



Universiteit
Leiden
The Netherlands

Explorative data analysis of time series based algorithm features of CMA-ES variants

Nobel, J.P. de; Wang, H.; Bäck, T.H.W.; Chicano, F.

Citation

Nobel, J. P. de, Wang, H., & Bäck, T. H. W. (2021). Explorative data analysis of time series based algorithm features of CMA-ES variants. *Gecco '21: Proceedings Of The Genetic And Evolutionary Computation Conference*, 510-518. doi:10.1145/3449639.3459399

Version: Publisher's Version

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3277010>

Note: To cite this publication please use the final published version (if applicable).

Explorative Data Analysis of Time Series based Algorithm Features of CMA-ES Variants

Jacob de Nobel
Leiden University
Leiden, Netherlands

Hao Wang
Leiden University
Leiden, Netherlands

Thomas Baeck
Leiden University
Leiden, Netherlands

ABSTRACT

In this study, we analyze behaviours of the well-known CMA-ES by extracting the time-series features on its dynamic strategy parameters. An extensive experiment was conducted on twelve CMA-ES variants and 24 test problems taken from the BBOB (Black-Box Optimization Bench-marking) testbed, where we used two different cutoff times to stop those variants. We utilized the `tsfresh` package for extracting the features and performed the feature selection procedure using the Boruta algorithm, resulting in 32 features to distinguish either CMA-ES variants or the problems. After measuring the number of predefined targets reached by those variants, we contrive to predict those measured values on each test problem using the feature. From our analysis, we saw that the features can classify the CMA-ES variants, or the function groups decently, and show a potential for predicting the performance of those variants. We conducted a hierarchical clustering analysis on the test problems and noticed a drastic change in the clustering outcome when comparing the longer cutoff time to the shorter one, indicating a huge change in search behaviour of the algorithm. In general, we found that with longer time series, the predictive power of the time series features increase.

CCS CONCEPTS

• Theory of computation → Bio-inspired optimization; • Mathematics of computing → Time series analysis; Exploratory data analysis.

1 INTRODUCTION

In the field of continuous black-box optimization, a large variety of optimization algorithms are proposed in the past [7, 8] and some prominent ones, e.g., the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [15] have been successfully applied to solve real-world problems. More recently, many interesting scientific efforts have been devoted to these algorithms to boost up their performance even further. Particularly, for the CMA-ES, many variants have been devised based on the standard algorithm [15], e.g., the so-called BIPOP-CMA-ES [11] (please see [28] for an overview). Also, it is known that their empirical performance is often hard to distinguish with a statistical meaning when aggregated over a

large number of different test problems [28]. This given, provides the need to tackle the well-known *algorithm selection* problem [24], which targets to find a proper optimization algorithm among a portfolio on some problem class [20]. As an effective approach into this direction, Explorative Landscape Analysis (ELA) [19] claims that it is vital for the algorithm selection task to obtain knowledge on optimization problems. It contrives to acquire such knowledge on the properties of test problems by designing numerical features. Rather than focusing on the problem side, in this paper, we argue that it is equally important to construct a reliable and useful quantification of algorithm behaviour. Answering the question; *how different optimization algorithms would behave on different optimization problems* since conceivably, an appropriate algorithm could be chosen more effective if we have sufficient knowledge from both the algorithm and problem side.

Here, the first step in this direction is taken by monitoring the strategy parameters of the CMA-ES variants during the optimization runs and extracting time series features from the dynamics of those parameters. Specifically, we adopted a wide set of pre-defined time series features as in the `tsfresh` library, e.g., the entropy or the absolute sum of changes of the series, and apply them on the strategy parameters of the CMA-ES, e.g., the step-size and the eigenvalues of the covariance matrix. For instance, we could compute the so-called “absolute sum of changes” feature for the step-size, which measures how much the step-size oscillates over time. In this manner, time series analysis could provide insight into how the different CMA-ES variants behave over time by quantifying the differences in their features.

Facilitated by this approach, we analyse the differences between several different CMA-ES variants on single or multiple test problems. We also delve into in the dynamics of a single variant’s strategy parameters on different benchmark problems. This also provide us knowledge about the objective function a variant is optimizing, allowing us to classify objective functions in terms of algorithm dynamics. Furthermore, this could also open the possibility of quantifying problem properties, serving as an alternative to the standard ELA methodology. We further extend our investigation by exploring the feasibility of predicting the performance of those variants using the its time series features.

As for the experimental study, we took the well-know BBOB test problems from the COCO platform [13] and conducted clustering analysis and predictive modeling for the time series features.

This paper is organized as follows; section 2 introduces the necessary preliminaries for the proposed time series feature based approach, described in section 3, including a hierarchical clustering experiment. Section 4 includes a supervised learning experiment with the computed time series features. This is followed by a regression approach in section 5, investigating the relation between

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, July 10–14, 2021, Lille, France

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

ACM ISBN 978-1-4503-8350-9/21/07...\$15.00

<https://doi.org/10.1145/3449639.3459399>

the time series features and the benchmark performance of the CMA-ES. The paper is concluded in section 6.

2 PRELIMINARIES

In this work, we focus on optimizing real-valued objective functions of the form $f: [-5, 5]^d \rightarrow \mathbb{R}$ and take the well-known Black-Box Optimization-Benchmarking (BBOB) noiseless problem set [13]. We take instance number 1 and set the search dimension to $d = 5$.

CMA-ES. In this work, we analyse the CMA-ES algorithm empirically, a nature inspired continuous black box optimizer. In the CMA-ES, candidate solutions are sampled according to a multivariate normal distribution $\mathcal{N}(\vec{m}, \sigma^2 \mathbf{C})$. Through an iterative procedure, the CMA-ES adapts the parameters of this distribution gradually using selected solutions in the current iteration as well as historical information on the trajectory of its parameters. Specifically, the *rank- μ -update* (see [12] for a detailed explanation) yields the maximum likelihood estimate of \mathbf{C} based on the selected solutions while the *rank-one-update* utilizes the so-called cumulative evolution path \vec{p}_c , which summarizes the trajectory of the algorithm and preserves the sign of consecutive steps therein. Similarly, the step size σ is controlled dynamically according to the length of the conjugate evolution path \vec{p}_σ , which accumulates consecutive steps transformed by the square root of the inverse covariance matrix $\mathbf{C}^{-\frac{1}{2}}$. Hereinafter, we shall denote by \mathbf{D} the diagonal matrix holding the eigenvalues of \mathbf{C} .

The Modular CMA-ES. The CMA-ES is a state of the art continuous black-box optimization algorithm, for which over the years several variants have been proposed in the literature [1, 2, 10, 16, 23, 25, 29]. We build upon the Modular EA (ModEA) framework [27], which contains eleven structural variants of the CMA-ES algorithm and implements those as configurable modules¹. Taken as algorithmic building blocks, these modules can be used to create 4 608 unique CMA-ES variants, which allows the user to easily construct custom variants of the CMA-ES algorithm. In [27] it was shown that this includes variants that outperform any classical CMA-ES variants from the literature.

