

Sequential vs. Integrated Algorithm Selection and Configuration: A Case Study for the Modular CMA-ES

Diederick Vermetten¹ and Hao Wang¹ and Carola Doerr² and Thomas Bäck¹

Abstract. When faced with a specific optimization problem, choosing which algorithm to use is always a tough task. Not only is there a vast variety of algorithms to select from, but these algorithms often are controlled by many hyperparameters, which need to be tuned in order to achieve the best performance possible. Usually, this problem is separated into two parts: *algorithm selection* and *algorithm configuration*. With the significant advances made in Machine Learning, however, these problems can be integrated into a combined algorithm selection and hyperparameter optimization task, commonly known as the CASH problem.

In this work we compare sequential and integrated algorithm selection and configuration approaches for the case of selecting and tuning the best out of 4,608 variants of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) tested on the Black Box Optimization Benchmark (BBOB) suite. We first show that the ranking of the modular CMA-ES variants depends to a large extent on the quality of the hyperparameters. This implies that even a sequential approach based on complete enumeration of the algorithm space will likely result in sub-optimal solutions. In fact, we show that the integrated approach manages to provide competitive results at a much smaller computational cost.

We also compare two different mixed-integer algorithm configuration techniques, called *irace* and Mixed-Integer Parallel Efficient Global Optimization (MIP-EGO). While we show that the two methods differ significantly in their treatment of the exploration-exploitation balance, their overall performances are very similar.

1 Introduction

In computer science, optimization has become an important field of study over the past decades. Because of its rising popularity and its high practical relevance, many different techniques have been introduced to solve particular types of optimization problems. As these methods are developed further, small modifications might lead the algorithm to behave better on specific problem types. However, it has long been known that no single algorithm variant can outperform all others on all functions, as stated in the *no-free-lunch theorem* [46]. This fact leads a new set of challenges for practitioners and researchers alike: How to choose which algorithm to use for which problem?

Even when limiting the scope to a small class of algorithms, the choice of which variant to choose can be daunting, leading practitioners to resort to a few standard versions of the algorithms, which might not be particularly well suited to their problem. The problem of selecting an algorithm (variant) from a large set is commonly referred

to as the *algorithm selection* problem [23]. However, the algorithm variant is not the only factor having an impact on performance. The setting of the variable hyperparameters can also play a very important role [27, 7]. The problem of choosing the right hyperparameter setting for a specific algorithm is commonly referred to as the *algorithm configuration* problem [11].

Naturally, the *algorithm selection* and *algorithm configuration* problems are highly interlinked. Because of this, it is natural to attempt to tackle both problems at the same time. Such an approach is commonly referred to as the *Combined Algorithm Selection and Hyperparameter optimization* (CASH) task, which was introduced in [35] and later studied in [13, 14, 25]. Note, though, that the by far predominant approach in real-world algorithm selection and configuration is still a sequential approach, in which the user first selects one or more algorithms (typically based on previous experience) and then tunes their parameters (either manually or using one of the many existing software frameworks, such as [12, 20, 26, 28, 6]), before deciding which algorithmic variant to use for their problem at hand. In fact, we observe that the tuning step is often neglected, and standard solvers are run with some default configurations which have been suggested in the research literature (or happen to be the defaults in the implementation). Efficiently solving the CASH problem is therefore far from easy, and far from being general practice.

In this work, we address the CASH problem in the context of selecting and configuring variants of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES). The CMA-ES family [19] is an important collection of heuristic optimization techniques for numeric optimization. In a nutshell, the CMA-ES is an iterative search procedure, which updates after each iteration the covariance matrix of the multivariate normal distribution that is used to generate the samples during the search, effectively learning second-order information about the objective function. Important contributions to the class of CMA-ES have been made over the years, which all reveal different strengths and weaknesses in different optimization contexts [5, 16]. While most of the suggested modifications have been proposed in isolation, [38] suggested a framework which modularizes eleven popular CMA-ES modifications such that they can be combined to create a total number of 4,608 different CMA-ES variants. It was shown in [39] that some of the so-created CMA-ES variants improve significantly over commonly used CMA-ES algorithms. This modular CMA-ES framework, which is available at [36], provides a convenient way to study the impact of the different modules on different optimization problems [38].

