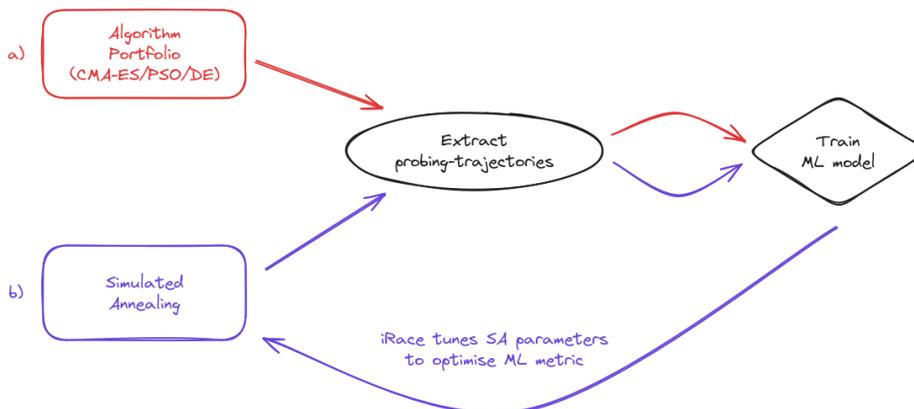




# Improving Algorithm-Selection and Performance-Prediction via Learning Discriminating Training Samples

Quentin Renau  
Edinburgh Napier University  
Edinburgh, Scotland, UK  
q.renau@napier.ac.uk

Emma Hart  
Edinburgh Napier University  
Edinburgh, Scotland, UK  
e.hart@napier.ac.uk



**Figure 1: Two approaches to training an algorithm selector are compared. In (a), short trajectories are extracted by running solvers from the portfolio of interest which are then used to train a classifier to predict the best solver. In (b), trajectories are produced from running a simulated annealing algorithm whose parameters are tuned via *irace* [25] to optimise a machine-learning metric.**

## ABSTRACT

The choice of input-data used to train algorithm-selection models is recognised as being a critical part of the model success. Recently, feature-free methods for algorithm-selection that use short trajectories obtained from running a solver as input have shown promise. However, it is unclear to what extent these trajectories reliably discriminate between solvers. We propose a *meta* approach to generating discriminatory trajectories with respect to a portfolio of solvers. The algorithm-configuration tool *irace* is used to tune the parameters of a simple Simulated Annealing algorithm (SA) to produce trajectories that maximise the performance metrics of ML models trained on this data. We show that when the trajectories obtained from the tuned SA algorithm are used in ML models for algorithm-selection and performance prediction, we obtain significantly improved performance metrics compared to models trained both on raw trajectory data and on exploratory landscape features.

## KEYWORDS

Algorithm Selection, Performance Prediction, Black-Box Optimisation, Algorithm Trajectory

### ACM Reference Format:

Quentin Renau and Emma Hart. 2024. Improving Algorithm-Selection and Performance-Prediction via Learning Discriminating Training Samples. In *Genetic and Evolutionary Computation Conference (GECCO '24)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3638529.3654025>

## 1 INTRODUCTION

*Algorithm-selection* – the process of selecting the best algorithm to solve a given problem instance – is motivated by the fact that algorithms in a portfolio deliver complementary performance on diverse problem instances and was first recognised by Rice [40] in 1976. Algorithm-selectors usually take the form of a model that given an input describing an instance, either predict the label of the best solver (classification) or its performance (regression) [46]. Typically models are trained using a feature-vector that describes the instance. In continuous domains, Exploratory Landscape Analysis [27] is often used to create features, whereas in combinatorial domains, it is more typical to use hand-designed features that are specific to the domain.

However, feature-based selection poses a number of issues. Firstly, in some domains, specifying features is not intuitive, and it can be difficult to create a sufficient number to train a model. Conversely, in other domains, feature-selection methods are required



This work is licensed under a Creative Commons Attribution International 4.0 License. *GECCO '24*, July 14–18, 2024, Melbourne, VIC, Australia  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0494-9/24/07.  
<https://doi.org/10.1145/3638529.3654025>

to select the most informative ones [18, 26, 44]. Secondly, instance features often do not correlate well with algorithm performance data. For instance, Sim *et. al.* [43] demonstrate in the TSP domain that instances that are close in the feature-space can be very distant in the performance space (i.e., the Euclidean distance between their feature-vectors is small while the distance between the performance of two algorithms on the instances is very large). Thirdly, while features can capture nuances of a fitness landscape [27], they are usually defined *independently* of any solver [15], which appears problematic if the goal is to inform solver prediction.

To address this, some recent work proposes to use information extracted from fitness trajectories obtained by running a solver for a short period of time to train selectors [15, 22, 39]: these methods directly capture something about the search process of a solver on an instance. Time-series features can be extracted from the trajectories and used to train a selector [8, 9, 39] or alternatively, the trajectory can be used directly [39], avoiding the need for feature-calculation. However, for trajectory-based data to be useful to train a model, it must be *discriminatory* – despite the success of some of the approaches just mentioned above, it is clear that many meta-heuristic solvers might produce similar trajectories on an instance, particularly given their stochastic nature.

In this article, we propose a method that harnesses the benefits that trajectory-based approaches can bring (i.e. that they reflect algorithm behaviour) but also addresses the weaknesses just raised. Our method is inspired by research in other fields where *meta-algorithms* have been used to learn a transformation that enables an existing algorithm or model to improve its performance. For example, in the domain of classification, [24] use genetic programming to learn a low-dimensional embedding of ML datasets such that when a classifier is applied in the learned space, classification metrics are improved. In robotics, [5, 6] describe a meta-approach to solving control problems in which an EA learns a definition of an archive such that when a MAP-Elites algorithm [29] is applied in the defined archive, objective performance is improved. Also in robotics, Meyerson *et. al.* [28] propose a method for meta-learning behaviour descriptors which are then used by novelty-search algorithm [23] to solve a maze-navigation task. As described in Figure 1, our proposed method uses the well-known algorithm-configuration tool irace [25] to tune the parameters of a simple Simulated Annealing (SA) algorithm [21] such that it produces *short* trajectories that when used to train an algorithm-selector or regressor, optimise the model metrics (e.g. regression or classification). The goal of tuning the SA algorithm is therefore to produce short, discriminatory trajectories that can be cheaply calculated. Following training, when new test instances arrive, the tuned SA algorithm is run to create an input trajectory to a model which is trained to either predict performance or the label of the best solver from a portfolio of interest.

