



HAL
open science

Automated Algorithm Selection for Radar Network Configuration

Quentin Renau, Johann Dreo, Alain Peres, Yann Semet, Carola Doerr,
Benjamin Doerr

► **To cite this version:**

Quentin Renau, Johann Dreo, Alain Peres, Yann Semet, Carola Doerr, et al.. Automated Algorithm Selection for Radar Network Configuration. GECCO '22: Genetic and Evolutionary Computation Conference, Jul 2022, Boston, United States. 10.1145/3512290.3528825 . hal-03718883

HAL Id: hal-03718883

<https://hal.science/hal-03718883>

Submitted on 13 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Automated Algorithm Selection for Radar Network Configuration

Quentin Renau
renau.quentin@gmail.com
Thales Research & Technology, École
Polytechnique, Institut Polytechnique
de Paris
Palaiseau, France

Johann Dreo
johann.dreo@pasteur.fr
CSB lab., Institut Pasteur[†]
Paris, France

Alain Peres
alain.peres@thalesgroup.com
Thales Land & Air Systems
Rungis, France

Yann Semet
yann.semet@thalesgroup.com
Thales Research & Technology
Paris, France

Carola Doerr
Carola.Doerr@lip6.fr
Sorbonne Université, CNRS, LIP6
Paris, France

Benjamin Doerr
Laboratoire d'Informatique (LIX),
CNRS, École Polytechnique, Institut
Polytechnique de Paris
Palaiseau, France

ABSTRACT

The configuration of radar networks is a complex problem that is often performed manually by experts with the help of a simulator. Different numbers and types of radars as well as different locations that the radars shall cover give rise to different instances of the radar configuration problem. The exact modeling of these instances is complex, as the quality of the configurations depends on a large number of parameters, on internal radar processing, and on the terrains on which the radars need to be placed. Classic optimization algorithms can therefore not be applied to this problem, and we rely on “trial-and-error” black-box approaches.

In this paper, we study the performances of 13 black-box optimization algorithms on 153 radar network configuration problem instances. The algorithms perform considerably better than human experts. Their ranking, however, depends on the budget of configurations that can be evaluated and on the elevation profile of the location. We therefore also investigate automated algorithm selection approaches. Our results demonstrate that a pipeline that extracts instance features from the elevation of the terrain performs on par with the classical, far more expensive approach that extracts features from the objective function.

CCS CONCEPTS

• **Applied computing** → *Industry and manufacturing*; • **Computing methodologies** → **Randomized search**.

KEYWORDS

radar network configuration, evolutionary computation, algorithm selection, exploratory landscape analysis

1 INTRODUCTION

Radar networks are placed to form detection barriers with the goal to detect incoming objects. The efficiency of these barriers, i.e., the network coverage or the detection probability has to be optimized taking into account several systemic and environmental constraints such as the areas to defend and exclusion areas in the domain (areas in which radars cannot be placed).

A radar network is composed of several radars. The radars may differ in type (e.g., with respect to the area or the distance they can sense). Radars are typically configurable, i.e., they have one or several parameters that determine their behavior. A common parameter of a radar device is the *tilt*, which is the angle between the horizontal plane and the direction of sensing. Optimizing a radar network therefore comprises the choice of the locations of the radars and their specific configuration. The number of parameters causes the problem to be intractable, and its modeling shows non-linear constraints, non-convexity and non-differentiability, making it difficult to be analytically described.

Solving the radar network configuration problem therefore relies on *black-box optimization*, where the problem instances do not need to be explicitly modeled but where it suffices that the quality of potential solutions can be assessed through a simulator that evaluates how suitable the suggested configuration is.

In the literature and in industrial contexts, metaheuristics have been widely investigated and used to solve numerical black-box optimization problems. Metaheuristics are stochastic methods that do not need a priori information on the function to solve. Even if these methods have no optimality guarantee, they have been well performing on different types of problems, from traffic congestion [3] to RNA design [19].

In this paper, we seek to maximize the coverage of a terrain using a radar network. The network is composed of four radars of two types that differ from their parameters.

In order to find the locations and configuration of the radars, we compare the performance of 13 metaheuristics on 153 problem instances. We also investigate the performance of an automated algorithm selection procedure by building two selectors. One selector follows the traditional approach and is built using landscape feature extracted on the radar objective function. The second selector is built extracting landscape features on the physical landscape only.

In this paper, we exhibit the complementarity of metaheuristics performances on the radar network configuration problem. We observe that different algorithms perform well on different instances. We also observe that different algorithms perform well on different budget of function evaluations.

We also compare the two selectors performances and find that they have similar performances while one is computationally much cheaper to build than the other.

[†] Full affiliation for J. Dreo: *Institut Pasteur, Université Paris Cité, Bioinformatics and Biostatistics Hub, Department of Computational Biology, Computational Systems Biomedicine laboratory, F-75015 Paris, France.*

Paper organization: The radar configuration problem is introduced in Sec. 2. Specifics about our experimental setup are provided in Sec. 3, and the results of our experiments are presented in Sec. 4. The comparison with the human expert is made in Sec. 5. Finally, Sec. 6 presents the two automated algorithm selection pipelines and analyzes their performance. Our paper is concluded in Sec. 7 with directions for future work.

Reproducibility: While we cannot make the implementation of the radar configuration problem available in open-access format, we provide binaries of the radar network configuration use-case [27]. It contains the objective function presented in this paper, thus it is possible to recreate the 153 instances presented in Sec. 3.1 and to generate any instance on any terrain. The detailed performance of algorithms for 500 and 2,500 function evaluations are available at [27]. They follow the data format of the IOHprofiler [5] environment and can hence be easily analyzed and visualized by the IOHanalyzer module [32] of the IOHprofiler project [5], available online at <https://iohanalyzer.liacs.nl/>. Some of the plots in this paper were also created with IOHanalyzer.

2 PROBLEM DESCRIPTION

Radar configuration problems differ with respect to the object to detect, the number and types of the radars that can be placed, their parameters, the terrain they shall protect, and with respect to constraints on the locations in which the radars can be placed. We briefly describe these aspects in the following paragraphs.