The modularity of this framework allows us to integrate the *algorithm selection* and *configuration* into a single mixed-integer search space, where we can optimize both the algorithm variant and the corresponding hyperparameters at the same time. We show that such an

¹ Leiden University, The Netherlands

² Sorbonne Université, CNRS, LIP6, Paris, France

integrated approach is competitive with sequential approaches based on complete enumeration of the algorithm space, while requiring significantly less computational effort. We also investigate the differences between two algorithm configuration tools, irace [28] and MIP-EGO [44] (see Section 2.4 for short descriptions). While the overall performance of these two approaches is comparable, the balance between algorithm selection and algorithm configuration shows significant differences, with irace focusing much more on the configuration task, and evaluating only few different CMA-ES variants. MIP-EGO, in turn, shows a broader exploration behavior, at the cost of less accurate performance estimates.

2 Experimental Setup

We summarize the algorithmic framework, the benchmark suite, the performance measures, and the two configuration tools, irace and MIP-EGO, which we employ for the tuning of the hyperparameters.

2.1 The Modular CMA-ES

Table 1 summarizes the eleven modules of the modular CMA-ES from [38]. Out of these, nine modules are binary and two are ternary, allowing for the total number of 4,608 different possible CMA-ES variants.

So far, all studies on the modular CMA-ES framework have used default hyperparameter values [37, 38, 41]. However, it has been shown that substantial performance gains are possible by tuning these hyperparameters [1, 7], raising the question how much can be gained from combining the tuning of several hyperparameters with the selection of the CMA-ES variant. In accordance with [7], we focus on only a small subset of these hyperparameters, namely c_1 , c_c and c_μ , which control the update of the covariance matrix. It is well known, though, that other hyperparameters, and in particular the population size [3] have a significant impact on the performance as well, and might be much more critical to configure as the ones chosen in [7]. However, we will see that the efficiency of the CMA-ES variants is nevertheless strongly influenced by these three hyperparameters. In fact, we show that the ranking of the algorithm variants with default and tuned hyperparameters can differ significantly, indicating that a sequential execution of *algorithm selection* and *algorithm configuration* will not provide optimal results.

#	Module name	0	1	2
1	Active Update [22]	off	on	-
2	Elitism	(μ, λ)	$(\mu + \lambda)$	-
3	Mirrored Sampling [8]	off	on	-
4	Orthogonal Sampling [43]	off	on	-
5	Sequential Selection [8]	off	on	-
6	Threshold Convergence [32]	off	on	-
7	TPA [15]	off	on	-
8	Pairwise Selection [2]	off	on	-
9	Recombination Weights [4]	$r_i(\mu)$	$\frac{1}{\mu}$	-
10	Quasi-Gaussian Sampling	off	Sobol	Halton
11	Increasing Population [3, 16]	off	IPOP	BIPOP

Table 1: Overview of the CMA-ES modules available in the used framework. The entries in row 9 specify the formula for calculating each weight w_i , with $r_i := \log(\mu + \frac{1}{2}) - \frac{\log(i)}{\sum_j w_j}$.

2.2 Test-bed: the BBOB Framework

For analyzing the impact of the hyperparameter tuning, we use the Black-Box Optimization Benchmark (BBOB) suite [17], which is a standard environment to test black-box optimization techniques. This testbed contains 24 functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$, of which we use the five-dimensional versions. Each function can be transformed in objective and variable space, resulting in separate instances with similar fitness landscapes. A large part of our analysis is built on data from [41], which uses the first 5 instances of all functions, for which 5 independent runs were performed on each instance, for each algorithm variant, and each function. This data is available at [40].

2.3 Performance Measures

We next define the performance measures by which we compare the different algorithms. First note that CMA-ES is a stochastic optimization algorithm. The number of function evaluations needed to find a solution of a certain quality is therefore a random variable, which we refer to as the *first hitting time*. More precisely, we denote by $t_i(v, f, \phi)$ the number of function evaluations that the variant v used in the i -th run before it evaluates a solution x satisfying $f(x) \leq \phi$ for the first time. If target ϕ is not hit, we define $t_i(v, f, \phi) = \infty$. To be consistent with previous work, such as [41], we decide to use two estimators of the mean of the hitting time distribution:

Definition 2.1. For a set of functions $\mathcal{F} = \{f^{(1)}, \dots, f^{(K)}\}$, the *average hitting time (AHT)* is defined as:

$$\tilde{T}(v, \mathcal{F}, \phi) = \frac{1}{nK} \sum_{i=1}^n \sum_{j=1}^K \min\{t_i(v, f^{(j)}, \phi), P\}$$

When a run does not succeed in hitting target ϕ , we have $t_i(v, f, \phi) = \infty$. In this case, a penalty $P \geq B$ (where B is the maximum budget) is applied. Usually, this penalty is set to ∞ , in which case this value is called AHT. Otherwise, it is commonly referred to as *penalized AHT*.