We evaluate the proposed approach on the well-known BBOB test-suite [11] to learn models for both classification and regression. Models are trained either using the trajectory from the tuned SA algorithm directly or by extracting time-series features from the tuned trajectory. As a baseline, we compare the results of the meta-algorithm to models that directly use ELA features and to previous results that use trajectory data obtained by running all solvers in a portfolio. We show the meta-algorithm outperforms

ELA-based models in all experiments; obtains better RMSE than single trajectory-based methods relying on the solver portfolio, and provides similar or better performance to a multi-trajectory input method but at a considerable less computational cost.

To the best of our knowledge, this is the first time such an approach has been used to improve algorithm-selectors. It addresses several weaknesses in current practice, such as recognising the need to use training data which reflects the behaviour of a solver (i.e. is ‘algorithm-centric’; the need to have a cheap method of computing input to a model; ensuring the inputs are discriminatory and finally, that scales with the size of the portfolio of solvers (unlike the approach described in [39] which requires trajectories to be computed for every solver in the portfolio to create input to a model).

*Reproducibility:* code, data, and additional plots are available at [38].

## 2 RELATED WORK

We focus on related work in algorithm-selection in the continuous domain given this is the subject of our study. A more general overview of algorithm-selection (which includes models for both classification and performance prediction) can be found in [17].

In continuous optimisation, the vast majority of previous work focuses on the use of *landscape-features* to train selectors. Exploratory Landscape Analysis (ELA) [27] is a popular method for extracting features, which has grown over the years with a gradual introduction of new features [10, 19, 30]. The reader is referred to [34] for formal definitions of the most used features and their properties. The use of ELA features has been shown to be successful in both algorithm configuration [3] and algorithm selection [17] on benchmark data as well as real-world optimisation problems [37]. The main drawback of ELA however is the overhead cost induced by the feature computation: typically this involves sampling a large number of points in order to compute features then discarding the samples, which can be wasteful. In combinatorial optimisation, feature-free approaches to algorithm-selection are beginning to emerge [1, 2, 42, 50] that use deep-learning (DL) models but have the disadvantage that deep-learning models can be very difficult to explain.

All of the approaches just described take an instance-centric view: that is, the input to a selector is independent of the execution of any algorithm. As already noted in 1, this can be problematic, given that previous work has suggested that there is not necessarily a strong correlation between the distance of two instances in a feature-space and the distance in the performance space according to a chosen portfolio of solvers [43].

Some recent work has begun to address this, using information derived from running a solver as input to a selector. For example, recent work from Jankovic *et. al.* [14, 15] proposes extracting ELA features from the search trajectory of an algorithm. In [14], trajectories are obtained using half of the available budget and combined with the state variables of the algorithm to train a performance predictor. However, their approach was outperformed by classical ELA features computed on the full search space. In [15], ELA features were computed from short trajectories of 30d points: this information was then used to train a model that predicted an algorithm to

warm-start a search process. Cenikj *et. al.* [8] log algorithm trajectories using a large budget which is used to construct time-series of statistics derived from both the population fitness to train a classifier that predicts which of the 24 BBOB functions the trajectory belongs to. Their approach successfully outperforms a model which is trained on ELA features extracted from algorithm trajectories rather than on statistical information. Kostovska *et. al.* [22] also use time-series features to train a model: the features are extracted from state variables of CMA-ES to perform a per-run algorithm selection with warm-starting, which demonstrates similar performance to ELA features on the per-run algorithm selection task. Most recently, in [39], probing-trajectories capturing the best-so-far or best-per-generation fitness over a small number of generations are used directly to train a selector without extracting any features. Despite the promise of these trajectory approaches, there remain some issues. Firstly, there can be wide variance in trajectories obtained from running the same solver on an instance multiple times, hence randomly sampled trajectories might not discriminate between solvers. Secondly, in the approach described in [39], models are trained using a concatenated list of short trajectories obtained from each solver in the portfolio as input. Clearly, this does not scale with increasing sizes of portfolios, where the cost of computing even short trajectories over multiple solvers could be prohibitive.

Although we are unaware of other work that attempts to improve discrimination between model inputs depending on solvers, it is worth mentioning research from other domains which uses a meta-algorithm approach in which a process is nested inside an outer learning loop that learns something that improves the performance of the inner process. As noted in Section 1, these approaches are common both in Evolutionary Robotics [5, 6, 28] and particularly in the Quality-Diversity literature [33]. The latter requires a space to be defined in which an algorithm searches for diverse but high-quality solutions. A meta-algorithm searches for the most appropriate space in which to run the optimiser of interest. Our work builds on this idea by using a meta-algorithm to tune a simulated annealing algorithm to produce trajectories which are discriminatory with respect to solvers in a portfolio of interest and therefore enable better algorithm-selection models to be trained.

### 3 METHODS

In this section, we describe the methods used to generate data, the portfolio of solvers of interest and the machine-learning models used.

#### 3.1 Searching for Discriminatory Trajectories

As shown in Figure 1, the meta-approach we propose tunes a Simulated Annealing [21] algorithm to produce discriminating trajectories. We chose SA both for its simplicity and for the ability to easily tune its hyper-parameters. To minimise the computational effort required to compute trajectories which are used as input to a model but are not used to solve an instance, we minimise the length of the trajectories optimised by SA. We use the dual annealing algorithm from `scipy` [49]<sup>1</sup> to provide the trajectories. The search over the configurations of the dual annealing algorithm is performed

using the default `irace` [25] package<sup>2</sup>, setting the maximum number of experiments to 5,000. The parameters involved in creating differentiating trajectories and their possible values are:

- The *length of the trajectory*  $n \in [5, 100]$ . The length is restricted to 100 to provide short trajectories with a low budget of function evaluations.
- The SA *initial temperature*  $T \in [0.02, 5e4]$ .
- The SA *visiting distribution parameter*  $v \in [1.5, 2.5]$ .
- The SA *acceptance parameter*  $a \in [-1.1e4, -5]$ .

We extract two kinds of trajectories: the ‘best’ trajectory records the best fitness seen so far on an instance during a run of the SA algorithm, while the ‘current’ trajectory records the current fitness value (given that SA can accept worse solutions). The fitness function used in `irace` is the classification accuracy of a model trained on a set of trajectories for a model that predicts the label of the best solver, and the RMSE in a performance prediction task. The method for determining the labels/performance data is described in the next section.

