- [9] Leonardo C. T. Bezerra, Manuel López-Ibáñez, and Thomas Stützle. 2016. Automatic Component-Wise Design of Multi-Objective Evolutionary Algorithms. *IEEE Transactions on Evolutionary Computation* 20, 3 (2016), 403–417. <https://doi.org/10.1109/TEVC.2015.2474158>
- [10] Leonardo C. T. Bezerra, Manuel López-Ibáñez, and Thomas Stützle. 2020. Automatically Designing State-of-the-Art Multi- and Many-Objective Evolutionary Algorithms. *Evolutionary Computation* 28, 2 (2020), 195–226. https://doi.org/10.1162/evco_a_00263
- [11] Sébastien Cahon, Nordine Melab, and El-Ghazali Talbi. 2004. ParadisEO: A Framework for the Reusable Design of Parallel and Distributed Metaheuristics. *J. Heuristics* 10, 3 (2004), 357–380. <https://doi.org/10.1023/B:HEUR.0000026900.92269.ec> Latest release available on <https://nojhan.github.io/paradisEO/>.
- [12] Borja Calvo, Ofer M. Shir, Josu Ceberio, Carola Doerr, Hao Wang, Thomas Bäck, and Jose A. Lozano. 2019. Bayesian performance analysis for black-box optimization benchmarking. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'19, Companion)*. ACM, 1789–1797. <https://doi.org/10.1145/3319619.3326888>
- [13] Eduardo Carvalho Pinto and Carola Doerr. 2018. Towards a More Practice-Aware Runtime Analysis of Evolutionary Algorithms. *CoRR abs/1812.00493* (2018). [arXiv:1812.00493](https://arxiv.org/abs/1812.00493) <http://arxiv.org/abs/1812.00493>
- [14] T. Cloete, Andries Petrus Engelbrecht, and Gary Pampara. 2008. Cilib: A collaborative framework for Computational Intelligence algorithms - Part II. In *Proc. of the International Joint Conference on Neural Networks (IJCNN'08)*. IEEE, 1764–1773. <https://doi.org/10.1109/IJCNN.2008.4634037>
- [15] Viviane Grunert da Fonseca, Carlos M. Fonseca, and Andreia O. Hall. 2001. Inferential Performance Assessment of Stochastic Optimisers and the Attainment Function. In *Evolutionary Multi-Criterion Optimization, First International Conference, EMO 2001, Zurich, Switzerland, March 7-9, 2001, Proceedings (Lecture Notes in Computer Science)*, Eckart Zitzler, Kalyanmoy Deb, Lothar Thiele, Carlos A. Coello Coello, and David Corne (Eds.), Vol. 1993. Springer, 213–225. https://doi.org/10.1007/3-540-44719-9_15
- [16] Marcelo de Souza and Marcus Ritt. 2018. Automatic Grammar-Based Design of Heuristic Algorithms for Unconstrained Binary Quadratic Programming. In *Evolutionary Computation in Combinatorial Optimization - 18th European Conference, EvoCOP 2018, Parma, Italy, April 4-6, 2018, Proceedings (Lecture Notes in Computer Science)*, Arnaud Liefooghe and Manuel López-Ibáñez (Eds.), Vol. 10782. Springer, 67–84. https://doi.org/10.1007/978-3-319-77449-7_5
- [17] Marcelo de Souza and Marcus Ritt. 2018. An Automatically Designed Recombination Heuristic for the Test-Assignment Problem. In *2018 IEEE Congress on Evolutionary Computation, CEC 2018, Rio de Janeiro, Brazil, July 8-13, 2018*. IEEE, 1–8. <https://doi.org/10.1109/CEC.2018.8477801>
- [18] Bilel Derbel, Arnaud Liefooghe, Sébastien Vérel, Hernán E. Aguirre, and Kiyoshi Tanaka. 2019. New features for continuous exploratory landscape analysis based on the SOO tree. In *Proc. of Foundations of Genetic Algorithms (FOGA'19)*. ACM, 72–86. <https://doi.org/10.1145/3299904.3340308>
- [19] Benjamin Doerr, Carola Doerr, and Jing Yang. 2020. Optimal parameter choices via precise black-box analysis. *Theoretical Computer Science* 801 (2020), 1–34. <https://doi.org/10.1016/j.tcs.2019.06.014>
- [20] Benjamin Doerr, Huu Phuoc Le, Régis Makhlama, and Ta Duy Nguyen. 2017. Fast genetic algorithms. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'17)*. ACM, 777–784. <https://doi.org/10.1145/3071178.3071301>
- [21] Carola Doerr, Hao Wang, Furong Ye, Sander van Rijn, and Thomas Bäck. 2018. IOHprofiler: A Benchmarking and Profiling Tool for Iterative Optimization Heuristics. *CoRR abs/1810.05281* (2018). <http://arxiv.org/abs/1810.05281> Available at <http://arxiv.org/abs/1810.05281>. A more up-to-date documentation of IOHprofiler is available at <https://iohprofiler.github.io/>.