In contrast, the *expected running time (ERT)* equals

$$\text{ERT}(v, \mathcal{F}, \phi) = \frac{\sum_{i=1}^n \sum_{j=1}^K \min\{t_i(v, f^{(j)}, \phi), B\}}{\sum_{i=1}^n \sum_{j=1}^K \mathbb{1}\{t_i(v, f^{(j)}, \phi) < \infty\}}$$

Previous work has shown ERT, as opposed to AHT, to be a consistent, unbiased estimator of the mean of the distribution hitting times [3]. However, it is good to note that ERT and AHT are equivalent when all runs of variant v manage to hit target ϕ .

In the context of the modular CMA-ES, the CASH problem is adopted as follows. Given a set of CMA-ES variants V , a common hyperparameter space H , a set of function instances \mathcal{F} , and a target value ϕ , the CASH problem aims to find the combined algorithm and hyperparameter setting that solves the problem below:

$$v^*, \mathbf{h}^* = \underset{v \in V, \mathbf{h} \in H}{\operatorname{argmin}} \text{ERT}(v, \mathcal{F}, \phi).$$

Note here that we aim at finding the best (variant, hyperparameter)-pair for each of the 24 BBOB functions individually and we consider as \mathcal{F} the set of the first five instances of each function. We do not aggregate over different functions, since the benchmarks can easily be distinguished by exploratory landscape approaches [7].

2.4 Hyperparameter Tuning

In this work, we compare two different off-the-shelf tools for mixed-integer hyperparameter tuning: irace and MIP-EGO.

Irace [28, 29] is an algorithm designed for hyperparameter optimization, which implements an iterated racing procedure. irace is implemented in R and is freely available at [30]. For our experiments, we use the elitist version of irace with adaptive capping, which we briefly describe in the following.

irace works by first sampling a set of candidate parameter settings, which can be any combination of discrete, continuous, categorical, or ordinal variables. These parameters are empirically evaluated on some problem instances, which are randomly selected from the set of available instances. After running on x instances, a statistical test is performed to determine which parameter settings to discard. The remaining parameter settings are then run on more instances and continuously tested every ℓ iterations until either only a minimal number of candidates remain or until the budget of the current iteration is exhausted. The surviving candidates with the best average hitting times are selected as the elites.

After the racing procedure, new candidate parameter settings are generated by selecting a parent from the set of elites and “mutating” it, as described in detail in [28]. After generating the new set of candidates, a new race is started with these new solutions, combined with the elites. Since we use an elitist version of irace, these elites are not discarded until the competing candidates have been evaluated on the same instances which the elites have already seen. This is done to prevent the discarding of candidates which perform well on the previous race based on only a few instances in the current race.

Apart from using elitism and statistical tests to determine when to discard candidate solutions, we also use another recently developed extension of irace, the so-called *adaptive capping* [9] procedure. Adaptive capping helps to reduce the number of evaluations spent on candidates which will not manage to beat the current best. Adaptive capping enables irace to stop evaluating a candidate once it reaches a mean hitting time which is worse than the median of the elites, indicating that this candidate is unlikely to be better than the current best parameter settings.

Mixed-Integer Parallel Efficient Global Optimization (MIP-EGO) [44, 45] is a variant of Efficient Global Optimization (EGO, a sequential model-based optimization technique), which can deal with mixed-integer search-spaces. Because EGO is designed to deal with expensive function evaluations, and this variant has the ability to deal with continuous, discrete, and categorical parameters, it is also well suited to the hyperparameter tuning task. It uses a much different approach as irace, as we will describe in the following.