Object Characteristics. For the sake of simplicity, we consider a simple object model to be detected by the radars. We suppose the speed of the object to be constant. The object altitude above ground level h is also supposed to be constant, i.e., if the ground altitude above sea level is z , then the object is flying above the ground at altitude $z + h$.

We define the object angle θ as the angle between the object direction and the North. This angle is called the *azimuth*.

Terrain Definition. The terrain to cover is a square of 50 km \times 50 km and for each point, the object can face in any direction. Therefore, the domain Δ to cover is represented by voxels (x, y, θ) , where θ is the object azimuth. For our model, we split each axis of Δ into 30 equal parts, resulting in a total number of $30^3 = 27,000$ voxels that need to be protected.

The geographical data used to model the terrain is coming from the NASA Shuttle Radar Topography Mission [9] (SRTM). These digital elevation models (DEM) cover the entire globe and are provided in mosaics of 5 degrees by 5 degrees tiles. In order to use complete data, we used DEM that were post-processed with interpolation in [12] (Data available at <https://srtm.csi.cgiar.org>).

Radar Network Modeling. We consider four radars R of two types: one rotating and three staring radars. These two types only differ by the angular domain that the radar can sense. In our simulations, we assume that radars with rotating antennas can sense all around them at any time, whereas radars with staring antennas can only sense predefined regions around their *staring direction*.

The number of tunable parameters depends on the type of the radar: the staring radars have four parameters that can be modified:

- (1) the x (\approx longitude) location in the terrain;
- (2) the y (\approx latitude) location in the terrain;

- (3) the *tilt*, i.e., the angle between the horizontal plane and the direction of sensing;
- (4) the *staring angle*, i.e., the direction the radar will look to.

For rotating radars, the staring angle is not needed and we therefore have three parameters to tune. The dimension of our radar configuration problem is thus $d = 3 + 3 \times 4 = 15$.

After normalization, our decision space is modeled as $[0, 1]^{15}$, i.e., each choice ξ of radar location and configuration (*network* in the following), is a 15-dimensional vector with entries in $[0, 1]$.

Detecting is not a situation where the object is either detected or not detected. The detection relies on a performance indicator which is called the *probability of detection*. Computing the probability of detection of one radar is very complex and depends on the radar type and its processing along as the object to detect.

Multiple radars are gathered into a *network*. A network is used to aggregate the probabilities of detection for each radars. In this paper, the probability of detection of the network is given by multiplying the probabilities of non-detection for each radar and returning the corresponding probability of detection. We assume that the radars are independent and compute:

$$P_d^\xi(x, y, \theta) := 1 - \prod_{i=1}^4 (1 - P_d^{R_i}(x, y, \theta)),$$

where $P_d^{R_i}(x, y, \theta)$ is the probability of detection of the radar R_i given its location x, y and the azimuth θ of the object.

Objective Function. The objective function aims at maximizing the coverage of Δ . For a given problem instance (i.e., for a given area T to cover), our objective is to find the location and the configuration of the radars such that the number of covered voxels is as large as possible. For a given radar network ξ , we consider a voxel (x, y, θ) covered if the probability of detection $P_d^\xi(x, y, \theta)$ is greater than or equal to a predefined fixed threshold τ .

For a given problem instance T , the objective function for the coverage scenario is then defined as $f_T : [0, 1]^{15} \rightarrow [0 \dots 27,000]$:

$$f_T(\xi) := \sum_{(x, y, \theta) \in [1 \dots 30]^3} \mathbb{1} \left(P_d^\xi(x, y, \theta) \geq \tau \right).$$

Our objective is thus to find the optimal radar network ξ^* maximizing $f_T(\xi)$:

$$\xi^* = \arg \max_{\xi \in \Xi} f_T(\xi).$$

3 DESIGN OF EXPERIMENT

In this section, we briefly describe the 153 problems instances that we selected for our experiments and the 13 algorithms that we used to solve these instances.

3.1 Problem Instances

Each instance is defined by an individual terrain (see Sec. 2). Overall, 17 different tiles of the digital elevation model (DEM) from all around the world were selected manually, so as to represent a large variety of terrains.

The size of our use-case (see Sec. 2) is smaller than the area of the DEM tiles. To match the dimensions of our use-case, we therefore subsampled the original tiles domains. As there is no particular reason to favor one area over another, we applied the

Table 1: Difference between the highest and lowest altitudes (in meters) of each instance. Columns represent the subsample in the DEM tile and rows corresponds the region of the DEM tile.

	Instance								
	0	1	2	3	4	5	6	7	8
afghan.	1616	1266	1152	1161	1534	1177	1085	1144	1023
argentina	24	24	24	24	27	22	26	26	27
australia	189	181	117	67	376	93	203	384	90
belarus	166	196	145	158	135	187	154	166	147
brasil	61	79	85	108	89	90	129	58	107
canada	1494	1571	1419	1452	1528	1194	1474	1515	1387
chile	1408	1311	1555	2625	1252	2469	1869	1330	2523
china	1417	1439	1462	893	1398	1039	924	1271	956
congo	66	70	68	84	64	74	88	64	87
france	161	329	300	325	331	177	189	186	307
india	3828	2111	3821	2224	1983	3195	3175	2070	3755
iran	914	487	841	929	788	788	986	319	979
moldavia	111	245	241	235	257	166	170	187	229
nepal	1439	1753	1232	1656	1723	1367	1529	1644	1264
russia	341	302	341	343	302	327	355	330	322
sahara	25	25	24	43	20	28	25	27	29
usa	1161	1486	1518	1300	1321	1049	1173	1393	1228

same downsampling mask to each DEM tile. The mask was created by choosing the upper left $50 \text{ km} \times 50 \text{ km}$ square and by sampling eight additional squares of the same area uniformly at random, obeying to the condition that the areas shall be pairwise disjoint (i.e., non-overlapping). This mask provides us with nine subsamples for each terrain. The result is a total number of $9 \times 17 = 153$ instances for our radar network configuration problem.