#### 3.2 Data

We consider the first 5 instances of the 24 noiseless Black-Box Optimisation Benchmark (BBOB) functions in dimension  $d = 10$  from the COCO platform [12] as a test-bed. For each instance, we collect data from running three algorithms: CMA-ES [13], Particle Swarm Optimisation (PSO) [16], and Differential Evolution (DE) [45]. Each algorithm is run 5 times per instance. Following the method described in [39], we record the best median target value obtained from running CMA-ES, PSO and DE after 100,000 function evaluations. Note that although this is a very large number of evaluations, its only purpose is to provide a ‘true’ label or performance value for each algorithm to train and evaluate models. The data obtained is used to label each instance with the best-performing algorithm to train a classification model to predict the best solver and as the performance value for a prediction task.

We compare the results of our meta-trajectories approach to training models that use (1) probing-trajectories directly to train models, (2) time-series features extracted from the trajectories and (3) state-of-the-art ELA feature-vectors as input. We obtain trajectory data directly from [47] which records search-trajectories per run. Note that some automated algorithm configuration was performed by [47] before they collected this data, see [48] for further details. ELA feature-data for the BBOB suite is extracted from [35]. For each feature, 100 independent values are available per function instance which were sampled using the Sobol’ low-discrepancy sequence. We use data from the 10-dimensional functions with a feature computation budget of  $30 \times d$  and the general recommendation for feature computation  $50 \times d$  [20]. We select 10 cheap features based on their expressiveness and invariance to transformation, described in [36].

### 4 MACHINE LEARNING TASKS

We use the trajectories obtained from the tuned SA algorithm to train two types of ML models: (a) a classifier that predicts the best

<sup>1</sup>version 1.10.1

<sup>2</sup>version 3.5

**Table 1: Function evaluations budgets for each algorithm in our portfolio for 2 and 7 generations.**

Algorithm Trajectory	2 generations	7 generations
CMA-ES	20	70
DE	60	210
PSO	80	280
ALL	160	560

solver and (b) a regressor that predicts performance. The following section describes the model inputs, the models used, and the validation procedures that are applied.

#### 4.1 Inputs for ML models

Models can be trained with three types of input:

- *Raw Probing-Trajectories*: these consist of a time-series of fitness values from the first  $n$  function evaluations of an algorithm. Two types of trajectory are recorded. The ‘current’ trajectory simply logs the fitness obtained at every function evaluation in the order that they occur. On the other hand, the ‘best’ trajectory records the *best-so-far* fitness seen during a run after each function evaluation. We use the labels ‘best’ and ‘current’ to refer to these trajectories from this point onwards.

These trajectories are obtained from the SA algorithm or directly from the data published by [47]. This archive contains data describing individual trajectories of CMA-ES, PSO, and DE. Following the approach of [39], we also consider ‘concatenated’ trajectories from each of three algorithms which we refer to as the ‘ALL’ trajectory (i.e. it combines individual trajectories into a single new trajectory that can be used to train a model). For the portfolio solvers CMA-ES/PSO/DE, we use trajectories obtained over 2 generations and separately over 7 generations. The data we obtained from [47] was generated using a different population size for each solver (due to parameter tuning) hence the length of the trajectories generated for each solver are slightly different. The values are provided in Table 1.

The input to a model is always a time-series representing one algorithm run. As just described, the length of the time-series depends on the number of generations and population size used in the algorithms.

- *Time-series Features*: As in [9, 39], we also extract time-series features from all trajectories defined in the previous section, using the *tsfresh* Python package<sup>3</sup>, and perform feature selection using the *Boruta* Python package<sup>4</sup>. Input is a vector of features, representing one algorithm run.
- *ELA Features* From the 100 10–dimensional feature vectors that are available for each BBOB instances (see Section 3.2), we randomly sample 5 vectors. This ensures fair comparison between features and trajectories as data from 5 runs is

available for calculating trajectories. We compare two budgets for features: the recommended budget of  $50d = 500$  sample points and  $30d = 300$  sample points. In this case, the input is a 10–dimensional vector of features, representing one function instance.

#### 4.2 Models

The *classification* task is a typical algorithm-selection task, i.e., given an input, output the best algorithm to use from a given portfolio. As noted previously, the best algorithm is defined as the algorithm having the lowest median performance after 100,000 function evaluations. There is no under-represented label: out of the 24 functions, CMA-ES is the best performing algorithm for 11 functions, DE for 7, and PSO for 6. We use the *classification accuracy* as the metric to evaluate the model. The model used depends on the type of input:

- *Feature-based*: inputs can be derived from calculating either ELA features or time-series features extracted from the probing-trajectories. We train a default Random Forests [7] from the *scikit-learn* package [32]<sup>5</sup> as performed in [4, 37]. We train separate models using ELA features or time-series features.
- *Time-series-based*: for trajectory-based input, we use a specialised time-series classifier, specifically the default Rotation Forests [41] from the *sktime* package<sup>6</sup>. Although other time-series classifiers could have been chosen we selected this as it is closely related to the classifier used for features.

The *regression* task consists of a performance prediction of one algorithm from our portfolio after 100,000 function evaluations. Here, the performance measure used for the regression task is the *Root Mean Squared Error (RMSE)*. Note the performances of algorithms are not normalised.

- *Feature-based*: we use the default Random Forests [7] from the *scikit-learn* package [32], following the method used in [14, 34]. Models were trained using ELA features and time-series features. As the latter did not lead to any improvements on the RMSE, results are not shown for this model for space reasons.
- *Time-series-based*: we use a specialised time-series regressor, specifically the Time Series Forests from the *sktime* package<sup>7</sup>. As for the classification, it is chosen for its closeness to the classifier used for features.

#### 4.3 Validation

Both classification and regression share the same validation procedure. Following the method described in [22], we perform a *leave-one-instance-out (LOIO) cross-validation*. All Classifiers and regressors are trained using data from all runs of the 24 functions on all except one instance. The runs from the 24 functions on the instance left out are used as the validation set. Overall,  $24 \times (5-1) \times 5 = 480$  inputs are used to train the model while the remaining  $24 \times 1 \times 5 = 120$  inputs are used for validation.

<sup>3</sup>version 0.20.1

<sup>4</sup>version 0.3

<sup>5</sup>version 1.1.3

<sup>6</sup>version 0.16.1

<sup>7</sup>version 0.16.1

**Table 2: SA parameters for the best configuration found by *irace* for each ML task.**

Task	# samples	T	visit	acceptance
Default	100	5230	2.62	-5
Classification current	99	23,318.74	2.185	-8,848.35
Classification best	99	15,912.03	1.831	-5,110.81
Regression CMA-ES current	62	19,578.93	1.998	-1,053.54
Regression CMA-ES best	85	21,838.23	2.299	-7,878.54
Regression PSO current	66	42,832.28	2.492	-3,601.99
Regression PSO best	34	49,602.83	1.722	-10,325.21
Regression DE current	36	45,778.60	2.053	-5,925.62
Regression DE best	64	4,315.52	1.938	-6,135.95

## 5 RESULTS

In this section, we present the results obtained for both ML tasks, i.e., algorithm selection performed with a classifier (Section 5.2) and performance prediction of the three algorithms in our portfolio (Section 5.3).