EGO works by initially sampling an set of solution candidates from some specified probability distribution, specifically a Latin hypercube sampling in MIP-EGO. Based on the evaluation of these initial points, a meta-model is constructed. Originally, this was done using Gaussian process regression, but MIP-EGO uses random forests to be able to deal with mixed-integer search spaces. Based on this model, a new point (or a set of points) is proposed according to some metric, called the *acquisition function*. This can be as simple as selecting the point with the largest *probability of improvement* (PI) or the largest *expected improvement* (EI). More recently, acquisition functions based on moment-generating function of the improvement have also been introduced [45]. For this paper, we use the basic EI acquisition function, which is maximized using a simple evolution strategy. After selecting the point(s) to evaluate, the meta-model is updated according to the quality of the solutions. The process is re-

peated until a termination criterion (budget constraint in our case) is met.

3 Baseline: Sequential methods

To establish a baseline of achievable performance of tuned CMA-ES variants, we propose a simple sequential approach of algorithm selection and hyperparameter tuning. Since the ERT for all variants on all benchmark functions is available, a complete enumeration technique would be the simplest form of algorithm selection. Then, based on the required robustness of the final solution, either one of several algorithm variants can be selected to undergo hyperparameter tuning. More precisely, we define two sequential methods as follows:

- **Naïve sequential:** Perform hyperparameter tuning (using MIP-EGO) on the one CMA-ES variant with the lowest ERT
- **Standard sequential:** Perform hyperparameter tuning (using MIP-EGO) on a set of 30 variants. We have chosen to consider the following set of variant in order to have a wide representation of module settings, and to be able to fairly compare the impact of hyperparameter tuning across functions:
 - The 10 variants with lowest ERT.
 - The 10 variants ranked 200-210 according to ERT.
 - 10 ‘common’ variants, i.e., CMA-ES variants previously studied in the literature (see Table V in [38]).

For both of these methods, the execution of MIP-EGO has a budget of 200 ERT-evaluations, each of which is based on 25 runs of the underlying CMA-ES variant (i.e., 5 runs per each of the five instances). Since the observed hitting times show high variance, we validate the ERT values by performing 250 additional runs (50 runs per each instance). All results shown will be ERT from these verification runs, unless stated otherwise. The variant selection and hyperparameter tuning is done separately for each function.

3.1 First Results

While the two sequential methods introduced are quite similar, it is obvious that the naïve one will always perform at most equally well as the standard version, since the algorithm variant tuned in the naïve approach is always included in the set of variants tuned by the standard method (the same tuned data is used for both methods to exclude impact of randomness). In general, the standard sequential method achieves ERTs which are on average around 20% lower than the naïve approach.

To better judge the quality of these sequential methods, we compare their performance to the default variant of the CMA-ES, which is the variant in which all modules are set to 0. This can be done based on ERT, for each function, but that does not always show the complete picture of the performance. Instead, the differences between the performances of the sequential method and the default CMA-ES are shown in a *Empirical Cumulative Distribution Function (ECDF)*, which aggregates all runs on all functions and shows the fraction of runs and targets which were hit within a certain amount of function evaluations. This is shown in Figure 1 (targets used available at [42]). From this, we see that the sequential approach completely dominates the default variant. When considering only the ERT, this improvement is on average 73%.

As well as comparing performance against the default CMA-ES, it can also be compared against the best modular variant with default hyperparameters, i.e. the result of pure algorithm selection. For




Figure 1: ECDF curve over all benchmark functions of both the **standard sequential** method as well as the default CMA-ES. Figure generated using IOHprofiler [10].

this, the standard sequential approach manages to achieve a 24.7% improvement in terms of average ERT, as opposed to 6.3% for the naïve version. Of note for the naïve version is that not all comparisons against the pure algorithms selection are positive, i.e. for some (5) functions it achieves a larger ERT. This might seem counter-intuitive, as one would expect hyperparameter tuning to only improve the performance of an algorithm. However, this is where the inherent variance of evolution strategies has a large impact. In short, because ERTs seen by MIP-EGO are based on only 25 runs, it may happen that a sub-optimal hyperparameter setting will be selected. This is explained in more detail in the following sections.

3.2 Pitfalls

The sequential methods described here have the advantage of being based on algorithm selection by complete enumeration. In theory, this would be the perfect way of selecting an algorithm variant. However, since CMA-ES are inherently stochastic, variance has a large effect on the ERT, and thus on the algorithm selection. This might not be an issue if one assumes that hyperparameter tuning has an equal impact on all CMA-ES variants. Unfortunately, this is not the case in practice.