The instances are labeled by the region name of the DEM tile and by the number of its sample in the mask, i.e., the first sample in the mask for the Brasil region has the instance name *brasil0*. Each problem instance was generated using the objective function available in [27].

The difference between the highest and the lowest points in each instance is reported in Tab. 1. This information gives a hint of the actual landscape of the terrain. It also gives a first idea about which problems might be easier to solve than others, as, intuitively, flat areas should be easier to cover than mountainous ones.

We define *flat instances* to be instances where the difference between the highest and lowest point is below 100 meters. *Mountainous instances* have a difference between the highest and lowest points greater than 1,000 meters. We have a total number of 36 flat instances, 60 mountainous instances, and 57 intermediate ones.

3.2 Algorithm Portfolio

The portfolio of algorithms is composed of 12 of the algorithms with their default implementation parameters:

- Differential Evolution [29] with *scipy* [31] (version 1.5.4) implementation;
- Nelder-Mead [22] with *scipy* implementation;
- Powell’s method [24] with *scipy* implementation;
- L-BFGS-B [6] with *scipy* implementation and a finite differences step of 0.01;
- Particle Swarm Optimization (PSO) [14] with global best topology and *pyswarms* [21] (version 1.3.0) implementation;
- Random Search with uniform sampling;

Table 2: Differential Evolution population size, mutation and recombination parameters. Comparison of default algorithm parameters with tuned algorithm parameters.

	DE (default)	DE_2500_chile
popsize	15	6
mutation	$\mathcal{U}(0.5,1)$	0.1
recombination	0.7	0.58

- Quasi-Random Search with Sobol’ [28] low-discrepancy sequences;
- Five CMA-ES variants from the ModCMA framework [30]:
 - Vanilla CMA-ES [8] (configuration 0000000000 in the Mod-CMA framework);
 - CMA-ES with mirrored sampling [4] and pairwise selection [1] (configuration 00100001000);
 - CMA-ES with elitism (configuration 0100000000);
 - CMA-ES with active update [13] (configuration 1000000000);
 - CMA-ES with active update, elitism and BIPOP increasing population size [7] (configuration 11000000002).

We expand our portfolio of algorithms by performing algorithm configuration on Differential Evolution on one mountainous instance, the *chile0* instance, in order to create specific solvers. To this end, we use the *irace* [18] (version 3.4.1) configurator with default parameters and the lowest budget possible of runs: 180. This choice of number of runs is motivated by the expensive computation time of the objective function. Running *irace* on this single use-case takes around 24 hours on our computer hardware. We also tune other algorithms on single or multiple instances, but the results of tuned algorithms were comparable to their default configuration.

The parameters of Differential Evolution to tune are the population size (between 4 and 100), the mutation (between 0 and 2) and, recombination constants (between 0 and 1). Compared to the default setting, *irace* suggests a much smaller population size and a smaller mutation constant (see Tab. 2). The tuned version of the DE algorithm enters our portfolio of algorithms under the name *DE_2500_chile*.

3.3 Experimental Setup

We run each algorithm 30 independent times on each problem instance. For each of these runs, we log the whole performance trajectory, i.e, we log all the improvements achieved by the algorithms up to a budget of 2,500 function evaluations. For the analyses presented in this paper, we mainly focus on one **low-budget setting with 500 function evaluations** and one **large-budget setting with 2,500 function evaluations**.

As most default optimization algorithm implementations assume minimization, we seek to find a network configuration that minimizes the non-covered areas.

4 ALGORITHM PERFORMANCES

In this section, we discuss the performances of the algorithms from our portfolio when applied to the use-case defined in Sec. 2. We analyze the performance distribution over the 30 runs per instance in Sec. 4.1 and regard the median performance over time in Sec. 4.2.

A statistical comparison in Sec. 4.3 underlines that these runtime distributions are significantly different, which in principle suggests to use algorithm configuration and selection. In Sec. 4.4, we give the performance of the different algorithms aggregated over all 30 runs on all 153 instances and derive from this a single best solver.

4.1 Individual Runs

Figures 1 show the histograms of the objective values reached in the 30 runs of the 13 algorithms on two terrains and using the small budget of 500 evaluations. Within a terrain type, the scaling of the x axis and the bins used in the histograms are identical each algorithm; they differ between the terrain types because of the significantly different objective values obtained – naturally, the radar network configuration problem is much harder in the rugged *nepal3* terrain than in the flat *sahara0* terrain.

At a first glance, all algorithms have a reasonable performance. The simple approaches of random and quasirandom search, mostly given for comparison, naturally are often beaten by the more advanced algorithms, in particular, on the less rugged terrain.

On both instance, Powell’s method found the best solution most often. However, given the dispersion of solution values, Powell’s method is not the best performing algorithm in median on either instance. Here, some DE variant, Nelder-Mead (on the *sahara0* instance) or, DE_2500_chile are superior.

4.2 Influence of the Budget

The relative performance of the algorithms not only depends on the instance type, but also crucially on the computational budget. This is visible from Fig. 2, showing for each algorithm how the median objective values out of 30 runs develops over time. On both instances, the CMA-ES variant *00100001000* performs very well from around 1,000 function evaluations on, whereas for smaller budgets other algorithms are significantly better, e.g., DE_2500_chile.

The differences between the algorithms are less pronounced for the *sahara0* instance, i.e., around a 300 voxels difference between the best and 5-th best algorithm after 2,500 iterations over a domain size of 27,000 voxels. For the more rugged *nepal3* landscape, however, the differences are noteworthy as, apparently, the different algorithms get stuck in local optima of very different solution quality.

Budget and Elitism. The results presented in Sec. 4.2 seem to show that some CMA-ES variants are preferable for small budgets whereas others perform well on large budgets.

The main differences between these variants are their respective selection mechanism. Variants that perform well for the low budget are elitist, i.e., they use plus-selection, whereas variants that perform well for the larger budget are non-elitist, i.e., they use comma-selection.