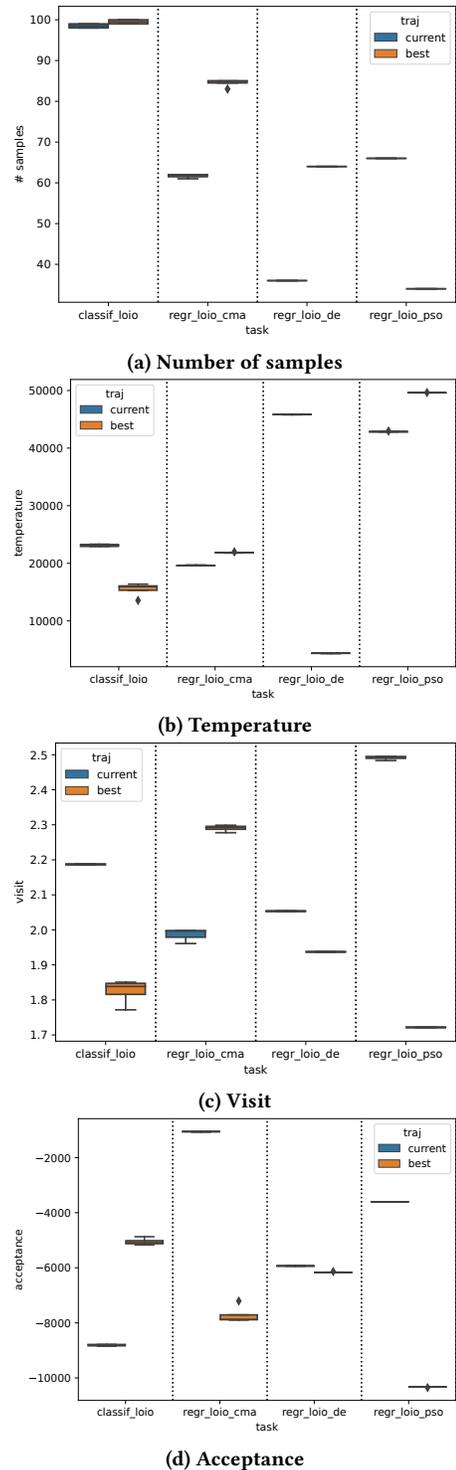
### 5.1 Tuned SA Parameters

The motivation for the meta-algorithm proposed is to learn hyper-parameter values for an SA algorithm that produces discriminatory trajectories. Using *irace*, we tune SA separately on eight different tasks: classification using best/current trajectories; regression for each of three solvers, using best/current trajectories).

The parameters obtained in each case are shown in Table 2 and used to create test trajectories for all results shown below. It is clear from this table that for each scenario, a unique set of hyper-parameters is obtained. By default, *irace* outputs the four best configurations found. However in two tasks, *irace* returns fewer configurations: three for the prediction of CMA-ES performances using the ‘current’ trajectory and two for the prediction of DE performances using the ‘current’ trajectory. Figure 2 shows the distributions of the four tuned parameters (i.e., number of samples, temperature, visit, and acceptance) for each task across all output configurations. We observe that the tuned parameters vary widely with respect to the default configuration of SA that uses 100 samples, a temperature of 5,230, a visit of 2.62 and an acceptance of -5. The variance of the distribution of configurations found by *irace* is small for each task indicating that there do not seem to be multiple optima in the search-space.

### 5.2 Algorithm Selection

Figure 3 compares the classification accuracies obtained with ELA features and trajectories (individual trajectories from algorithms in the portfolio, concatenated trajectories from algorithms in the



**Figure 2: Distributions of SA parameters found by *irace* for the different ML tasks: classification, performance prediction of CMA-ES, PSO, and DE.**

portfolio, default SA trajectories used as baseline (*SA*), and tuned SA trajectories (*SA – best*, *SA – current*). Following [39], we display results for trajectories obtained over 2 and 7 generations for each of the portfolio algorithms. ‘Best’ trajectories are displayed for 2 generations (Figure 3a) and ‘current’ trajectories for 7 generations (Figure 3b) to save space (additional plots are available at [38]). For each of the different methods of obtaining trajectories, we show box-plots for classifiers trained on trajectories and time-series features (with/without selection), and indicate median ELA performance as a line.

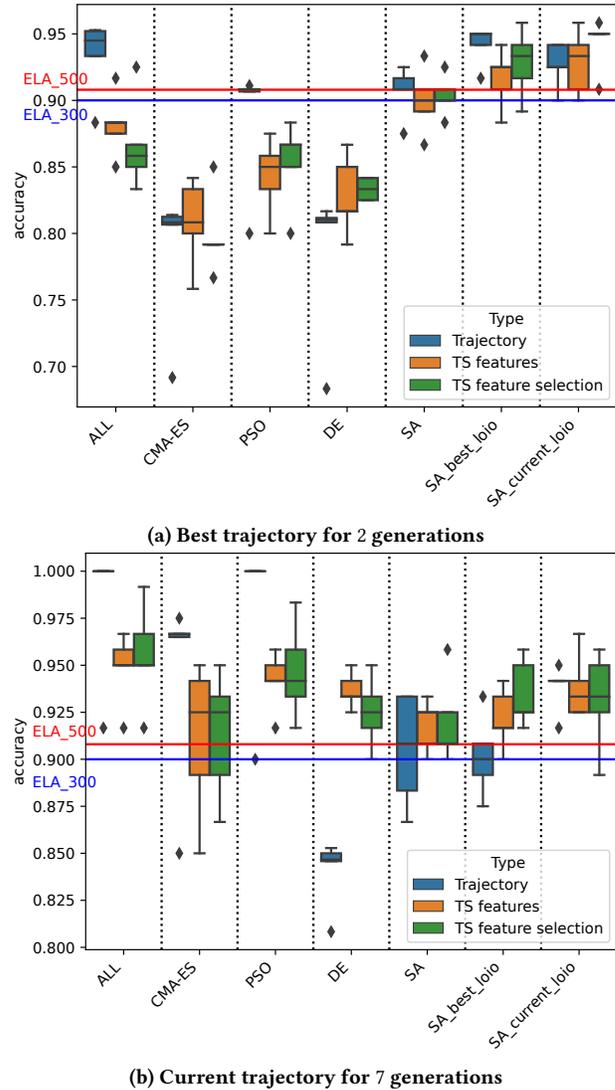
Figure 3a where trajectories are obtained from only 2 generations shows that classifiers trained on SA trajectories (tuned or by default) and on the ‘ALL’ trajectories outperform classifiers trained on ELA features at both budgets with respect to median fitness. Recall however that the SA trajectories are obtained from only 100 or fewer samples depending on the tuned sample-size, i.e. they use three to five times less budget than ELA features. *SA\_best\_loio* has a median accuracy of 95% compared to 94.5% for the ‘ALL’ classifier. *SA\_current\_loio* has a 92.5% median accuracy using the raw-trajectory, rising to 95% accuracy when time-series feature selection is used. The tuned SA trajectories result in better accuracy than the untuned version (as expected). Classifiers trained on the raw trajectory data outperform the feature-based approaches as a rule of thumb. Training on individual trajectories performs poorly in comparison in each case.