3.2.1 Curse of High Variance

The inherent variance present in the ERT-measurements does not only cause potential issues for the algorithm selection, it also plays a large role in the hyperparameter configuration. As previously mentioned, the ERT after running MIP-EGO can be larger than the ERT with the default hyperparameters, even though the default hyperparameters are always included in the initial solution set explored by MIP-EGO. Since this might seem counter-intuitive, a small-scale experiment can be designed to show this phenomenon in more detail.

This experiment is set up by first taking the set of 50 hitting times for each instance as encountered in the verification runs. Then, sample x runs per instance from these hitting times and calculate the resulting ERT. Repeat this 10 times, and take the minimal ERT. Then we can compare the original ERT to this new value. This is similar to the internal data seen by MIP-EGO, if we assume that 5% of the variants it evaluated have a similar hitting time distribution. When performing this experiment on a set of 100 algorithm variants on F21,

we obtain the results as seen in Figure 2, which shows that the actual differences between ERTs given by MIP-EGO and those achieved in the verification runs matches the difference we would expect based on this experiment.

3.2.2 Differences in Improvement

Even when accounting for the impact of stochasticity on algorithm selection, there can still be large differences in the impact of hyperparameter tuning for different variants. This can be explained intuitively by the notion that some variants have hyperparameters which are already close to optimal for certain problems, while others have very poor hyperparameter settings. Hyperparameter tuning might then lead to some variants, which perform relatively poorly with default hyperparameters, to outperform all others when the hyperparameters have been sufficiently optimized.

This can be shown clearly by looking at one function in detail, F12 in this case, and studying the impact of hyperparameter tuning on two sets of algorithm variants. The two groups are selected as follows: the top 50 according to ERT, and a set of 50 variants. Then, for each of these variants, the hyperparameters are tuned using MIP-EGO. The resulting ERTs are shown in Figure 3. From this figure, it is clear that the relative improvements are indeed much larger for the group of random variants. There are even some variants which start with very poor ERT, which after tuning become competitive with the variants from the first group. In this first group the effects noted in

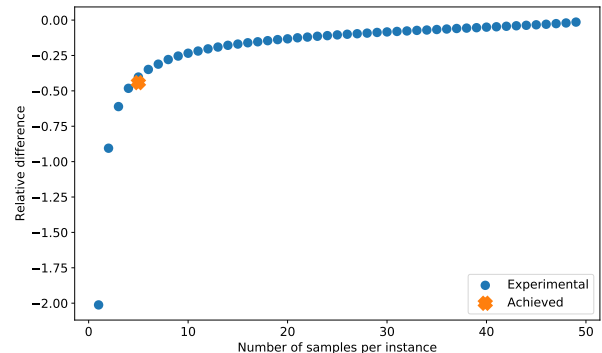


Figure 2: Average improvement of ERT obtained from 250 runs vs the value obtained after running MIP-EGO (25 runs) in orange, vs experimental improvement. Experimental improvement obtained over 100 repetitions of selecting k samples per instance for each variant and calculating their respective ERT.




Figure 3: ERT for two groups of CMA-ES variants. The ERT before tuning is shown in light color (based on 25 runs), while the ERT after tuning and verification is shown in a darker shade.

Section 3.2.1 are also clearly present, with some variants performing worse after tuning than before.




Figure 4: Evolution of ERT-based ranking (lower rank is better) of 100 algorithm variants on F12. Default refers to the ERT using the default hyperparameters while optimized is the best ERT using the tuned hyperparameters as found by MIP-EGO. Darker lines correspond to larger changes in ranking. Colors correspond to the grouping of variants.

We can also rank these CMA-ES variants based on their ERT, both with the default and tuned hyperparameters, both for the 25 runs seen during the tuning as the 250 verification runs. The resulting differences in ranking are then shown in Figure 4. This figure shows both the impact of variance on the 25-run rankings as the much larger differences present between the rankings with default versus tuned hyperparameters.

These differences in improvement after hyperparameter tuning are also highly dependent on the underlying test function. When executing the sequential approach mentioned previously, 30 variants are tuned for each function, and the ERTs are verified using 250 runs. The resulting data can then give some insight into the difference in terms of relative improvement possible per function, as is visualized in Figure 5. This shows that, on average, a relatively large performance improvement is possible for the selected variants. However, the distributions are have large variance, and differ greatly per function. This highlight the previous findings of different variants receiving much greater benefits from tuned hyperparameters than others, thus confirming the results from Figure 4.