Fig. 3 illustrates this behavior on the *sahara0* (Fig. 3a) and the *nepal3* instances (Fig. 3b) of two CMA-ES variants. One variant is the non-elitist vanilla CMA-ES (0000000000) and the other variant uses the elitism module (0100000000). On both instances, we observe that the elitist variant performs better on lower budgets. In contrast, the non-elitist variant is performing better for larger budgets. On all instances, there is a number of function evaluations where the non-elitist performance curve crosses the elitist one.

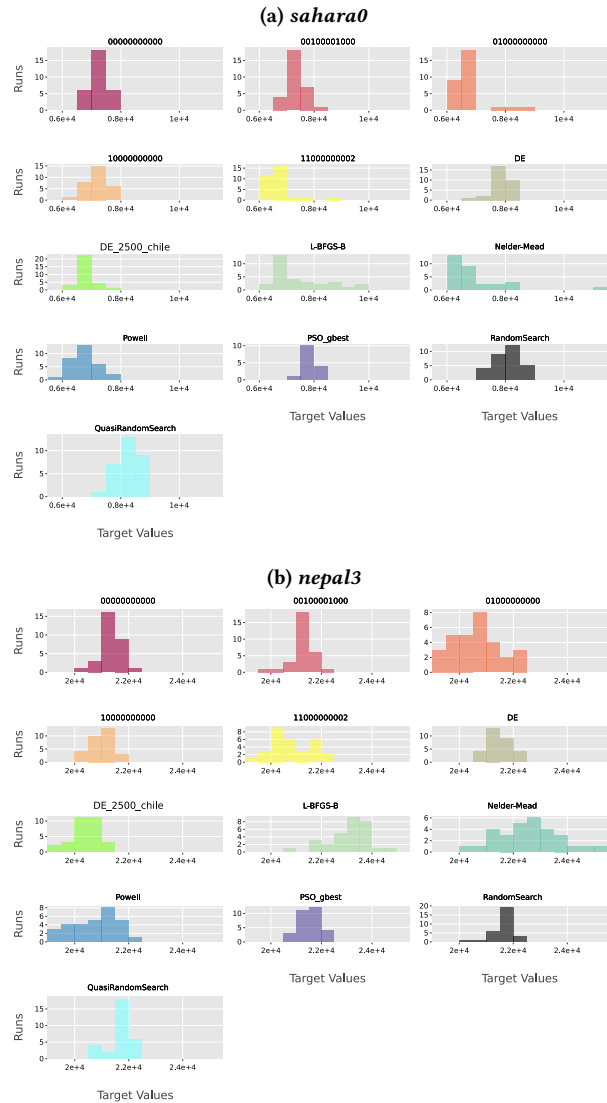


Figure 1: Histograms (over 30 runs) of the objective values reached on the unconstrained use-case in 500 function evaluations.

The crossing point occurs around the 900 evaluations on the *sahara0* and around 1,300 evaluations on *nepal3*. Overall, this crossing occurs on all instances between 600 and 1,300 evaluations. Nevertheless, the crossing seems not to depend of the landscape of the terrain, i.e., the variants performances may cross each other later on some flat instances than mountainous instances and vice versa. As an example, the crossing occurs around 1,300 evaluations on the *brasil1* instance which is rather flat and around 750 evaluations on the *chile6* instance.

4.3 Statistical Significance

All algorithms seem to have similar performances. To confirm this behavior, we perform a Kolmogorov-Smirnov test to determine if

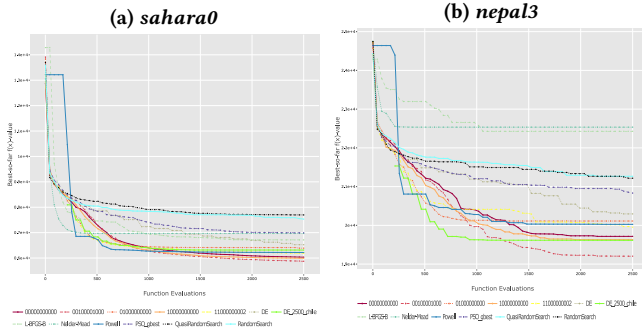


Figure 2: Median objective values over time, for 30 runs, for *sahara0* and *nepal3* instances, for the 13 algorithms of our portfolio.

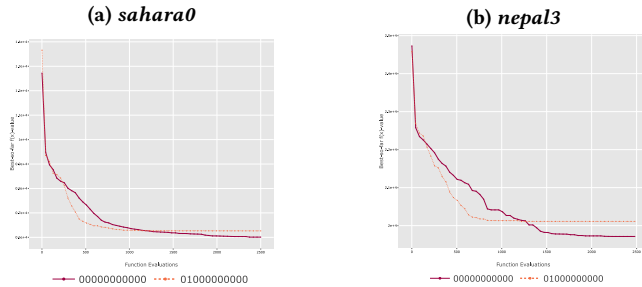


Figure 3: Median objective values over time, for 30 runs of vanilla (red) and elitist (orange) CMA-ES variants with different selection mechanism.

one algorithm outperforms all others on each instance. We perform this test for each pair of algorithms with a confidence level $\alpha = 0.01$.

In the *nepal3* instance, for a budget of 500 function evaluations, the best performing algorithm, *DE_2500_chile*, is statistically better than the majority of algorithms except for *Powell*, and the *11000000002* and *01000000000* CMA-ES variants.

For a larger budget of 2,500 function evaluation, *Powell*, *DE_2500_chile* and, all CMA-ES variants perform better than other algorithms but are do not show statistically different performances between them.

This tendency of multiple algorithms having similar performances can also be seen for *sahara0* instance and all other instances in the benchmark (see [27])

4.4 Single Best Solvers (SBS)

The results of the previous subsections show that the different algorithms have substantially different strengths. This is what motivated us to use an automated algorithm selection approach. Before doing so, we nevertheless try to derive some insight on which of the algorithms in some general sense are superior.

For this comparison, we again regard the median performance of an algorithm on an instance. This defines, for each instance and each of the two budgets a best-performing algorithm.