Figure 3b shows results when trajectories are collected over 7 generations using ‘current’ trajectories. In contrast to the previous results, all methods except the DE classifier have a median accuracy better than those obtained using ELA features. While the SA trajectories outperform ELA features, they are outperformed by other classifiers: *SA\_current\_loio* has a median accuracy of 94.2% which is lower than all other classifiers except for the one trained on DE trajectories. Nevertheless, we highlight that the SA methods require significantly fewer function evaluations to achieve this result than every other algorithm presented here except CMA-ES. SA is also more robust than CMA-ES which has one outlier with 85% accuracy, while SA consistently performs above 90%.

Overall, tuning SA to find discriminating trajectories to train an algorithm-selector is a good low-budget approach in that it requires significantly fewer function evaluations than ELA features and some of the trajectories approaches. It also provides some confidence the trajectories are discriminatory: this is not guaranteed when calculating trajectories from the three solvers in the portfolio due to the stochastic nature of the methods. Furthermore, if the portfolio of algorithms is large, this introduces an additional burden due to the number of combinations of trajectories that could be concatenated, making it more difficult to select the best combination to train with. The budget required for this may be more expensive than calculating ELA features for a large portfolio.

### 5.3 Performance Prediction

Figure 4 shows the RMSE of the performance prediction of CMA-ES, PSO, and DE using trajectories and ELA features for 7 generations on both ‘best’ and ‘current’ trajectories. Note that this should be minimised. For each solver in the portfolio, we show the RMSE of



**Figure 3: Accuracy of classification on the LOIO cross-validation for best-so-far and current probing-trajectories, time series features and time series feature selection for 2 and 7 generations. Median ELA feature accuracy is represented by lines for 300 and 500 function evaluations.**

the predicted performance obtained from the set of regressors<sup>8</sup>. For each regressor, we omit the results from training using time-series features as they do not improve results.

We observe that for CMA-ES (Figure 4a) and DE (Figure 4c) performance predictions, all of the trajectory-based regressors outperform or match ELA performances. However, for PSO, only the tuned SA-based regressors outperform ELA. All of the other methods result in an RMSE that is worse than predicted by the ELA regressors.

<sup>8</sup>Similar results are obtained using 2 generations – results are not included for space reasons

The default SA configuration ‘current’ trajectory outperforms all regressors trained on individual trajectories from the original portfolio (CMA-ES/DE/PSO/ALL) with a median RMSE of 5.91. However, the best result is obtained from the tuned SA approach *SA\_current\_CMA* with a median RMSE of 5.28. Again, we remind the reader that this is also the approach using the smallest budget of function evaluations with 62 (close to the CMA-ES trajectory 70 evaluations and nine times less than the ‘ALL’ trajectory).

For predicting DE, the default configuration of SA with the ‘best’ trajectory also outperforms other trajectory-based regressors where the algorithm comes from the initial portfolio and also outperforms the regressor trained on ELA features (Figure 4c). For this task, both specifically tuned SA methods have similar performance with a median RMSE of 23.95 for *SA\_current\_DE* on the ‘current’ trajectory and 24.6 for *SA\_best\_DE* on the ‘best’ trajectory. For this task, it is more advantageous to use *SA\_current\_DE*: its performance is slightly better but its computational cost is lower, i.e., 36 function evaluations against 64 for *SA\_best\_DE*. Using only 36 function evaluations (i.e. the tuned value from *irace*), *SA\_current\_DE* uses almost half CMA-ES budget to compute trajectories and 15 times less evaluations than the ‘ALL’ trajectory.

Prediction of performances of PSO (Figure 4b) is the only regression task where only the specifically tuned SA methods outperform ELA features. All other models trained using trajectories obtain worse RMSE. Both tuned SA methods achieve similar RMSE with a median of 54.4 for *SA\_current\_PSO* and 60.6 for *SA\_best\_PSO*. We suggest *SA\_best\_PSO* seems more beneficial from a computational point of view as it requires only 34 function evaluations against 66 for *SA\_current\_PSO*.

To summarise, it is clear that trajectory-based algorithm selection and trajectory-based performance prediction provide performances that are equal to or better than the metrics obtained using ELA features as input, using far less computational budget. Tuning SA to obtain discriminating trajectories is the only trajectory-based approach to consistently outperform ELA features, and also outperforms other trajectory-based models.

### 5.4 Transfer Learning to reduce computation

In previous sections, we showed that tuning SA to produce discriminatory trajectories outperforms the state-of-the-art. For the regression task, SA was tuned separately to create trajectories to train regressors for each of the three solvers (CMA-ES, DE, PSO). The ML literature suggests that often there is potential for transfer-learning (TL) [31] between models – this aims to improve the performance of learners on target domains by transferring the knowledge contained in different but related source domains. Therefore we now assess whether TL can be exploited to reduce the computational burden of tuning per solver. For example, we examine whether the hyper-parameters used to produce discriminatory trajectories for PSO and DE can also be used to predict CMA-ES performances.