Figure 5: Distribution of relative improvement in ERT between the default and tuned hyperparameters. For each function, 30 variants are tuned with MIP-EGO, and the resulting (variant, hyperparameters)-pairs are rerun 250 times to validate the results. The same is done for the default hyperparameters, and then the relative improvement in ERT is calculated.

3.2.3 Scalability

The final, and most important, issue with the sequential methods lies in their scalability. Because these methods rely on complete enumeration of all variants based on ERT, the required number of function evaluations grows as the algorithm space increases. If just a single new binary module is added, the size of this space doubles. This exponential growth is unsustainable for the sequential methods, especially if the testbed will also be expanded to include higher-dimensional functions (requiring more budget for the runs of the CMA-ES).

4 Integrated Methods

To tackle the issue of scalability, we propose a new way of combining algorithm selection and hyperparameter tuning. This is achieved by viewing the variant as part of the hyperparameter space, which is easily achieved by considering the module activations as hyperparameters. This leads to a mixed-integer search space, which both MIP-EGO and irace can easily adapt to. Thus, we will use two integrated approaches: MIP-EGO and irace. Both will get a total budget of 25,000 runs, which irace allocates dynamically while MIP-EGO allocates 25 runs to calculate ERTs for its solution candidates.

4.1 Case Study: F12




Figure 6: Comparison of ERTs achieved by the integrated approach using irace and the tuning a set of 56 tuned variants (using MIP-EGO).

The viability of this integrated approach can be established by looking at a single function and comparing the results from the integrated approach to the previously established baselines. This is done for F12, since for this function, data for the top 50 variants is available, as shown in Figure 3. We run irace 4 times on instance 1 of this function, and compare the result to those achieved by the best tuned variants. This is done in Figure 6. From this figure, it can be seen that two of the runs from irace are very competitive with the best tuned variants, while the other two still manage to outperform most variants with default hyperparameters. This shows that this integrated approach is quite promising, and worth to study in more detail.

4.2 Results

The results from running the integrated and sequential approaches on all 24 benchmark functions are shown in Figure 7. This figure shows

that, in general, the ERT achieved by irace and MIP-EGO is comparable. Irace has a slight advantage, beating MIP-EGO on 14 out of 24 functions. However, both methods still manage to outperform the naïve sequential approach while using significantly fewer runs, and are only slightly worse than the more robust version of the sequential approach. As expected, all methods manage to outperform pure algorithm selection quite significantly.




Figure 7: Relative ERTs against the best algorithm variant with default hyperparameters (targets chosen as in [41]) from running MIP-EGO and irace on the integrated selection and configuration space, as well as from the two sequential approaches. The ‘predicted’ relative ERT (based on 25 runs, with the exception of irace) is shown as a small black bar, whereas all other shown ERTs are based on 250 verification runs. y -axis cut at 1.5 (full data set available at [42]).

4.3 Comparison of MIP-EGO and Irace

From the results presented in Figure 7 it can be seen that the performance of the two integrated methods, MIP-EGO and irace, is quite similar. However, when introducing these methods, it was clear that their working principles differ significantly. To gain more understanding about how these results are achieved, three separate principles were studied: prediction error, balance between exploration and exploitation, and stability.

4.3.1 Prediction Error

The bars in Figure 7 seem to indicate that the prediction error for irace is smaller than the one for MIP-EGO. This is indeed the case: the average prediction error is 10.6% for irace, compared to 17.4% for MIP-EGO, suggesting that the AHT values reported by irace are more robust than the ERTs given by MIP-EGO. However, we also note that there exist some outliers, for which the prediction error of irace is relatively large (up to 35% for function 4). This happens because irace reports penalized AHT instead of ERT during the prediction-phase (see Definition 2.1). However, these prediction errors for irace can be positive (i.e. overestimating the real ERT), whereas MIP-EGO always underestimates the actual ERT.