In this study, an extended version of ModEA, specifically designed for creating CMA-ES variants is used². This Modular CMA-ES includes a number of extra options, specific to the CMA-ES. This includes an extra module for performing step size adaptation, the so-called Median Success Rule (MSR) [1] step size adaptation rule.

In this study, twelve CMA-ES variants are compared. The variants consist mostly of the all single-module³ variants available within the Modular CMA-ES, with a few exceptions. Most notably, these are modules that cause every generation to have a different number of fitness function evaluations, such as modules that employ a restart strategy (e.g. IPOP) or the module for sequential selection. These modules are left out because strategy parameters are updated only once every generation and if a variant is allowed to use more (or less) fitness function evaluations per generation than other variants, this would create a bias in the observations of

¹nine of which takes binary options and the remaining two are ternary. Please see [27] for the detail

²The extended framework can be found here: Anonymized for review purposes.

³CMA-ES configurations with only one active module.

the strategy parameters, which is why such variants are left out. The list includes:

- (1) *Standard*: The standard CMA-ES, as defined in [15].
- (2) *Active*: CMA-ES with active covariance matrix update [16] penalizes bad candidate solutions with negative weights in the covariance matrix update.
- (3) *Mirrored* sampling [3]: For each new sampled point, a mirrored pair is produced.
- (4) *Mirrored Pairwise* sampling [2, 29]: Mirrored sampling, however, only the best of the two mirrored points is selected.
- (5) *Orthogonal* sampling: The sampled candidate solutions are orthonormalized using a Gram-Schmidt procedure [29].
- (6) *Elitist*: $(\mu + \lambda)$ - selection.
- (7) *Equal Weights*: Recombination weights of $\frac{1}{\lambda}$
- (8) *MSR*: Median success rule step size adaptation [1].
- (9) *TPA*: Two-point step size adaptation [10].
- (10) *Halton* sampling: Quasi-random Halton sampling [25].
- (11) *Sobol* sampling: Quasi-random Sobol sampling [25].
- (12) *Threshold Convergence*: Threshold Convergence of mutations [23]

CMA-ES' dynamic parameters. Many internal parameters of the CMA-ES algorithms are adapted throughout the optimization process, e.g., the center of mass and the covariance matrix. Among all such parameters, we consider, in this work, the following ones for our analysis:

- $\sigma \in \mathbb{R}_{\geq 0}$: The step size.
- $f_{\text{best}} \in \mathbb{R}$: The best-so-far objective value.
- $\vec{v} := (\sqrt{\lambda_1}, \dots, \sqrt{\lambda_d})^\top$: The square root of the diagonal of \mathbf{D}
- $\vec{p}_c \in \mathbb{R}^d$: The evolution path.
- $\vec{p}_\sigma \in \mathbb{R}^d$: The conjugate evolution path.

The covariance matrix \mathbf{C} and the basis of its eigenvectors \mathbf{B} are not considered in this experiment, as they would immensely increase the number of observations stored, while the information gained in storing these variables is small in comparison. Also, the center of mass \vec{m} is also not taken into account, as we consider it less of a variable telling something about the state of the algorithm than about a location in the search space.

Time series features. Here, we consider a multivariate real-valued time series with a fixed sampling frequency [4], which represents, in this context, the dynamic strategy parameters of the CMA-ES algorithm. In each iteration of CMA-ES, we compute one data point of this time series as follows: $\forall t \in [L]$,

$$\vec{x}_t = (\sigma, f_{\text{best}}, \|\vec{v}\|, \|\vec{p}_\sigma\|, \|\vec{p}_c\|, \text{ave}(\vec{v}), \text{ave}(\vec{p}_\sigma), \text{ave}(\vec{p}_c))^\top, \quad (1)$$

where $\|\vec{v}\| = \sqrt{\vec{v}^\top \vec{v}}$ and L is the total number of iterations (running length) executed in our experiments and $\text{ave}(\cdot)$ stands for the average of the components of its argument. We compute the component-wise average and norm of each vector variable in order to store their information invariant to problem dimension. Hereafter, we shall denote this multivariate time series as $\mathbf{X} = [\vec{x}_1, \vec{x}_2, \dots, \vec{x}_L] \in \mathbb{R}^{8 \times L}$.

Given the time series data, we often probe into the characteristic thereof by means of extracting numerical features in the time domain, which give rise to meaningful interpretations and also embed the high-dimensional time series data into a lower-dimensional

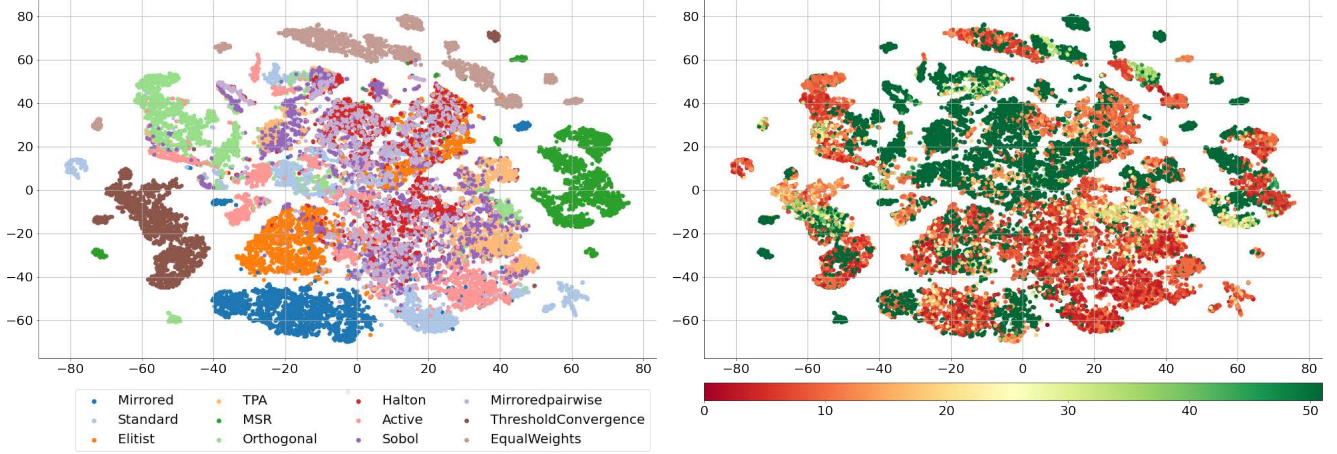


Figure 1: For the $L = 500$ case, the two-dimensional embedding (obtained from t-SNE) of the selected time series features over all test problems are color-coded according to the variant (left) and the number of predefined targets $\{10^{2-(x-1)/5} : x \in [0..51]\}$ being hit in each independent run (right).