Figure 5 shows the RMSE for performance prediction of CMA-ES using regressors trained on trajectories tuned on DE, PSO and CMA-ES respectively. As expected, the specialists *SA\_current\_CMA* and *SA\_best\_CMA* outperform all other configurations (median performance 5.91, 5.28 respectively for the ‘best’ and ‘current’ trajectory models). Nevertheless, using trajectories obtained from SA with

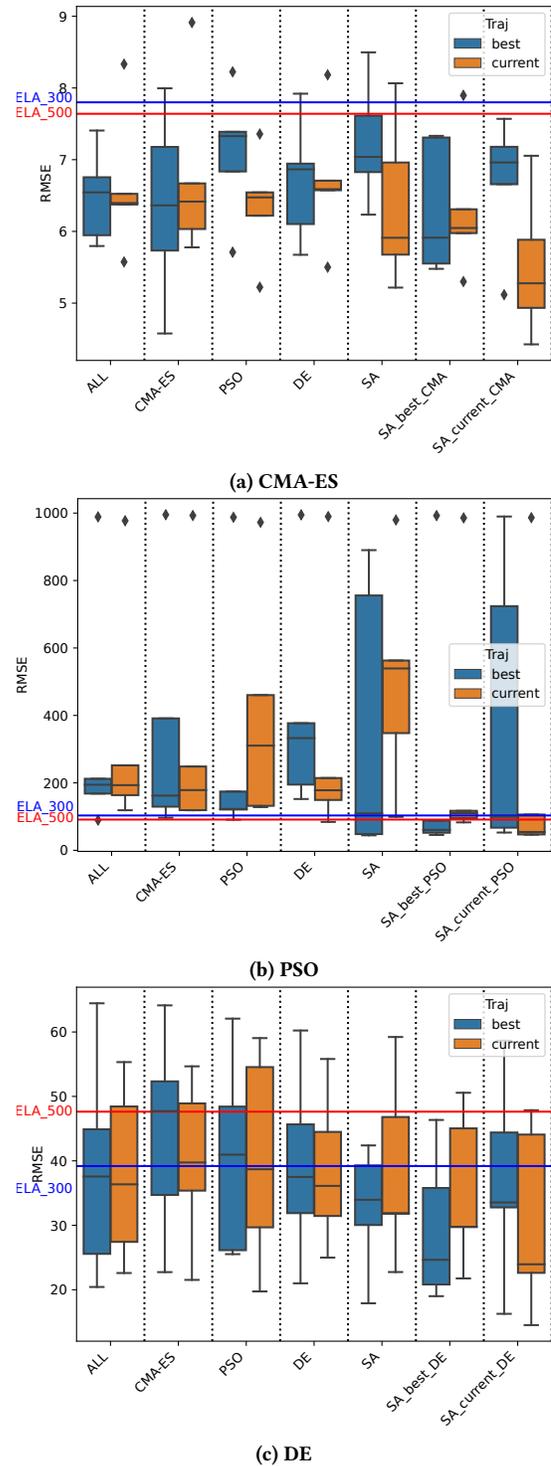
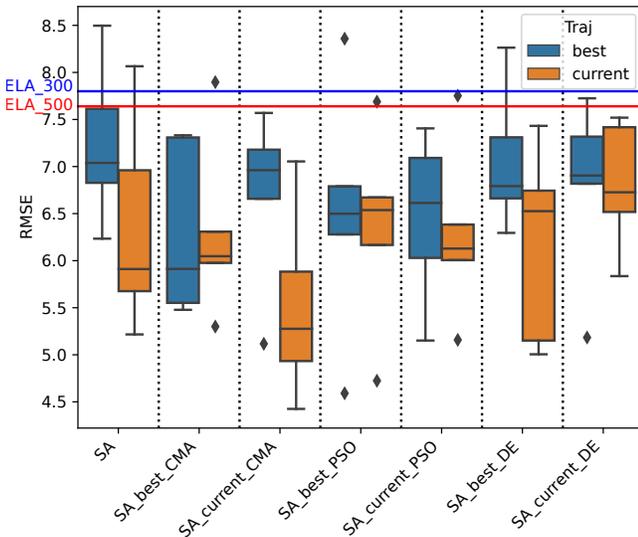


Figure 4: RMSE of performance prediction of CMA-ES, PSO, and DE on the LOIO cross-validation for best-so-far and current probing-trajectories for 7 generations. Median ELA feature RMSE is represented by lines for 300 and 500 function evaluations.



**Figure 5: RMSE of performance prediction of CMA-ES on the LOIO cross-validation for best-so-far and current probing-trajectories for all tuned SA (including other tasks). Median ELA features RMSE are represented by lines for 300 and 500 function evaluations.**

hyper-parameters tuned on DE or PSO to predict CMA-ES performance is surprisingly robust: using parameters tuned for PSO gives a median RSME of 6.5 (compared to 5.91 from the specialist), while parameters tuned for DE have median RMSE 6.79. We also note that tuning on a ‘best’ trajectory and re-using the parameters to generate a ‘current’ trajectory for a solver often results in a noticeable loss in RSME: for example, using parameters tuned on the CMA-ES ‘current’ trajectories to obtain trajectories for CMA-ES ‘best’ increases the median RMSE from 5.91 to 6.96. Finally, it is worth observing that *all* methods provide better results than regressors trained on ELA features: that is, tuning SA to generate discriminatory trajectories for *any* of the solvers in the portfolio provides a set of hyper-parameters that can be re-used to create trajectories for other solvers in the portfolio that can be used to predict performance more accurately than ELA features.

Given that the ‘best’ and ‘current’ trajectories have different configurations of SA that can be very different, it seems not possible to transfer the tuned parameters to the other type of trajectory without having a loss in performance.

## 6 CONCLUSION

We addressed the issue of algorithm-selection and performance prediction in a continuous optimisation setting. Previous work [39] has shown that using data that is algorithm-centric (i.e. trajectories) to train models can outperform ML models trained using features (e.g., ELA), with the additional advantage of being low-budget in terms of the number of function evaluations needed to create a trajectory. However there are some weaknesses with this approach: (1) it is difficult to ensure that trajectories are sufficiently discriminatory to train a high-performing model and (2) as a trajectory

needs to be generated per solver in a portfolio, the approach does not scale well.