We then compare each algorithm on each instance with the performance of this best-performing algorithm. Fig. 4 displays the distribution of the solution quality loss (in percent) of each algorithm. We see, for example, that for a computational budget of 500, *DE_2500_chile* obtained the best result on 46% of the instances (that is, on 69 out of 153 instances), and that it never encountered a loss of more than 15%. Other generally good algorithms for this budget are the CMA-EA variants *11000000002* and *01000000000* as well as *Powell's method*.

To define a *single best solver*, denoted by *SBS*₅₀₀ and *SBS*_{2,500}, we compute the median (over 153 instances) performance loss of each algorithm. This loss corresponds to the difference between the performance of the best algorithm on a given instance with the performance of the target algorithm. For the budget of 500 evaluations, this number is smallest for *DE_2500_chile*, namely as little as 0.39%. For a budget of 2,500, apparently the ES variant *00100001000* with a median performance loss of 0% is the *SBS*_{2,500}. This algorithm performed best on 86 of the 153 instances.

5 COMPARISON WITH MANUAL OPTIMIZATION

In practice, the configuration of radar networks is often hand-designed by experts with the help of a simulator. The objective of this section is to compare hand-designed solutions with those of the 13 algorithms from our portfolio.

Problem Instance. The problem considered to be solved by hand is the *canada0* instance. The landscape can be divided into two areas: a rather flat area on the left and a peaked area with mountains on the right.

Algorithm Performances on *canada0*. The best performing algorithm on this problem is vanilla CMA-ES (00000000000) with a median number of 15,169 voxels that are not covered and a best run of 14,125 voxels not covered for 2,500 function evaluations.

For a smaller budget of 500 function evaluations, the CMA-ES variant 11000000002 achieves the best median performance with a fitness of 16,518. *Powell's method* performs the best run with a value of 15,841.

For comparison, the median value of random search for the large budget is 17,865 and it is 18,545 for the small budget.

Radar Network Configuration Contest. In order to gather a sufficient amount of data to compare hand-designed solutions to algorithms, we launched a contest to solve the *canada0* instance. The competition had 13 participants of diverse expertise in radar network configuration, ranging from engineers with significant radar & optimization expertise, to first-year PhD students.

At first, we presented the landscape of the instance to the contestant. Then, we introduced them the goal, i.e., covering as much space as possible and the resources to reach that goal, i.e., the four radars and their corresponding parameters.

Participants has 30 minutes to solve the problem, and they could evaluate as many solutions as they wanted within this time frame. For comparison, the average algorithm wallclock time for one run of 2,500 function evaluations is around five minutes. For each contestant, we recorded the best value found and the improvements done during the optimization.

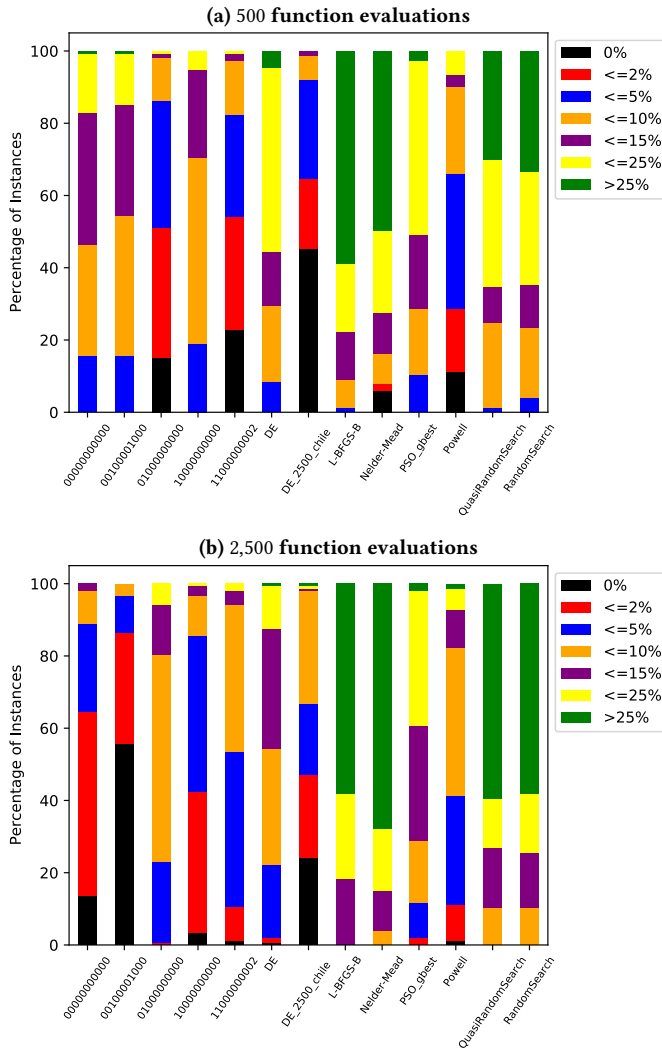


Figure 4: Percentages of instances solved within $X\%$ of the best algorithm performance.

Comparison of Hand-Designed Configurations with Algorithms. The median fitness value obtained by the group is 17,495 voxels not covered. The best result is a configuration that reached fitness value 16,425. A few statistics comparing the human-designed results with those of the algorithms are summarized in Fig. 5.

In these results, the median performance of the human-designed configuration is around 2% better than random search. Overall, we see that the optimization algorithms tend to perform much better than the participants of our competition. The median performance of our participants is 6% worse than the median of best performing algorithm on the budget of 500 function evaluations and 15% worse on the budget of 2,500 function evaluations. Moreover, the *canada0* instance seems to be a hard instance for manual optimization, as human performance is close to that of random search.

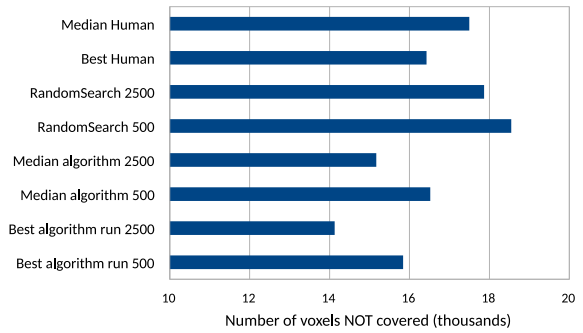


Figure 5: Comparison of performance between solvers and human results.