Euclidean space [6]. Here, we adopted the widely used `tsfresh` package [5], which implements a variety of feature functions for univariate time series, such as the auto-correlation and energy functions. Given a set of m feature functions $\{\phi_i\}_{i=1}^m$ from `tsfresh` (where $\phi_i: \mathbb{R}^L \rightarrow \mathbb{R}$), we apply each feature function over each time series in \mathbf{X} (i.e., each row vector of \mathbf{X}), yielding an overarching feature map, $\Phi: \mathbb{R}^{8 \times L} \rightarrow \mathbb{R}^{8m}$:

$$\mathbf{X} \mapsto (\phi_1(\mathbf{X}_{1*}), \dots, \phi_1(\mathbf{X}_{8*}), \dots, \phi_m(\mathbf{X}_{1*}), \dots, \phi_m(\mathbf{X}_{8*}))^\top,$$

where \mathbf{X}_{i*} represents the i th row of \mathbf{X} . Hereafter, we shall denote a feature vector as $\Phi(\mathbf{X})$, or simply Φ if the input time series is clear from the context. The feature functions applied in our experiment can be found in table 2. When applied in our case, such a feature extraction procedure, results in 6352 features (see section 3.1 for detail), which ought to undergo a proper feature selection process (see section 3.2) to reduce its dimensionality, hence facilitating the modelling tasks later.

3 THE PROPOSED APPROACH

Given a benchmarking problem set $\mathcal{F} = \{f_1, f_2, \dots\}$, a set of CMA-ES variants $\mathcal{A} = \{A_1, A_2, \dots\}$, and a set of time series feature functions, we aim to distinguish those variants for each $f \in \mathcal{F}$ by computing the aforementioned time series features of every variant on f and building a classifier based on the features. Specifically, we repeat the feature computation N times independently for each A and f , since the time series features from a single run exhibits large variability. We shall denote by $\Phi(A, f, d, i)$ the time series features computed from the i th independent run of a variant A on the d -dimensional function f . Taking this notation, we could express the data set for classifying (see section 4) CMA-ES variants as $\{(\Phi(A, f, d, i), A)\}_{i \in [1..N], A \in \mathcal{A}, f \in \mathcal{F}}$, and similarly the one for discriminating test problems, $\{(\Phi(A, f, d, i), f)\}_{i \in [1..N], A \in \mathcal{A}, f \in \mathcal{F}}$. As previously mentioned, we consider 12 CMA-ES variants for \mathcal{A} , 24 BBOB problems for \mathcal{F} with $N = 100$ and $d = 5$. For all the

classification tasks presented in the sequel, we use a variant of the random forest model, called Extremely Randomized Trees Classifier (ERTC) [9], which constructs each constituting decision tree without bootstrapping based on random splitting of internal nodes (implemented by the Python library `scikit-learn` [22]). We took the default hyper-parameter setting of this algorithm since it is not the focus of this work to perform a hyperparameter tuning task and the default hyperparameter settings are empirically robust. We choose this model because it treats all features more equally compared to the random forest, since as an aggressive feature selection was applied before training this model (see Sec. 3.2).

3.1 Data generation

To generate the time series data, we conducted $N = 100$ independent runs for each pair of CMA-ES variant and test function in 5D, using the first instance (`iid = 1`) of each problem. In addition, for every run, we start all CMA-ES variants with the same initial mean, the step-size, and covariance matrix (set to the identity matrix) to remove randomness incurred by the initialization, making the time series yielded from different variants/problems more comparable. Each independent run was terminated with a fixed number of 500 generations, from which two sets of features were extracted, one using only the first 100 generations (represented by $L = 100$) of the data and the other using the complete length of 500 generations ($L = 500$). In detail, we utilized the `tsfresh` [5] python library (version 0.16) to extract the features of time series \mathbf{X} , where we took all 794 default features⁴, resulting in $794 \times 8 = 6352$ raw feature values after applied over each time series.

Due to the fact that different feature values vary largely from one feature to the other and also from functions to functions, we re-scale the feature value to the unit interval for each time series on each function. Besides, we also recorded the number of predefined targets

⁴This huge set of features are obtained from the combination of 63 default feature functions and several different parameterizations for each of them. Please see https://tsfresh.readthedocs.io/en/latest/text/list_of_features.html for the detail.

$N_s = \{10^{2-(x-1)/5} : x \in [0..51]\}$ being hit in each independent run, for the purpose of predicting the performance of each CMA-ES variants based on its time series features.

3.2 Feature selection

The vast number of raw features cast any learning task that takes them as input as a high-dimensional problem, and hence it is necessary to reduce its dimensionality by selecting more informative features. For this purpose, we picked the Boruta [18] algorithm⁵, which works in conjunction with the Random Forests Classifier (RFC) and tries to capture all the important features in the data set w.r.t. a target variable (here: the CMA-ES variant label). It creates so-called shadow features by randomly permuting the values for each feature, then trains a random forest⁶ on the combination of the original features and the shadow features, and finally determines to retain a feature if its feature importance is significantly higher than its shadow counterpart. We took the default parameter settings of Boruta except that we set its maximal number of iterations to 150 and took the “auto”⁷ option for the number of trees of the random forest model employed therein. Practically, we grouped the feature data by unique pairs of test problems and two running lengths $L \in \{100, 500\}$, which leads to 48 groups, and performed the feature selection procedure on each group separately with Boruta, where a random forest model is trained to classify the algorithm variants on each test problem based on their selected time-series features.

We only keep features that are selected by Boruta in all aforementioned 48 groups of data. If multiple parameterizations of the same feature and time series combination are selected, only the combination that shows the highest feature importance is retained. The feature importance is retrieved from the random forest model underpinning the Boruta algorithm. This selection procedure managed to reduce the dimensionality drastically from 6 352 to 32, which is summarized in table 1 in terms of the number of selected ones for each raw time series, accompanied by the total and average feature importance (FI) values from the random forest model employed in Boruta. From the table, we can observe that 1) none of the features originating from f_{best} is selected; 2) features from $\|\vec{p}_\sigma\|$ are selected the most frequently; 3) features based on σ show the highest average feature importance. Table 2 shows a specific overview of the selected feature functions and their parameter sets, as well as the time series they were applied over.

Furthermore, we visualize the selected features of all test problems in Fig. 1, which is obtained by embedding the features to a two-dimensional Euclidean space using the t-SNE [26] dimensionality reduction technique⁸. From the figure (left), we can see that many variants form distinct clusters in this embedded space, while some variants (e.g., Standard and Mirrored-Pairwise), however, are hardly distinguishable from each other. This is to be expected, since some variants are functionally more similar than other variants.

⁵It is available in Python: <https://github.com/blue-yonder/tsfresh>.

⁶Essentially, any other feature importance calculation can also be applied here, e.g., from Gradient Boosting.

⁷Please see https://github.com/scikit-learn-contrib/boruta_py for more information on the Boruta package

⁸We took the implementation from the sklearn [22] library with the perplexity parameter set to 100.