To address this, we proposed a meta-algorithm that tunes the hyper-parameters of a simple solver (Simulated Annealing) to generate trajectories that when used as input data to an ML model, improve its performance metric (either classification or regression). This directly addresses both weaknesses identified above. Firstly, the meta-algorithm learns to create discriminatory trajectories. Secondly, rather than generating trajectories from each solver in a portfolio as in [39], only one trajectory needs to be generated using the tuned SA algorithm. Furthermore, the SA trajectories only use 100 or fewer samples and, therefore are obtained at a very low budget. We highlight the following findings:

- The ML performance-metrics obtained using models that use trajectories from a tuned SA algorithm as input outperform the same metrics obtained from models that use ELA features as input, using considerably less budget (at minimum,  $\approx 1/3$  of the ELA budget).
- For classification models that predict the label of the best solver, at low-budget (2 generations), models using SA trajectories from the tuned SA algorithm as input have similar median accuracy to those obtained using an ‘ALL’ trajectory, but use  $\approx 62\%$  of the budget used by ‘ALL’.
- For all three regression models that predict performance (CMA-ES, DE, PSO), the best RMSE is obtained by a model that uses an SA trajectory as input.
- Although using the meta-algorithm to tune each model individually leads to the best results, the loss in RMSE when using a set of hyper-parameters tuned for one solver to obtain trajectories for a different solver is very small, suggesting that budget could be saved with the trade-off of a small reduction in performance.

Obvious next steps include testing the approach on a larger portfolio of solvers and on different benchmarks. With respect to the latter, it would be particularly interesting to consider combinatorial optimisation domains here although this may exclude a comparison to ELA-based methods as the use of landscape features is much less common in combinatorial domains.

Given that the method relies on creating a time-series of data, then a more thorough evaluation of state-of-the-art time-series classifiers (and tuning of their hyper-parameters) is also likely to improve results. We investigated tuning the hyper-parameters of a Simulated Annealing algorithm to create the discriminatory trajectories but clearly, other algorithms could replace this. We stress however that the motivation behind this work is to generate data at a low-cost, therefore this might restrict the type of algorithm that could be used. It would also be interesting to design a specialist algorithm whose only purpose would be to generate discriminatory trajectories. Finally, further work in understanding the extent to which transfer-learning could be exploited to further reduce computation would be beneficial.