6 AUTOMATED ALGORITHM SELECTION

In this section, we present how we create our training and testing sets, how we extract landscape features to build our selectors, how we create the mapping between feature data and algorithms performances presented in Sec. 4.4, and how we assess the selectors performances.

6.1 Training and Testing Sets

Training and testing sets are created randomly using all 153 available problem instances. The training set is composed of 80% of the instances, and the remaining 20% compose the testing set. The instances are selected uniformly at random without replacement. As this random selection can be biased and not representative of the full set of instances, we perform five independent runs to assess the selectors.

6.2 Feature Computation

We describe the two ways that we used to compute problem instance features. Each feature is a numerical value that measures a certain aspect of the problem landscape. The features are computed using *exploratory landscape analysis* [20]. To this end, sample points need to be drawn in the domain and the fitness of the sample points evaluated. Features are approximated using the pairs of sample points and the associated fitness values.

We experimented with two different ways of extracting features for our problem instances: the traditional one that computes features on the objective function, and a second one that computes the features on the DEM, i.e., without involving the actual radar configuration simulator. In both cases, we use the *flacco* package [17] to extract the problem instances features. We computed all cheap features available in *flacco*, except for the PCA features which do not provide sufficient information when used with non-adaptive sampling (see [26] for a discussion). We also excluded NBC features [15] as they are not robust to the grid sampling of the DEM. All experiments are thus based on 33 features in total.

Objective Function. For extracting the features on the objective function, we used Sobol low discrepancy sampling [28], following a suggestion made in [25]. Following common recommendation from the literature [16], we base the feature computation on $50d = 750$ samples. To compensate for the randomness in the sampling, we

perform this feature extraction step 100 independent times. The 100 runs of the features computed on the objective function are aggregated using their *median value*. Hence, for each problem, one vector corresponding to the median values of each feature is used to construct the mapping.

Digital Elevation Model (DEM). Building on the idea that the structure of the physical landscape is the factor that has the biggest impact on the objective function, the features can also be computed on the DEM, rather than via the objective function.

In order to compute features, we need to have an objective function. In this case, we take the altitude as a function, i.e., we define a function $g(x, y) = z$, where z is the altitude of the point (x, y) in the domain. Hence, unlike the radar network configuration problem, the dimension of this problem is $d = 2$. On top of that, the domain is discretized, which implies that the number of different altitudes that we can use for the feature computation is at most the number of pixels which is equal to $30 \times 30 = 900$.

This approach has two main *benefits*: first, it avoids the rather expensive evaluation of the objective function. Another advantage, as the domain is discretized, is that we can fully sample the search space. Doing so, we can remove the randomness of the sampling.

The *drawback* of this approach is the loss of information. Even though the altitude seems to be the factor that has the biggest impact on the objective function, it is not the only one. While computing features using the objective function consider all radar parameters (*tilt*, *staring angle*, and internal processing), computing features on DEM has an exclusive focus on the elevation profile of the instance. Disregarding radar parameters may imply some loss of information on the problem to solve.

In the following, the features computed on the DEM are referred to as DEM features.

6.3 Building the Mapping between Feature Data and Algorithm Performances

Using the features extracted by the above-described procedure, we build selectors for the low-budget case and for the large-budget case. To this end, we link the algorithm performance to the feature representation of the instances. We build our model using the median performance of the 30 independent runs (see Sec. 4.4).

We denote with S_{radar} the selector built with features computed on the objective function and by S_{DEM} , we denote the selector built with features extracted from the DEM.

As previously done in the literature [2, 10], we build the mapping between feature data and algorithm performances using the default scikit-learn Random Forests *regression* [23]. For each algorithm, we create one Random Forests regression model.

Each regression model learns the performances of the associated algorithm for a given feature vector. The selector is composed of all the regression models. When the selector has an unknown feature vector as input, every regression model predicts the performance of its associated algorithm, and the selector then ranks the algorithms by their predicted performance.

6.4 Definition of the SBS

To perform a landscape-aware algorithm selection, we need to split our 153 instances into a training and a testing set (see Sec. 6.1). The

splitting can have an impact on the definition of the SBS. Training splits composed of a majority of mountainous instances may have a different SBS than training splits composed of a majority flat instances.

To visualize the impact of the splits, we create 1,000 independent splits of training instances and record the SBS on both budgets. On the lower budget of 500 evaluations, three algorithms were SBS on the splits. 11000000002 was SBS on 503 splits, DE_2500_chile was SBS on 394 splits, and 01000000000 on 103 splits. On the use-cases, we will consider as SBS_{500} the 11000000002 CMA-ES variant as this algorithm was more often SBS on the different splits.

Concerning the larger budget, two algorithms were SBS on the splits. The CMA-ES variant 00100001000 performed best for 986 splits and DE_2500_chile on the remaining 14 splits. The $SBS_{2,500}$ is the same for almost all the different splits.

6.5 Selector Performances

We will show in this section the results of our two selectors and how they compare to the virtual best solver (VBS, i.e., the hypothetical optimal selector that always selects the best algorithm for each instance), the SBS, and a common baseline: the vanilla CMA-ES (00000000000).

We saw in Sec. 4.4 that algorithms have similar performances. This implies that the VBS-SBS gap is small and that the margin to perform landscape-aware algorithm selection is small.

On the low budget, the selectors outperform the baseline vanilla CMA-ES by around 7% in median performance. On the larger budget, selectors and vanilla CMA-ES have equivalent performances, i.e., at around 2.5% of the VBS in median performance (see Fig. 6).

This could be expected from what was mentioned in Sec. 4.4. Vanilla CMA-ES was not performing well for low budget settings and needed more function evaluations to converge. On the larger budget, it was one of the best performing algorithm so we expected its performances to be close to those of our selectors.

Overall, the selector perform worse than the SBS for both budgets at 2.6% from the VBS in median. We recall that for this use-case, the median gap between the SBS and the VBS is 0.39% on low budget and 0% for the large budget.