Series	#selected features	Total FI	Average FI
$\ \vec{p}_\sigma\ $	9	1.1071	0.1230
$\text{ave}(\vec{p}_\sigma)$	1	0.0398	0.0398
$\ \vec{p}_c\ $	8	0.6081	0.0760
$\text{ave}(\vec{p}_c)$	0	-	-
$\ \vec{v}\ $	5	0.5442	0.1088
$\text{ave}(\vec{v})$	6	0.8566	0.1428
σ	3	0.4893	0.1631
f_{best}	0	-	-

Table 1: Grouped by the time series, the number of selected features and their corresponding feature importance are listed with the highest value depicted in boldface.

Feature Function	Applied Over	Parameters
absolute_sum_of_changes	$\{\ \vec{p}_c\ \}$	\emptyset
approximate_entropy	$\{\ \vec{p}_\sigma\ \}$	$\{(2, .1)\}$
autocorrelation	$\{\ \vec{p}_\sigma\ \}$	$\{(1.)\}$
change_quantiles	$\{\ \vec{p}_c\ , \text{ave}(\vec{p}_\sigma), \ \vec{p}_\sigma\ , \sigma\}$	$\{(mean, F, .5, .6, .2), (mean, 1, .4, .6, 0), (mean, F, .4, .6, 0), (mean, F, .2, .6, 0)\}$
cid_ce	$\{\ \vec{p}_\sigma\ \}$	$\{(F)\}$
energy_ratio_by_chunks	$\{\ \vec{v}\ , \text{ave}(\vec{v})\}$	$\{(10, 0)\}$
fft_aggregated	$\{\text{ave}(\vec{v})\}$	$\{(centroid)\}$
fft_coefficient	$\{\ \vec{p}_c\ \}$	$\{(0, abs)\}$
index_mass_quantile	$\{\ \vec{v}\ , \text{ave}(\vec{v})\}$	$\{(1.)\}$
mean	$\{\ \vec{p}_c\ \}$	\emptyset
median	$\{\ \vec{v}\ , \text{ave}(\vec{v}), \ \vec{p}_c\ , \ \vec{p}_\sigma\ \}$	\emptyset
minimum	$\{\ \vec{v}\ , \text{ave}(\vec{v})\}$	\emptyset
number_of_crossing_m	$\{\ \vec{p}_\sigma\ \}$	$\{(1.)\}$
number_peaks	$\{\sigma\}$	$\{(1.)\}$
partial_autocorrelation	$\{\ \vec{p}_\sigma\ \}$	$\{(1.)\}$
quantile	$\{\ \vec{v}\ , \text{ave}(\vec{v}), \ \vec{p}_c\ , \ \vec{p}_\sigma\ , \sigma\}$	$\{(1.)\}$
sum_count	$\{\ \vec{p}_c\ , \ \vec{p}_\sigma\ \}$	$\{(1., -1.)\}$
sum_values	$\{\ \vec{p}_c\ \}$	\emptyset

Table 2: The selected time series features shown as the combination of the feature function, the time series on which the feature function is applied, and the parameterization therein. We express the parameterization as a set of tuples (empty for the parameterless case), where each tuple of parameter values is only applied on the time series appearing in the same order as with the tuple, e.g., for the “change_quantile” function, (mean, 0, .2, .6, 0) is only applied on σ . Please see https://tsfresh.readthedocs.io/en/latest/text/list_of_features.html for the detail.

3.3 Clustering Analysis

Here, we propose to measure the similarity between CMA-ES variants by taking the mean of corresponding feature vectors over all test problems and runs, and then computing the mean Euclidean distance among the resultant. We render the distance matrix as a heatmap in Fig. 3 for both $L = 100$ (the lower triangle) and $L = 500$ (the upper triangle), which we observe that TPA and Threshold Convergence are the most dissimilar pair of variants, which are also quite distant from the other variants. For $L = 100$, the most similar ones are the Sobol and Halton variants. Even though they still show a very high similarity for $L = 500$, the most similar pair is the Standard CMA-ES and Mirrored Sampling with pairwise selection.

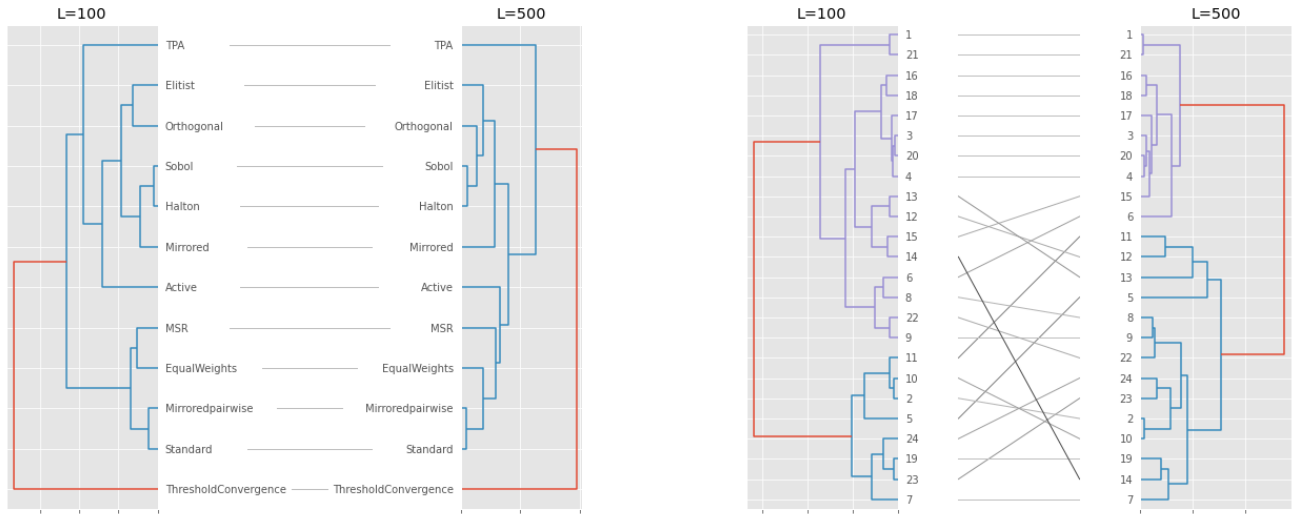


Figure 2: The hierarchical clustering results obtained on the pairwise distance of CMA-ES variants (left) and of test problems (right), in which we also contrast the $L = 100$ case with $L = 500$ to demonstrate the extent to which the clustering result would be affected by a prolonged running length. Also, we take the distance between the cluster to determine the color scale for the line connecting them. The dendrograms were produced for the mean vector of each variant/test function, using Euclidean distance and the Ward [21] hierarchical clustering.

In addition, we ran a more fine-grained hierarchical clustering on the distance matrix (shown in Fig. 2, left). The similarities between the variants invoked the following consideration: *variants which are seemingly different in terms of algorithmic structure, can exhibit similar behaviour*. This suggests that it can be helpful to probe into the dynamics of newly designed variants using time-series features, if such a variant does not perform as expected.