- [22] Carola Doerr, Furong Ye, Naama Horesh, Hao Wang, Ofer M. Shir, and Thomas Bäck. 2020. Benchmarking discrete optimization heuristics with IOHprofiler. *Applied Soft Computing* 88 (2020), 106027. <https://doi.org/10.1016/j.asoc.2019.106027>
- [23] Tome Eftimov, Gasper Petelin, and Peter Korosec. 2020. DSCTool: A web-service-based framework for statistical comparison of stochastic optimization algorithms. *Appl. Soft Comput.* 87 (2020), 105977. <https://doi.org/10.1016/j.asoc.2019.105977>
- [24] Philippe Flajolet, Danièle Gardy, and Loys Thimonier. 1992. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Applied Mathematics* 39, 3 (1992), 207–229. [https://doi.org/10.1016/0166-218X\(92\)90177-C](https://doi.org/10.1016/0166-218X(92)90177-C)
- [25] Carlos M. Fonseca, Andreia P. Guerreiro, Manuel López-Ibáñez, and Luís Paquete. 2011. On the Computation of the Empirical Attainment Function. In *Proc. of Evolutionary Multi-Criterion Optimization (EMO'11) (LNCS)*, Vol. 6576. Springer, 106–120.
- [26] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. 2012. DEAP: Evolutionary Algorithms Made Easy. *Journal of Machine Learning Research* 13 (jul 2012), 2171–2175.
- [27] Christian Gagné and Marc Parizeau. 2006. Genericity in Evolutionary Computation Software Tools: Principles and Case Study. *International Journal on Artificial Intelligence Tools* 15, 2 (April 2006), 173–194.
- [28] V. Grunert da Fonseca and C.M. Fonseca. 2002. A link between the multivariate cumulative distribution function and the hitting function for random closed sets. *Statistics & Probability Letters* 57, 2 (2002), 179–182. [https://doi.org/10.1016/S0167-7152\(02\)00046-9](https://doi.org/10.1016/S0167-7152(02)00046-9)
- [29] Nikolaus Hansen, Anne Auger, Raymond Ros, Olaf Mersmann, Tea Tušar, and Dimo Brockhoff. 2020. COCO: a platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* (2020), 1–31. <https://doi.org/10.1080/10556788.2020.1808977>
- [30] Frank Hutter, Youssef Hamadi, Holger H. Hoos, and Kevin Leyton-Brown. 2006. Performance Prediction and Automated Tuning of Randomized and Parametric Algorithms. In *Proc. of Principles and Practice of Constraint Programming (CP'06) (LNCS)*, Vol. 4204. Springer, 213–228. https://doi.org/10.1007/11889205_17
- [31] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2011. Sequential Model-Based Optimization for General Algorithm Configuration. In *Proc. of Learning and Intelligent Optimization (LION'11)*. Springer, 507–523.