6.6 Performances of S_{radar} versus S_{DEM}

The performances of S_{DEM} and S_{radar} are similar. The median performance of S_{DEM} is within 0.2% of the median performance of S_{radar} on both budgets of function evaluations.

Given the small difference in performances, the DEM selector can be preferred to S_{radar} as its building and using are almost free in computation time. More precisely, the computation time of the selector S_{radar} is around 2.5 minutes when it only takes one or two seconds to use S_{DEM} . This difference is a consequence of the evaluation of the objective function samples and the feature computation.

6.7 Discussion

Performances of our selectors are slightly worse than the SBS. Nevertheless, their performances are quite close to the VBS, i.e., around 2.5% in median for both budgets.

The performance of our selectors may be improved by tuning the Random Forests or by using alternative regression techniques [11].

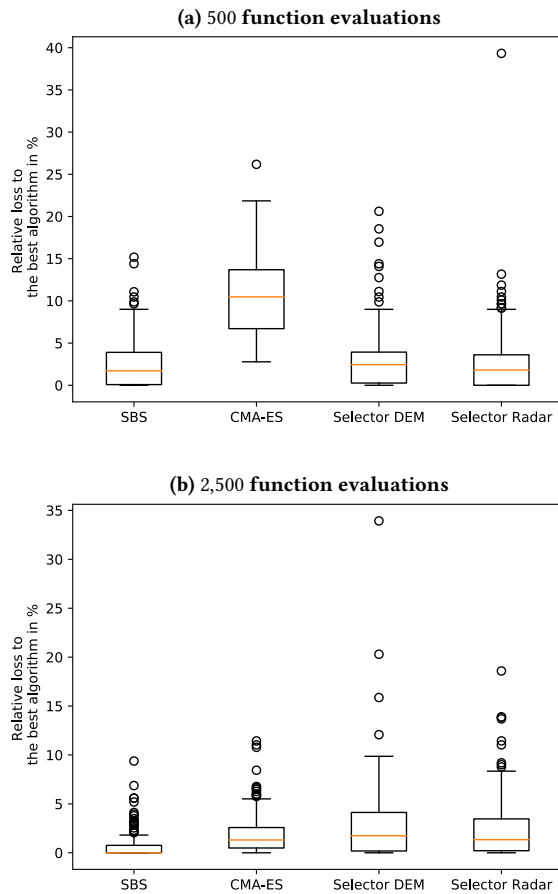


Figure 6: Relative loss in percentage of the SBS, vanilla CMA-ES, and the selectors with respect to the VBS, assuming that only the best algorithm is proposed.

Nevertheless, even if the performances of the selectors increase, the possible gain is relatively low.

Given the performances of the selectors compared to those of the SBS, it would be beneficial to use only the SBS to solve this use-case. Selectors are not performing badly but the VBS-SBS gap is too small and does not benefit an automated algorithm selection procedure. This may be different for different radar configuration problems or other algorithm portfolios. Overall, we consider our results as encouraging for the application of landscape-aware algorithm selection techniques in industrial contexts.

When there is time to run multiple algorithms, it could be beneficial to run both the SBS and the algorithm recommended by the selector. For some instances, the algorithm predicted by the selector may have better performances than the SBS.

7 CONCLUSION

We have analyzed in this work the performance of 13 metaheuristics on a radar configuration use-case with 4 radars. For this use-case, we provide binaries [27] that allow to reproduce our experiments

and that allow to create any benchmark of instances using the model presented.

We have shown that the heuristics outperform human radar configurations, while at the same time requiring much less time than the manual configuration. Our results encourage the use of standard black-box optimization heuristics to aid human experts in the design of radar networks.

We have also shown that the budget of function evaluations has a major influence on which algorithms to prefer: while elitist versions of the CMA-ES performed well for low budgets, they are outperformed by their elitist counterparts for large budgets.

We have also shown that a landscape-aware algorithm selection achieves good performance, indicating that the features provided by *flacco* provide useful information to distinguish the problem instances. We have also shown that it suffices to extract features on the DEM. Since this is computationally much cheaper than feature extraction on the objective function, this may pave a way towards real-world deployment of landscape-aware algorithm configuration and selection for radar network configuration.

Extending our pipeline towards more specialized solvers is therefore a straightforward next step of our research. We shall also include other meta-heuristics, e.g., surrogate-based or assisted ones.

Our experiments are moreover based on a research-style radar configuration model. In industrial contexts, the model complexity often increase with the maturity of the product. We suspect that the difference between the VBS and the SBS may increase with increasing model complexity, in which case a landscape-aware selection of algorithms may prove more beneficial than in our use-case studied above especially with its ability to better manage a limited evaluations budget.

ACKNOWLEDGMENTS

This work was supported by a public grant as part of the Investissements d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH, and by the Paris Ile-de-France Region.

REFERENCES

- [1] A. Auger, D. Brockhoff, and N. Hansen. 2011. Mirrored sampling in evolution strategies with weighted recombination. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '11*. ACM, 861–868. <https://doi.org/10.1145/2001576.2001694>
- [2] N. Belkhir. 2017. *Per Instance Algorithm Configuration for Continuous Black-Box Optimization*. Ph.D. thesis. Université Paris-Saclay. <https://hal.inria.fr/tel-01669527/document>
- [3] M. Böther, L. Schiller, P. Fischbeck, L. Molitor, M.S. Krejca, and T. Friedrich. 2021. Evolutionary minimization of traffic congestion. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*. ACM, 937–945. <https://doi.org/10.1145/3449639.3459307>
- [4] D. Brockhoff, A. Auger, N. Hansen, D. V. Arnold, and T. Hohm. 2010. Mirrored Sampling and Sequential Selection for Evolution Strategies. In *Parallel Problem Solving from Nature - PPSN XI, 11th International Conference, September 11-15, 2010, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 6238)*. Springer, 11–21. https://doi.org/10.1007/978-3-642-15844-5_2
- [5] 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/>.
- [6] R. Fletcher. 1987. *Practical Methods of Optimization; (2nd Ed.)*. Wiley-Interscience, USA.
- [7] N. Hansen. 2009. Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '09 (Companion)*. ACM, 2389–2396. <https://doi.org/10.1145/1570256.1570333>