The distance among all test problems is only computed from features of the standard CMA-ES in line with the classification setup (see section 4). In Fig. 4 a pairwise distance matrix is showing this result. We also drawn the result of a hierarchical clustering in Fig. 2 (on the right). From those two figures, we can see that the most closely related functions seem to be the triplet of functions f_3 , f_4 , and f_{20} , which is also very close to f_{15} , f_{16} , f_{17} , and f_{18} and hence forms a bigger cluster together. Besides, problem f_1 and f_{21} are also decently similar, opening an interesting question on interpreting it since f_1 is generally considered simple while f_{21} is conceivably difficult. We hypothesize that this is because the local topology of f_{21} resembles the global topology of f_1 .

Another intriguing observation is that problem f_{13} is noticeably different from the other problems for $L = 500$ while such a difference does not exist for the $L = 100$ case. This indicates that the dynamics of the standard CMA-ES when optimizing f_{13} are similar to the other functions in the first 100 generations of the optimization procedure and it that differentiates substantially from the others in the next 400 generations. This entails that there might be huge changes to the landscape of f_{13} between generations 100 and 500, attributing to this huge change in the CMA-ES's behavior. Interestingly, this connotation matches the design rationale behind f_{13} [14]: “Approaching the ridge is initially effective, but becomes

1	.26	.30	.04	.16	.20	.15	.23	.53	.12	.07	.53	
2	.36	.28	.41	.20	.33	.33	.32	.39	.31	.29	.87	
3	.19	.45	.24	.11	.19	.44	.31	.37	.22	.09	.69	
4	.05	.29	.31	.32	.22	.16	.22	.58	.14	.16	.48	
5	.27	.21	.26	.20	.14	.26	.27	.33	.12	.16	.79	
6	.30	.32	.14	.35	.15	.38	.31	.29	.21	.21	.77	
7	.11	.21	.26	.12	.31	.31	.13	.56	.19	.18	.56	
8	.14	.35	.20	.16	.26	.28	.28	.47	.20	.17	.60	
9	.53	.49	.45	.55	.42	.39	.58	.49	.47	.48	1.0	
10	.11	.27	.09	.16	.17	.17	.23	.17	.45	.02	.62	
11	.11	.24	.25	.10	.10	.16	.20	.22	.45	.05	.63	
12	.64	.84	.75	.60	.77	.73	.68	.69	1.0	.70	.69	
	1	2	3	4	5	6	7	8	9	10	11	12

Figure 3: Pairwise distance matrix for all CMA-ES variants. The numeric index used here refers to the number specified in the list in section 2. Euclidean distance is computed and scaled though min-max normalization. The lower triangle of the figure shows the pairwise distances for $L = 100$, the upper triangle shows $L = 500$.

ineffective close to the ridge... The necessary change in ‘search behavior’ close to the ridge...”, hinting that a solver needs to change its search behaviour as it approaches the ridge defined in the function landscape.

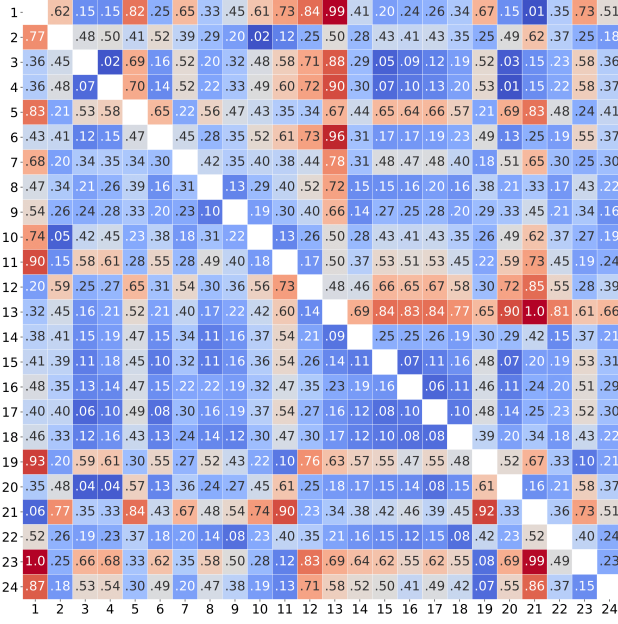


Figure 4: Pairwise distance matrices for all functions using data from the Standard CMA-ES. Euclidean distance is computed and scaled though min-max normalization. The lower triangle of the figure shows the pairwise distances for $L = 100$, the upper triangle shows $L = 500$.

4 CLASSIFYING CMA-ES VARIANTS AND PROBLEMS

Taking the 32 selected features, we constructed Extremely Randomized Trees Classifiers (ERTCs) to discriminate CMA-ES variances and the test problems for both cases when $L = 100$ and $L = 500$. This allows for investigating the impact of the length of raw time series, in terms of the classification power. Also, the resulting models are assessed using cross validation (CV) strategy (please see below for the detail). In addition to model accuracy we measure F1-score. Specifically, we calculate, for each fold, the precision (p) and recall r with respect to each label and take the harmonic mean, i.e. $2 \frac{p \times r}{p+r}$.

Classifying CMA-ES variants. Here, we consider a multi-class classification task to distinguish CMA-ES variants on

- each test problem separately, in which we have $|\mathcal{A}| \times N = 1200$ data points in each case, and
- on the whole data set from all 24 test problems.

For the former scenario, we perform a K-fold CV (with $K = 24$) and summarize the resulting classification accuracy in table 3, in which we observe an average accuracy of $.803 \pm .087$ for $L = 100$ and $.937 \pm .069$ for $L = 500$, indicating that a longer time series improves the classification performance. In detail, the lowest accuracy is observed on problem f_{23} for both $L = 100$ and $L = 500$ while the highest one is achieved on f_{21} for $L = 100$ and f_6 for $L = 500$.