4.3.2 Exploration-Exploitation Balance

While the prediction error is an important distinguishing factor between the two integrated methods, a much more important question to ask is how their search behaviour differs. This is best characterized by looking at the balance between exploration and exploitation,

which we analyze by looking at the complete set of evaluated candidate (variant, hyperparameter)-pairs, and noting how many unique variants were explored after the initialization phase. For MIP-EGO, this number is on average 565, while for irace it is only 112. This leads us to conclude that MIP-EGO is very exploitative in the algorithm space, while irace is much more focused on exploitation of the hyperparameters. On average, across all 24 benchmark functions, a fraction of 78.6% of all candidates evaluated by irace differ only in terms of the continuous hyperparameters, whereas only 2.6% of the evaluated (variant, hyperparameter) pairs contain unique variants. Even when including the initial random population, this value only increases to 9.7%, while MIP-EGO achieves an average fraction of 77.8% unique variants evaluated.

This difference in exploration-exploitation balance is expected to lead to a difference in variants found by irace and MIP-EGO, specifically in how these variants would rank with default hyperparameters. This is visualized in Figure 8. From this figure, the differences between irace and MIP-EGO are quite clear. While MIP-EGO usually has better ranked variants, the median ranking is only 108, as opposed to 428 for irace. This confirms the findings of Section 3.2.2, where we saw there can be quite large differences in ranking before and after hyperparameter tuning. However, we still find that a larger focus on exploration yields a selection of variants which are ranked better on average.




Figure 8: Original ranking (default hyperparameters) of the algorithm variant found by the integrated approaches.




Figure 9: Distributions of relative hitting times (divided by the maximal hitting time observed for the function) of 15 (variant, hyperparameters)-pairs, resulting from 15 independent runs of the integrated approaches, each of which is run 250 times on the corresponding benchmark function.

4.3.3 Stability

Finally, we study the variance in performance of the algorithm variants found by the two configurators. Since MIP-EGO is more exploitative, it might be more prone to variance than irace and thus less stable over multiple runs. To investigate this assumption, we select two benchmark functions and run both integrated methods 15 times. The resulting (variant, hyperparameters)-pairs are then rerun 250 times, the runtime distributions of which are shown in Figure 9. For F20, there is a relatively small difference between irace and MIP-EGO, slightly favoring irace. This indicates that the exploitation done by irace is indeed beneficial, leading to slightly lower hitting times. For F1, this effect is much larger, since for F1 most variants behave quite similarly, so the more benefit can be gained by tuning the continuous hyperparameters relative to exploring the algorithm space.

4.3.4 Summary

A summary of the differences between the four methods studied in this paper can be seen in Table 2. From this, we can see that the differences in terms of performance between the integrated and sequential methods is minimal, while they require a significantly lower budget. This budget value is in no way optimized, so an even lower budget than the one used in our study might achieve similar results. This might especially be true for irace, since it uses most of its budget to evaluate very small changes in hyperparameter values.

	Naïve Seq.	Seq.	MIP-EGO	irace
Best on # functions	0	9	9	6
Avg. Impr. over best modular CMA-ES	6.3%	24.7%	20.2%	20.7%
Avg. Impr. over default CMA-ES	67.4%	73.0%	72.9%	72.5%
Avg. Prediction Error	23.2%	18.8%	17.4%	10.6%
Budget (# function evaluations)/1,000	~ 120	150	25	25
% Unique CMA-ES variants explored	95.8%	76.8%	77.8%	9.7%

Table 2: Comparison of the four methods for determining (variant, hyperparameter)-pairs used in this paper. Seq.=sequential. Improvement over best modular CMA-ES refers to the relative improvement in ERT over the single best variant with default hyperparameters.

5 Conclusions and Future Work

We have studied several ways of combining *algorithm selection* and *algorithm configuration* of modular CMA-ES variants into a single integrated approach. We have shown that a sequential execution of brute-force algorithm selection and hyperparameter is sub-optimal because the large variance present in the observed ERTs. In addition, the sequential approaches require a large number of function evaluations, and quickly becomes prohibitive when new modules are added to the modEA framework. This clearly illustrates a need for efficient and robust combined algorithm selection and configuration (CASH) methods.

We have shown that both irace and MIP-EGO manage to solve the CASH problem for the modular CMA-ES. They outperform the results from the naïve sequential approach and show comparable performance to the more robust sequential method, and this at much smaller cost (up to a factor of 6 in terms of function evaluations).