- [32] Giorgos Karafotias, Mark Hoogendoorn, and A.E. Eiben. 2015. Parameter Control in Evolutionary Algorithms: Trends and Challenges. *IEEE Transactions on Evolutionary Computation* 19 (2015), 167–187.
- [33] Maarten Keijzer, J. J. Merelo, G. Romero, and M. Schoenauer. 2002. Evolving Objects: A General Purpose Evolutionary Computation Library. *Artificial Evolution* 2310 (2002), 829–888. <http://www.lri.fr/~marc/EO/EO-EA01.ps.gz> Latest release available on <https://nojhan.github.io/paradisEO/>.
- [34] Pascal Kerschke, Holger H. Hoos, Frank Neumann, and Heike Trautmann. 2019. Automated Algorithm Selection: Survey and Perspectives. *Evolutionary Computation* 27, 1 (2019), 3–45. https://doi.org/10.1162/evco_a_00242
- [35] Pascal Kerschke and Heike Trautmann. 2016. The R-Package FLACCO for exploratory landscape analysis with applications to multi-objective optimization problems. In *Proc. of IEEE Congress on Evolutionary Computation (CEC'16)*. IEEE, 5262–5269. <https://doi.org/10.1109/CEC.2016.7748359>
- [36] Per Kristian Lehre and Carsten Witt. 2012. Black-Box Search by Unbiased Variation. *Algorithmica* 64 (2012), 623–642.
- [37] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2016. Hyperband: A novel bandit-based approach to hyperparameter optimization. *arXiv preprint arXiv:1603.06560* (2016).
- [38] Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Mauro Birattari, and Thomas Stützle. 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* 3 (2016), 43–58.
- [39] Manuel López-Ibáñez, Marie-Eléonore Kessaci, and Thomas G Stützle. 2017. *Automatic Design of Hybrid Metaheuristic from Algorithmic Components*. Technical Report.
- [40] Manuel López-Ibáñez and Thomas Stützle. 2012. The Automatic Design of Multiobjective Ant Colony Optimization Algorithms. *IEEE Trans. Evol. Comput.* 16, 6 (2012), 861–875. <https://doi.org/10.1109/TEVC.2011.2182651>
- [41] Manuel López-Ibáñez and Thomas Stützle. 2014. Automatically improving the anytime behaviour of optimisation algorithms. *Eur. J. Oper. Res.* 235, 3 (2014), 569–582. <https://doi.org/10.1016/j.ejor.2013.10.043>
- [42] Nuno Lourenço, Francisco Pereira, and Ernesto Costa. 2012. Evolving Evolutionary Algorithms. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'12, Companion Material)*. ACM, 51–58. <https://doi.org/10.1145/2330784>

- 2330794
- [43] Marie-Eléonore Marmion, Franco Mascia, Manuel López-Ibáñez, and Thomas Stützle. 2013. Towards the Automatic Design of Metaheuristics. In *MIC 2013 - 10th Metaheuristics International Conference (Proceedings of the 10th Metaheuristics International Conference (MIC2013))*, Hoong Chuin Lau, Günther Raidl, , and Pascal Van Hentenryck (Eds.). Singapore, Singapore, 1–3. <https://hal.inria.fr/hal-01094692>
 - [44] Franco Mascia, Manuel López-Ibáñez, Jérémie Dubois-Lacoste, and Thomas Stützle. 2014. Grammar-based generation of stochastic local search heuristics through automatic algorithm configuration tools. *Comput. Oper. Res.* 51 (2014), 190–199. <https://doi.org/10.1016/j.cor.2014.05.020>
 - [45] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. 2011. Exploratory landscape analysis. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'11)*. ACM, 829–836.