FID	Acc. $L = 100$	Acc. $L = 500$	F1 $L = 100$	F1 $L = 500$
f_1	.939 ± .032	.991 ± .014	.939 ± .034	.991 ± .014
f_2	.742 ± .065	.980 ± .022	.741 ± .071	.980 ± .022
f_3	.847 ± .057	.983 ± .018	.843 ± .059	.984 ± .018
f_4	.821 ± .050	.983 ± .018	.819 ± .050	.983 ± .018
f_5	.793 ± .064	.822 ± .070	.789 ± .072	.825 ± .075
f_6	.844 ± .053	.995 ± .009	.842 ± .054	.995 ± .009
f_7	.829 ± .050	.878 ± .054	.829 ± .055	.877 ± .058
f_8	.771 ± .062	.971 ± .024	.764 ± .062	.971 ± .024
f_9	.788 ± .066	.982 ± .022	.782 ± .067	.982 ± .024
f_{10}	.771 ± .042	.986 ± .015	.765 ± .046	.986 ± .015
f_{11}	.796 ± .076	.992 ± .013	.795 ± .080	.992 ± .013
f_{12}	.877 ± .039	.945 ± .020	.877 ± .040	.946 ± .019
f_{13}	.828 ± .051	.881 ± .043	.826 ± .058	.882 ± .041
f_{14}	.870 ± .044	.935 ± .037	.868 ± .043	.936 ± .036
f_{15}	.809 ± .051	.975 ± .027	.812 ± .053	.976 ± .027
f_{16}	.800 ± .051	.944 ± .026	.804 ± .052	.944 ± .027
f_{17}	.788 ± .057	.878 ± .047	.788 ± .056	.877 ± .048
f_{18}	.748 ± .057	.850 ± .056	.748 ± .061	.846 ± .055
f_{19}	.680 ± .052	.845 ± .052	.681 ± .057	.842 ± .051
f_{20}	.834 ± .049	.978 ± .016	.834 ± .051	.979 ± .016
f_{21}	.958 ± .033	.992 ± .021	.958 ± .033	.993 ± .021
f_{22}	.767 ± .049	.939 ± .034	.767 ± .052	.941 ± .033
f_{23}	<u>.677 ± .072</u>	<u>.810 ± .063</u>	<u>.670 ± .077</u>	<u>.807 ± .066</u>
f_{24}	.705 ± .066	.952 ± .035	.704 ± .070	.953 ± .037
Avg.	.803 ± .067	.937 ± .058	.802 ± .068	.937 ± .059

Table 3: For classifying CMA-ES variants on each test problem separately, we list the accuracy and F1 scores resulted from a 24-fold CV procedure with the highest measured values marked in boldface and the lowest ones underlined.

For the latter scenario, we merged data points from all 24 test problems and conducted two different CV strategies: a standard randomized K-fold CV with $K = 24$ and a *Leave-One-Group-Out* (LOGO) CV, where we hold data points from a single problem for validation and train the model on the remaining 23 problems. Note that the LOGO procedure iterates over each test problem. We show the resulting accuracy and F1 score in table 4, in which we see that compared to both performance values in the K-fold case, the ones from LOGO CV drop slightly. This matches our expectation that the model would deteriorate if we exclude entirely the information of one test problem during the training and manage to generalize to it. We also notice that the variability of both performance values becomes much larger in the LOGO case, which is caused by substantial performance differences between each test problems. We support this statement with accuracy values in Fig. 5, where the model poorly performs on f_5 , f_6 , and f_7 while its accuracy exceeds 0.95 on problems f_1 - f_4 , f_{15} , f_{21} and f_{22} . Moreover, it is interesting that $L = 500$ does not always outperform $L = 100$ in terms of classification accuracy, e.g., on f_{12} , f_{13} , and f_{14} , $L = 100$ it is noticeably higher.

Classifying test problems. Furthermore, we delve into the question that whether the time series features are capable of discriminating test problems, which, if so, could serve as an alternative to the well-known Explorative Landscape Analysis (ELA) features [17] for identifying problem classes. For this purpose, we took only the

CV	Acc. $L = 100$	Acc. $L = 500$	F1 $L = 100$	F1 $L = 500$
K-Fold	.769 ± .013	.923 ± .005	.766 ± .014	.923 ± .005
LOGO	.691 ± .114	.833 ± .148	.682 ± .115	.822 ± .120

Table 4: Comparison of two classification approaches for CMA-ES variants: on each test problem individually (see table 3) and on all problems jointly (see Fig. 5).

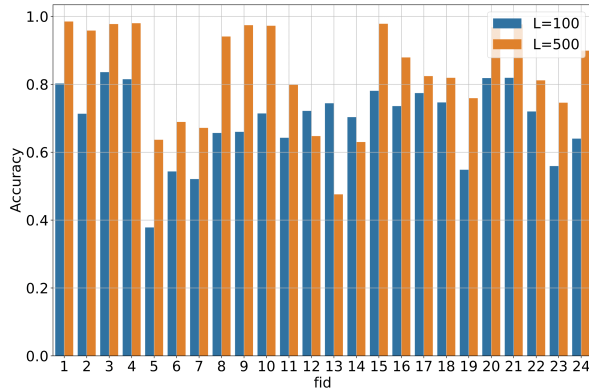


Figure 5: On classifying the CMA-ES variants over all test problems, Leave-One-Group-Out (LOGO) results show the testing accuracy on each single test problem, from a ERTC model trained on the remaining 23 problems. The $L = 500$ case is consistently better than $L = 100$ while those two cases exhibit quite similar trends over test problems.

time series feature data from the *Standard* CMA-ES to distinguish test problems using the ERTC method⁹. As with the previous task, we perform classification for both cases when $L = 100$ and $L = 500$, and use 24-fold cross validation.

This experiment, however, only achieves an accuracy of $.497 \pm .053$ for the $L = 100$ while it manages to reach $.716 \pm .048$ for $L = 500$. In contrast to the study on classifying algorithms, we encounter a substantial performance drop. In figure 6, a box plot is shown for this experiment, which shows the F1-score per classification label obtained for each cross validation fold. Again, we observe a large variance between the labels. Interestingly, we observe a pattern opposite of that of figure 5. Now, f_5 is one of the easiest functions to predict, while when using data from f_5 for testing in the LOGO experiment, rather poor classification performance was observed. This is in line with our earlier observation from the LOGO experiment; when a classifier is not able to learn function specific patterns in the data, it shows lower performance in classifying variants. The fact that f_5 is classified correctly more often than most other functions in this experiment, means that there is a pattern in the data of f_5 which distinguishes it from the data of other functions. This then also means, that when using data of f_5 for testing in the LOGO experiment, the classifier has a harder time giving a correct prediction, because it is now working on data

⁹This choice is made based on our observation that all CMA-ES variants give rise to very similar classification results. Hence, it is more natural to take the baseline variant.

Groups	Acc. $L = 100$	Acc. $L = 500$	F1 $L = 100$	F1 $L = 500$
BBOB	.643 ± .054	.758 ± .050	.640 ± .058	.750 ± .053
Mersmann	.774 ± .058	.840 ± .114	.756 ± .060	.820 ± .124

Table 5: Results on classifying function groups with a 24-fold CV procedure for both $L = 100$ and $L = 500$ cases, in which the feature extracted from the standard CMA-ES is used.

which is very dissimilar from the data it was trained on. The reverse can also be observed in f_3 . Because there are other functions in the data which express similar patterns, such as f_4 , the classifier is not able to clearly distinguish between them. However, when predicting variants (figure 5), the classifier is able to learn these function specific patterns during training from f_4 , which in turn allows it to more confidently predict variants when using data taken from f_3 for testing.

In addition to classify each test problem, we perform a classification task on problem groups, where the problems are partitioned according to the default groups in COCO/BBOB, or to the clustering results in [20] (referred as Mersman’s grouping hereafter), which is obtained on the expected running time (ERT) values of 51 different optimization algorithms on 5D. The latter grouping might be particularly suitable for this study, since it is obtained based on the consensus rankings of different optimization algorithms, which reflects, to some extent, the algorithmic behaviour and hence could be well related to the time series features extracted from dynamics of the algorithm. This conjecture is empirically supported by table 5, where the classification outcome of Mersman’s grouping is much better than that of the default BBOB grouping for $L = 100$ and $L = 500$.