## ACKNOWLEDGEMENTS

The authors are supported by funding from EPSRC award number(s): EP/V026534/1

## REFERENCES

- [1] M. Alissa, K. Sim, and E. Hart. 2019. Algorithm selection using deep learning without feature extraction. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 198–206.
- [2] M. Alissa, K. Sim, and E. Hart. 2023. Automated Algorithm Selection: from Feature-Based to Feature-Free Approaches. *J. Heuristics* 29, 1 (2023), 1–38. <https://doi.org/10.1007/s10732-022-09505-4>
- [3] N. Belkhir. 2017. *Per Instance Algorithm Configuration for Continuous Black Box Optimization*. phdthesis. Université Paris-Saclay. <https://hal.inria.fr/tel-01669527/document>
- [4] N. Belkhir, J. Dréo, P. Savéant, and M. Schoenauer. 2017. Per instance algorithm configuration of CMA-ES with limited budget. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17*. ACM, 681–688. <https://doi.org/10.1145/3071178.3071343>
- [5] D.M. Bossens, JB. Mouret, and D. Tarapore. 2020. Learning behaviour-performance maps with meta-evolution. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. 49–57.
- [6] D.M. Bossens and D. Tarapore. 2022. Quality-Diversity Meta-Evolution: Customizing Behavior Spaces to a Meta-Objective. *IEEE Transactions on Evolutionary Computation* 26, 5 (2022), 1171–1181.
- [7] L. Breiman. 2001. Random Forests. *Mach. Learn.* 45, 1 (2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [8] G. Cenikj, G. Petelin, C. Doerr, P. Korosec, and T. Eftimov. 2023. DynamoRep: Trajectory-Based Population Dynamics for Classification of Black-box Optimization Problems. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2023, Lisbon, Portugal, July 15-19, 2023*. ACM, 813–821. <https://doi.org/10.1145/3583131.3590401>
- [9] J. de Nobel, H. Wang, and T. Bäck. 2021. Explorative data analysis of time series based algorithm features of CMA-ES variants. In *GECCO '21: Genetic and Evolutionary Computation Conference, Lille, France, July 10-14, 2021*. ACM, 510–518. <https://doi.org/10.1145/3449639.3459399>
- [10] B. Derbel, A. Liefoughe, S. Vérel, H. Aguirre, and K. Tanaka. 2019. New features for continuous exploratory landscape analysis based on the SOO tree. In *Proceedings of Foundations of Genetic Algorithms (FOGA) '19*. ACM, 72–86. <https://doi.org/10.1145/3299904.3340308>
- [11] S. Finck, N. Hansen, R. Ros, and A. Auger. 2010. Real-Parameter Black-Box Optimization Benchmarking 2010: Presentation of the Noiseless Functions. <http://coco.gforge.inria.fr/downloads/download16.00/bbobdfnfunctions.pdf>.
- [12] N. Hansen, A. Auger, R. Ros, O. Mersmann, T. Tausar, and D. Brockhoff. 2021. COCO: a platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* 36, 1 (2021), 114–144. <https://doi.org/10.1080/10556788.2020.1808977>
- [13] 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>
- [14] A. Jankovic, T. Eftimov, and C. Doerr. 2021. Towards Feature-Based Performance Regression Using Trajectory Data. In *Applications of Evolutionary Computation - 24th International Conference, EvoApplications 2021, Held as Part of EvoStar 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12694)*. Springer, 601–617. [https://doi.org/10.1007/978-3-030-72699-7\\_38](https://doi.org/10.1007/978-3-030-72699-7_38)
- [15] A. Jankovic, D. Vermetten, A. Kostovska, J. de Nobel, T. Eftimov, and C. Doerr. 2022. Trajectory-based Algorithm Selection with Warm-starting. In *IEEE Congress on Evolutionary Computation, CEC 2022, Padua, Italy, July 18-23, 2022*. IEEE, 1–8. <https://doi.org/10.1109/CEC55065.2022.9870222>
- [16] 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>
- [17] P. Kerschke, H.H. Hoos, F. Neumann, and H. Trautmann. 2019. Automated Algorithm Selection: Survey and Perspectives. *Evolutionary Computation* 27, 1 (March 2019), 3–45.
- [18] P. Kerschke, L. Kotthoff, J. Bossek, H.H. Hoos, and H. Trautmann. 2018. Leveraging TSP Solver Complementarity through Machine Learning. *Evol. Comput.* 26, 4 (2018). [https://doi.org/10.1162/EVCO\\_a\\_00215](https://doi.org/10.1162/EVCO_a_00215)
- [19] 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>
- [20] 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>
- [21] S. Kirkpatrick, D. Gelatt Jr., and M.P. Vecchi. 1983. Optimization by Simulated Annealing. *Sci.* 220, 4598 (1983), 671–680. <https://doi.org/10.1126/science.220.4598.671>
- [22] A. Kostovska, A. Jankovic, D. Vermetten, J. de Nobel, H. Wang, T. Eftimov, and C. Doerr. 2022. Per-run Algorithm Selection with Warm-Starting Using Trajectory-Based Features. In *Parallel Problem Solving from Nature - PPSN XVII - 17th International Conference, PPSN 2022, Dortmund, Germany, September 10-14, 2022, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13398)*. Springer, 46–60. [https://doi.org/10.1007/978-3-031-14714-2\\_4](https://doi.org/10.1007/978-3-031-14714-2_4)
- [23] J. Lehman and K.O. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation* 19, 2 (2011), 189–222. [https://doi.org/10.1162/EVCO\\_a\\_00025](https://doi.org/10.1162/EVCO_a_00025)
- [24] A. Lensen, B. Xue, and M. Zhang. 2019. Can genetic programming do manifold learning too?. In *European Conference on Genetic Programming*. Springer, 114–130.
- [25] 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.
- [26] A. Loreggia, Y. Malitsky, H. Samulowitz, and V. A. Saraswat. 2016. Deep Learning for Algorithm Portfolios. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*. AAAI Press, 1280–1286. <https://doi.org/10.1609/AAAI.V30I1.10170>
- [27] 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>
- [28] E. Meyerson, J. Lehman, and R. Miikkilainen. 2016. Learning Behavior Characterizations for Novelty Search. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 (Denver, Colorado, USA) (GECCO '16)*. Association for Computing Machinery, New York, NY, USA, 149–156. <https://doi.org/10.1145/2908812.2908929>
- [29] JB. Mouret and F. Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [30] M.A. Muñoz, M. Kirley, and S.K. Halgamuge. 2015. Exploratory Landscape Analysis of Continuous Space Optimization Problems Using Information Content. *IEEE Transactions on Evolutionary Computation* 19, 1 (Feb. 2015), 74–87. <https://doi.org/10.1109/TEVC.2014.2302006>
- [31] S. Niu, Y. Liu, J. Wang, and H. Song. 2020. A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence* 1, 2 (2020), 151–166.
- [32] 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.
- [33] J.K. Pugh, L.B. Soros, and K.O. Stanley. 2016. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3 (2016), 40.
- [34] Q. Renau. 2022. *Landscape-Aware Selection of Metaheuristics for the Optimization of Radar Networks*. Ph.D. Dissertation. Polytechnic Institute of Paris, Palaiseau, France. <https://tel.archives-ouvertes.fr/tel-03593606>
- [35] Q. Renau, J. Dreó, C. Doerr, and B. Doerr. 2021. Exploratory Landscape Analysis Feature Values for the 24 Noiseless BBOB Functions. <https://doi.org/10.5281/zenodo.4449934>
- [36] Q. Renau, J. Dreó, C. Doerr, and B. Doerr. 2021. Towards Explainable Exploratory Landscape Analysis: Extreme Feature Selection for Classifying BBOB Functions. In *Applications of Evolutionary Computation - 24th International Conference, EvoApplications 2021, Held as Part of EvoStar 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12694)*. Springer, 17–33. [https://doi.org/10.1007/978-3-030-72699-7\\_2](https://doi.org/10.1007/978-3-030-72699-7_2)
- [37] Q. Renau, J. Dreó, A. Peres, Y. Semet, C. Doerr, and B. Doerr. 2022. Automated algorithm selection for radar network configuration. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '22*, Jonathan E. Fieldsend and Markus Wagner (Eds.). ACM, 1263–1271. <https://doi.org/10.1145/3512290.3528825>
- [38] Q. Renau and E. Hart. 2024. Improving Algorithm-Selectors and Performance-Predictors via Learning Discriminating Training Samples - Code and Data. <https://doi.org/10.5281/zenodo.10579661>
- [39] Q. Renau and E. Hart. 2024. On the Utility of Probing Trajectories for Algorithm-Selection. [arXiv:2401.12745](https://arxiv.org/abs/2401.12745)
- [40] J.R. Rice. 1976. The Algorithm Selection Problem. *Advance Computing* 15 (1976), 65–118. [https://doi.org/10.1016/S0065-2458\(08\)60520-3](https://doi.org/10.1016/S0065-2458(08)60520-3)
- [41] J.J. Rodríguez, L.I. Kuncheva, and C.J. Alonso. 2006. Rotation Forest: A New Classifier Ensemble Method. *IEEE Trans. Pattern Anal. Mach. Intell.* 28, 10 (2006), 1619–1630. <https://doi.org/10.1109/TPAMI.2006.211>
- [42] M. Seiler, J. Pohl, J. Bossek, P. Kerschke, and H. Trautmann. 2020. Deep learning as a competitive feature-free approach for automated algorithm selection on the traveling salesperson problem. In *International Conference on Parallel Problem Solving from Nature*. Springer, 48–64.
- [43] K. Sim and E. Hart. 2022. Evolutionary Approaches to Improving the Layouts of Instance-Spaces. In *Parallel Problem Solving from Nature - PPSN XVII - 17th International Conference, PPSN 2022, Dortmund, Germany, September 10-14, 2022, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13398)*. Springer, 207–219. [https://doi.org/10.1007/978-3-031-14714-2\\_15](https://doi.org/10.1007/978-3-031-14714-2_15)
- [44] K. Smith-Miles, D. Baatar, B. Wreford, and R. Lewis. 2014. Towards objective measures of algorithm performance across instance space. *Computers & Operations*

- Research* 45 (2014), 12–24.
- [45] 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>
- [46] R. Tanabe. 2022. Benchmarking Feature-Based Algorithm Selection Systems for Black-Box Numerical Optimization. *IEEE Trans. Evol. Comput.* 26, 6 (2022), 1321–1335. <https://doi.org/10.1109/TEVC.2022.3169770>
- [47] D. Vermetten, W. Hao, K. Sim, and E. Hart. 2022. To Switch or not to Switch: Predicting the Benefit of Switching between Algorithms based on Trajectory Features - Dataset. <https://doi.org/10.5281/zenodo.7249389>
- [48] D. Vermetten, H. Wang, K. Sim, and E. Hart. 2023. To Switch or Not to Switch: Predicting the Benefit of Switching Between Algorithms Based on Trajectory Features. In *Applications of Evolutionary Computation*. Springer Nature Switzerland, Cham, 335–350.
- [49] 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>
- [50] K. Zhao, S. Liu, Jeffrey X. Yu, and Y. Rong. 2021. Towards feature-free TSP solver selection: A deep learning approach. In *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.