- [8] N. Hansen and A. Ostermeier. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation* 9, 2 (2001), 159–195. <https://doi.org/10.1162/106365601750190398>
- [9] T. Hennig, J. Kretsch, C. Pessagno, P. Salamonowicz, and W. Stein. 2001. The Shuttle Radar Topography Mission. In *Proceedings of the First International Symposium on Digital Earth Moving (DEM '01)*. Springer, 65–77.
- [10] A. Jankovic and C. Doerr. 2020. Landscape-Aware Fixed-Budget Performance Regression for Modular CMA-ES Variants. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '20*. <https://doi.org/10.1145/3377930.3390183> To appear.
- [11] A. Jankovic, G. Popovski, T. Eftimov, and C. Doerr. 2021. The impact of hyperparameter tuning for landscape-aware performance regression and algorithm selection. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*. ACM, 687–696. <https://doi.org/10.1145/3449639.3459406>
- [12] A. Jarvis, H.I. Reuter, A. Nelson, and E. Guevara. 2008. Hole-filled seamless SRTM data V4. *International Centre for Tropical Agriculture (CIAT)* (2008). <https://srtm.csi.cgiar.org>
- [13] G.A. Jastrebski and D.V. Arnold. 2006. Improving Evolution Strategies through Active Covariance Matrix Adaptation. In *IEEE International Conference on Evolutionary Computation, CEC 2006, part of WCCI 2006*. IEEE, 2814–2821. <https://doi.org/10.1109/CEC.2006.1688662>
- [14] J. Kennedy and R. Eberhart. 1995. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, Vol. 4. 1942–1948 vol.4. <https://doi.org/10.1109/ICNN.1995.488968>
- [15] P. Kerschke, M. Preuss, S. Wessing, and H. Trautmann. 2015. Detecting Funnel Structures by Means of Exploratory Landscape Analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '15*. ACM, 265–272. <https://doi.org/10.1145/2739480.2754642>
- [16] P. Kerschke, M. Preuss, S. Wessing, and H. Trautmann. 2016. Low-Budget Exploratory Landscape Analysis on Multiple Peaks Models. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '16*. ACM, 229–236. <https://doi.org/10.1145/2908812.2908845>
- [17] P. Kerschke and H. Trautmann. 2019. Comprehensive Feature-Based Landscape Analysis of Continuous and Constrained Optimization Problems Using the R-Package Flacco. In *Applications in Statistical Computing: From Music Data Analysis to Industrial Quality Improvement*. Springer, 93–123.
- [18] M. López-Ibáñez, J. Dubois-Lacoste, L. Pérez Cáceres, M. Birattari, and T. Stütz. 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* 3 (2016), 43 – 58.
- [19] N. S. C. Merleau and M. Smerlak. 2021. A simple evolutionary algorithm guided by local mutations for an efficient RNA design. In *Genetic and Evolutionary Computation Conference, GECCO '21*. ACM, 1027–1034. <https://doi.org/10.1145/3449639.3459280>
- [20] O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, and G. Rudolph. 2011. Exploratory Landscape Analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '11*. ACM, 829–836. <https://doi.org/10.1145/2001576.2001690>
- [21] L.J.V. Miranda. 2018. PySwarms: a research toolkit for Particle Swarm Optimization in Python. *Journal of Open Source Software* 3, 21 (2018), 433. <https://doi.org/10.21105/joss.00433>
- [22] J. A. Nelder and R. Mead. 1965. A Simplex Method for Function Minimization. *Comput. J.* 7, 4 (01 1965), 308–313. <https://doi.org/10.1093/comjnl/7.4.308> arXiv:<https://academic.oup.com/comjnl/article-pdf/7/4/308/1013182/7-4-308.pdf>
- [23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [24] M. J. D. Powell. 1964. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comput. J.* 7, 2 (01 1964), 155–162. <https://doi.org/10.1093/comjnl/7.2.155>
- [25] Q. Renau, C. Doerr, J. Dreo, and B. Doerr. 2020. Exploratory Landscape Analysis is Strongly Sensitive to the Sampling Strategy. In *Parallel Problem Solving from Nature - PPSN XVI - 16th International Conference, PPSN 2020, Leiden, The Netherlands, September 5-9, 2020, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12270)*. Springer, 139–153. https://doi.org/10.1007/978-3-030-58115-2_10
- [26] Q. Renau, J. Dreo, C. Doerr, and B. Doerr. 2019. Expressiveness and robustness of landscape features. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '19 (Companion)*. ACM, 2048–2051. <https://doi.org/10.1145/3319619.3326913>
- [27] Q. Renau, J. Dreo, A. Peres, Y. Semet, C. Doerr, and B. Doerr. 2022. GECCO2022 Automated Algorithm Selection for Radar Network Configuration (Version V0) [Data set]. <https://doi.org/10.5281/zenodo.5963885>
- [28] I.M. Sobol'. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *U. S. S. R. Comput. Math. and Math. Phys.* 7, 4 (Jan. 1967), 86–112. [https://doi.org/10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9)
- [29] R. Storn and K.V. Price. 1997. Differential Evolution - A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization* 11, 4 (1997), 341–359. <https://doi.org/10.1023/A:1008202821328>
- [30] S. van Rijn, H. Wang, M. van Leeuwen, and T. Bäck. 2016. Evolving the structure of Evolution Strategies. In *2016 IEEE Symposium Series on Computational Intelligence, SSCI 2016*. IEEE, 1–8. <https://doi.org/10.1109/SSCI.2016.7850138>
- [31] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [32] Hao Wang, Diederick Vermetten, Furong Ye, Carola Doerr, and Thomas Bäck. 2022. IOHanalyzer: Detailed Performance Analysis for Iterative Optimization Heuristic. *ACM Transactions on Evolutionary Learning and Optimization* (2022). To appear. Free version available at <https://arxiv.org/abs/2007.03953>.