5 PREDICTING CMA-ES’S PERFORMANCE

From figure 1 (right), we see clearly that feature vectors that are close to each other in terms of its two-dimensional embedding appear to be associated with similar numbers of targets reaches, which motivates us to predict the performance of each run using the corresponding features. For this goal, we adopted an Extremely Randomized Trees Regressor (ERTR, from scikit-learn [22]) with its default parameter setting. The independent variables in this regression experiment consist of the 32 selected time series features, which are taken as input to predict the number of targets N_s . Also, we took the coefficient of determination $R^2(y, \hat{y}) = 1 - \frac{\sum (y - \hat{y})^2}{\sum (y - \text{mean}(y))^2}$ as the performance metric and applied a 24-fold CV procedure as before, for both $L = 100$ and $L = 500$ cases.

When training the ERTR model on all variants (on all problems), we observed an R^2 of $.557 \pm .018$ for $L = 100$, and $.778 \pm .014$ for $L = 500$. Alternatively, when training the model for each variants separately (see Fig. 7), we see the a huge improvement of the R^2 on the $L = 500$ case for most variants, which is also considerably stable in contrast to the $L = 100$ case. It is worth noting that we do not consider a scenario in which we took each problem separately, since the distribution of N_s value varies drastically from function to function. This is because the CMA-ES shows a different performance on each benchmark function, signifying that data from

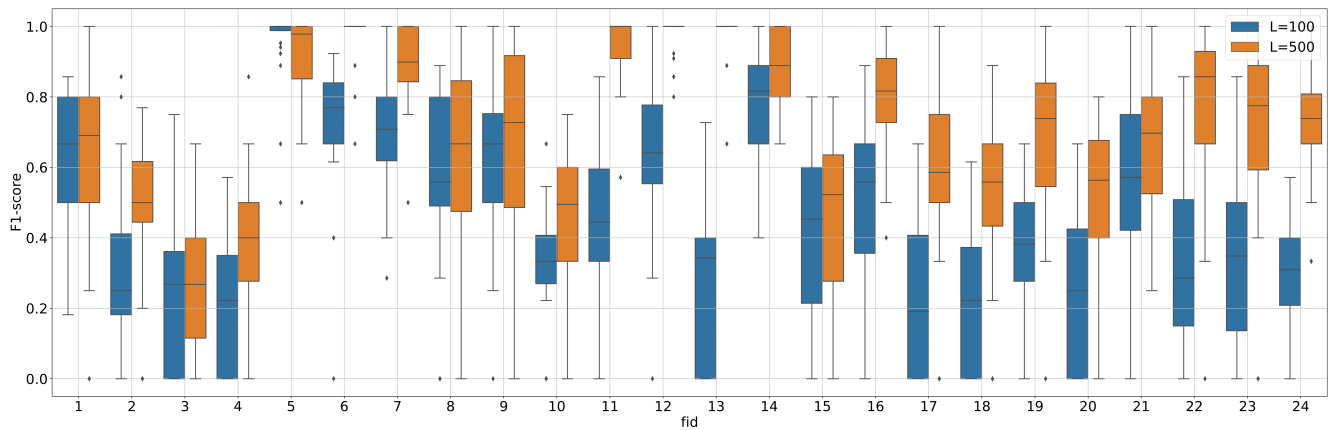


Figure 6: On classifying the test problems, we illustrate the resulting F1-score per each classification label, where the box plot indicates the variability of this score across a 24-fold cross validation. We also contrast the $L = 100$ case with the $L = 500$ one.

a given function can not be used to train a regression model for another function, and vice versa.

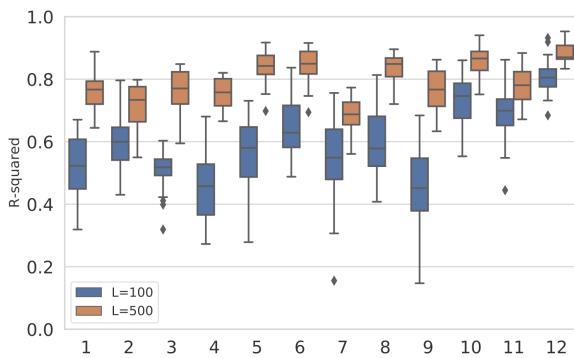


Figure 7: Regression, with a model trained on each CMA-ES variant separately, showing a box plot of the R^2 attained using K-fold cross validation, with $K=24$. The numeric index used here on the x-axis refers to the number specified in the list in section 2. Data for $L = 100$ is shown in blue, $L = 500$ is shown in orange.

6 CONCLUSIONS

In this work, we have used time series characterization methods to extract features from the dynamic strategy parameters of the CMA-ES variants during the optimization of 24 test problems from the Black-Box Optimization Benchmark (BBOB). After extracting the time series features on 5D with two different running length of each variant ($L = 100$ and $L = 500$ for 100 and 500 iterations, respectively), we conducted a proper feature selection to reduce the dimensionality of the resulting features. We built extremely randomized trees to classify the variants on a single and multiple

problems, as well as to distinguish those test problems. Furthermore, we performed in-depth explorative data analysis on the selected features, where we show the hierarchical cluster discovered in CMA-ES variants and in test problems and contrast such clustering results obtained in both $L = 100$ and $L = 500$ cases. Moreover, we explored the potential of directly predicting the performance of CMA-ES variants using their corresponding features.

We observed that the time series feature are capable of discriminating CMA-ES variants with nearly perfect accuracy, if we train the model on a single problem only. We also managed to achieve a decent accuracy when classifying the function groups. When it comes to predict the performance of those variants, we saw an indication that the feature might be reasonably related to the performance, which is subject to future researches.

From the clustering analysis of test problems, we observed an intriguing pattern on problem f_{13} , which resembles many other problems for the $L = 100$ case and becomes very distant from all other functions for $L = 500$. This phenomenon suggests that there might be a drastic change in the landscape which causes the change of the behavior of the CMA-ES, which is reflected by the feature. This is also supported by the fact that the landscape of f_{13} is very heterogeneous.

For the future research, we are planning to 1) apply this approach on more problems and algorithms; 2) to investigate the usefulness of the time series feature for the dynamic algorithm selection task; 3) to extend it by computing the features in a moving window; and 4) to dive more deeply into the data to investigate the relative contribution and values of each feature with the help of, for instance the Shapley value.

REFERENCES

- [1] Ouassim Ait Elhara, Anne Auger, and Nikolaus Hansen. 2013. A Median Success Rule for Non-Elitist Evolution Strategies: Study of Feasibility. *GECCO 2013 - Proceedings of the 2013 Genetic and Evolutionary Computation Conference* (07 2013). <https://doi.org/10.1145/2463372.2463429>
- [2] Anne Auger, Dimo Brockhoff, and Nikolaus Hansen. 2011. Mirrored sampling in evolution strategies with weighted recombination. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 861–868.