 - [46] Antonio J. Nebro, Juan J. Durillo, and Matthieu Vergne. 2015. Redesigning the JMetal Multi-Objective Optimization Framework. In *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation (GECCO Companion '15)*. Association for Computing Machinery, New York, NY, USA, 1093–1100. <https://doi.org/10.1145/2739482.2768462>
 - [47] Federico Pagnozzi and Thomas Stützle. 2019. Automatic design of hybrid stochastic local search algorithms for permutation flowshop problems. *Eur. J. Oper. Res.* 276, 2 (2019), 409–421. <https://doi.org/10.1016/j.ejor.2019.01.018>
 - [48] Jérémy Rapin and Olivier Teytaud. 2018. Nevergrad - A gradient-free optimization platform. <https://GitHub.com/FacebookResearch/Nevergrad>.
 - [49] Conor Ryan, John James Collins, and Michael O'Neill. 1998. Grammatical evolution: Evolving programs for an arbitrary language. In *European Conference on Genetic Programming*. Springer, 83–96.
 - [50] Eric O. Scott and Sean Luke. 2019. ECJ at 20: toward a general metaheuristics toolkit. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'19, Companion Material)*. ACM, 1391–1398. <https://doi.org/10.1145/3319619.3326865>
 - [51] Kate Smith-Miles and Simon Bowly. 2015. Generating new test instances by evolving in instance space. *Comput. Oper. Res.* 63 (2015), 102–113. <https://doi.org/10.1016/j.cor.2015.04.022>
 - [52] Stefan Wagner, Gabriel Kronberger, Andreas Beham, Michael Kommenda, Andreas Scheibenpflug, Erik Pitzer, Stefan Vonolfen, Monika Kofler, Stephan Winkler, Viktoria Dorfer, and Michael Affenzeller. 2014. Architecture and Design of the HeuristicLab Optimization Environment. *Topics in Intelligent Engineering and Informatics*, Vol. 6. Springer, 197–261. http://link.springer.com/chapter/10.1007/978-3-319-01436-4_10
 - [53] Hao Wang, Diederick Vermetten, Furong Ye, Carola Doerr, and Thomas Bäck. 2020. IOHanalyzer: Performance Analysis for Iterative Optimization Heuristic. *CoRR abs/2007.03953* (2020). arXiv:2007.03953 <https://arxiv.org/abs/2007.03953>
 - [54] Thomas Weise, Yan Chen, Xinlu Li, and Zhize Wu. 2020. Selecting a diverse set of benchmark instances from a tunable model problem for black-box discrete optimization algorithms. *Appl. Soft Comput.* 92 (2020), 106269. <https://doi.org/10.1016/j.asoc.2020.106269>
 - [55] Thomas Weise and Zijun Wu. 2018. Difficult features of combinatorial optimization problems and the tunable w-model benchmark problem for simulating them. In *Proc. of Genetic and Evolutionary Computation Conference (GECCO'18, Companion Material)*. ACM, 1769–1776.
 - [56] Lin Xu, Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2012. Evaluating Component Solver Contributions to Portfolio-Based Algorithm Selectors. In *Proc. of Theory and Applications of Satisfiability Testing (SAT'12) (Lecture Notes in Computer Science)*, Vol. 7317. Springer, 228–241. https://doi.org/10.1007/978-3-642-31612-8_18
 - [57] Furong Ye, Carola Doerr, and Thomas Bäck. 2019. Interpolating Local and Global Search by Controlling the Variance of Standard Bit Mutation. In *Proc. of IEEE Congress on Evolutionary Computation (CEC'19)*. IEEE, 2292–2299.
 - [58] Furong Ye, Hao Wang, Carola Doerr, and Thomas Bäck. 2020. Benchmarking a $(\mu + \lambda)$ Genetic Algorithm with Configurable Crossover Probability. In *Proc. of Parallel Problem Solving from Nature (PPSN'20) (LNCS)*, Vol. 12270. Springer, 699–713. https://doi.org/10.1007/978-3-030-58115-2_49
 - [59] Martin Zaefferer and Frederik Rehbach. 2020. Continuous Optimization Benchmarks by Simulation. In *Proc. of Parallel Problem Solving from Nature (PPSN'20) (Lecture Notes in Computer Science)*, Vol. 12269. Springer, 273–286. https://doi.org/10.1007/978-3-030-58112-1_19