- [3] Dimo Brockhoff, Anne Auger, Nikolaus Hansen, Dirk Arnold, and Tim Hohm. 2010. Mirrored Sampling and Sequential Selection for Evolution Strategies. *PPSN 2010, Proceedings*. https://doi.org/10.1007/978-3-642-15844-5_2
- [4] Chris Chatfield. 2004. *The analysis of time series: an introduction* (6th ed.). CRC Press, Florida, US.
- [5] Maximilian Christ, Nils Braun, Julius Neuffer, and Andreas W. Kempa-Liehr. 2018. Time Series Feature Extraction on basis of Scalable Hypothesis tests (tsfresh – A Python package). *Neurocomputing* 307 (2018), 72 – 77. <https://doi.org/10.1016/j.neucom.2018.03.067>
- [6] Maximilian Christ, Andreas W. Kempa-Liehr, and Michael Feindt. 2016. Distributed and parallel time series feature extraction for industrial big data applications. *CoRR* abs/1610.07717 (2016). arXiv:1610.07717 <http://arxiv.org/abs/1610.07717>
- [7] Sandor Csizsar. 2007. Optimization Algorithms (survey and analysis). *LINDI 2007 - International Symposium on Logistics and Industrial Informatics 2007, Proceedings*, 185–188. <https://doi.org/10.1109/LINDI.2007.4343536>
- [8] Ginnu George and Kumudha Raimond. 2013. A Survey on Optimization Algorithms for Optimizing the Numerical Functions. *International Journal of Computer Applications* 61 (01 2013), 41–46. <https://doi.org/10.5120/9935-4570>
- [9] Pierre Geurts, Damien Ernst, and Louis Wehenkel. 2006. Extremely Randomized Trees. *Machine Learning* 63 (04 2006), 3–42. <https://doi.org/10.1007/s10994-006-6226-1>
- [10] Nikolaus Hansen. 2008. CMA-ES with Two-Point Step-Size Adaptation. *CoRR* abs/0805.0231 (2008). arXiv:0805.0231 <http://arxiv.org/abs/0805.0231>
- [11] Nikolaus Hansen. 2009. Benchmarking a BI-Population CMA-ES on the BBOB-2009 Function Testbed. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers* (Montreal, Québec, Canada) (GECCO '09). Association for Computing Machinery, New York, NY, USA, 2389–2396. <https://doi.org/10.1145/1570256.1570333>
- [12] Nikolaus Hansen. 2016. The CMA Evolution Strategy: A Tutorial. *CoRR* abs/1604.00772 (2016). arXiv:1604.00772 <http://arxiv.org/abs/1604.00772>
- [13] Nikolaus Hansen, Dimo Brockhoff, Olaf Mersmann, Tea Tusar, Dejan Tusar, Ouassim Ait ElHara, Phillipe R. Sampaio, Asma Atamna, Konstantinos Varelas, Umut Batu, Duc Manh Nguyen, Filip Matzner, and Anne Auger. 2019. *COMparing Continuous Optimizers: numbb/COCO on Github*. <https://doi.org/10.5281/zenodo.2594848>
- [14] Nikolaus Hansen, Steffen Finck, Raymond Ros, and Anne Auger. 2009. *Real-Parameter Black-Box Optimization Benchmarking 2009: Noiseless Functions Definitions*. Research Report RR-6829. INRIA. <https://hal.inria.fr/inria-00362633>
- [15] Nikolaus Hansen and Andreas Ostermeier. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evol. Comput.* 9, 2 (June 2001), 159–195. <https://doi.org/10.1162/106365601750190398>
- [16] Grahame A Jastrebski and Dirk V Arnold. 2006. Improving evolution strategies through active covariance matrix adaptation. In *Evolutionary Computation, CEC 2006. IEEE Congress on. IEEE*, 2814–2821.
- [17] Pascal Kerschke and Heike Trautmann. 2019. Comprehensive feature-based landscape analysis of continuous and constrained optimization problems using the R-package flacco. In *Applications in Statistical Computing*. Springer, 93–123.
- [18] Miron B. Kursa and Witold R. Rudnicki. 2010. Feature Selection with the Boruta Package. *Journal of Statistical Software* 36, 11 (2010), 1–13.
- [19] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. 2011. Exploratory Landscape Analysis. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation* (Dublin, Ireland) (GECCO '11). Association for Computing Machinery, New York, NY, USA, 829–836. <https://doi.org/10.1145/2001576.2001690>
- [20] O. Mersmann, M. Preuss, H. Trautmann, B. Bischl, and C. Weihs. 2015. Analyzing the Bbob Results by Means of Benchmarking Concepts. *Evol. Comput.* 23, 1 (March 2015), 161–185. https://doi.org/10.1162/EVCO_a_00134
- [21] Fionn Murtagh and Pierre Legendre. 2011. Ward's Hierarchical Clustering Method: Clustering Criterion and Agglomerative Algorithm. *CoRR* abs/1111.6285 (2011). <http://dblp.uni-trier.de/db/journals/corr/corr1111.html#abs-1111-6285>
- [22] 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.
- [23] A. Piad-Morffis, S. Estévez-Velarde, A. Bolufé-Röhler, J. Montgomery, and S. Chen. 2015. Evolution strategies with threshold convergence. In *2015 IEEE Congress on Evolutionary Computation (CEC)*. 2097–2104. <https://doi.org/10.1109/CEC.2015.7257143>
- [24] John R. Rice. 1976. The Algorithm Selection Problem. *Advances in Computers* 15 (1976), 65–118. <http://dblp.uni-trier.de/db/journals/ac/ac15.html#Rice76>
- [25] Olivier Teytaud. 2015. Quasi-Random Numbers Improve the CMA-ES on the BBOB Testbed. In *Revised Selected Papers of the 12th International Conference on Artificial Evolution - Volume 9554*. Springer-Verlag, Berlin, Heidelberg, 58–70. https://doi.org/10.1007/978-3-319-31471-6_5
- [26] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9, 86 (2008), 2579–2605. <http://jmlr.org/papers/v9/vandermaaten08a.html>
- [27] 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)*. 1–8. <https://doi.org/10.1109/SSCI.2016.7850138>
- [28] Sander van Rijn, Hao Wang, Bas van Stein, and Thomas Bäck. 2017. Algorithm Configuration Data Mining for CMA Evolution Strategies. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Berlin, Germany) (GECCO '17). ACM, New York, NY, USA, 737–744. <https://doi.org/10.1145/3071178.3071205>
- [29] Hao Wang, Michael Emmerich, and Thomas Bäck. 2014. Mirrored Orthogonal Sampling with Pairwise Selection in Evolution Strategies. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing* (Gyeongju, Republic of Korea) (SAC '14). Association for Computing Machinery, New York, NY, USA, 154–156. <https://doi.org/10.1145/2554850.2555089>