We have also observed that, for the integrated approach, MIP-EGO has a heavy focus on exploring the algorithm space, while irace spends most of its budget on tuning the continuous hyperparameters of a single variant. These differences were shown to lead to a slight benefit for irace on the sphere-function, but in general the difference in performance was minimal across the benchmark functions. This indicates that there is still room for improvement by combining the best parts from both methods into a single approach. This could take advantage of the dynamic allocation of runs to instances and adaptive capping which irace uses, as well as the efficient generation of new candidate solutions using the working principles of efficient global optimization, as done in MIP-EGO.

Another way to improve the viability of the integrated approaches studied in this paper would be to tune their maximum budget, as this was set arbitrarily in our experiments, and might be reduced significantly without leading to a large loss in performance.

We have focused in this work on the three hyperparameters selected in [7]. A straightforward extension of our work would be to consider the configuration of additional hyperparameters — global ones that are present in all variants (such as the population size), but also conditional ones that appear only in some variants but not in others (i.e. the threshold value when the ‘threshold convergence’ module is turned on). While irace can deal with such conditional parameter spaces, MIP-EGO would have to be revised for this purpose.

Our long-term goal is to develop methods which adjust variant selection and configuration *online*, i.e., while optimizing the problem at hand. This could be achieved by building on exploratory landscape analysis [31] and using a landscape-aware selection mechanism. Relevant features could be local landscape features such as provided by the flacco software [24] (this is the approach taken in [7]), but also the state of the CMA-ES-parameters themselves, and approach suggested in [33]. We have analyzed the potential impact of such an online selection in [41]. Some initial work in determining how landscape features change during the search has been proposed in [21], but it was shown in [34] that some of the local features provided by flacco are not very robust, so that a suitable selection of features is needed for the use in a landscape-aware algorithm design.

Finally, we are interested in generalizing the integrated algorithm selection and configuration approach studied in this work to more general search spaces, and in particular to possibly much more unstructured algorithm selection problems. For example, one could consider to extend the CASH approach to the whole set of algorithms available in the BBOB repository (some of these are summarized in [18] but many more algorithms have been added in the last ten years since the writing of [18]). Note that it is an open question, though, how well the here-studied configuration tools irace and MIP-EGO would perform on such an unstructured, categorical algorithm selection space. Note also that here again we need to take care of conditional parameter spaces, since the algorithms in the BBOB data set have many different parameters that need to be set.

ACKNOWLEDGEMENTS

This work has been supported by the Paris Ile-de-France region, and by a public grant as part of the Investissement d’avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH.

REFERENCES

- [1] Martin Andersson, Sunith Bandaru, Amos HC Ng, and Anna Syberfeldt, ‘Parameter tuned CMA-ES on the CEC’15 expensive problems’, in *CEC*, pp. 1950–1957. IEEE, (2015).
- [2] Anne Auger, Dimo Brockhoff, and Nikolaus Hansen, ‘Mirrored Sampling in Evolution Strategies with Weighted Recombination’, in *GECCO*, pp. 861–868. ACM, (2011).
- [3] Anne Auger and Nikolaus Hansen, ‘A restart CMA evolution strategy with increasing population size’, in *CEC*, pp. 1769–1776, (2005).
- [4] Anne Auger, Mohamed Jebalia, and Olivier Teytaud, ‘Algorithms (x, sigma, eta): Quasi-random mutations for evolution strategies’, in *Artificial Evolution*, pp. 296–307. Springer, (2005).
- [5] Thomas Bäck, Christophe Foussette, and Peter Krause, *Contemporary Evolution Strategies*, Natural Computing Series, Springer, 2013.
- [6] Thomas Bartz-Beielstein, ‘SPOT: an R package for automatic and interactive tuning of optimization algorithms by sequential parameter optimization’, *CoRR*, abs/1006.4645, (2010).
- [7] Nacim Belkhir, Johann Dréo, Pierre Savéant, and Marc Schoenauer, ‘Per instance algorithm configuration of CMA-ES with limited budget’, in *GECCO*, pp. 681–688. ACM, (2017).
- [8] Dimo Brockhoff, Anne Auger, Nikolaus Hansen, Dirk V. Arnold, and Tim Hohm, ‘Mirrored Sampling and Sequential Selection for Evolution Strategies’, in *PPSN*, pp. 11–21. Springer, (September 2010).
- [9] Leslie Pérez Cáceres, Manuel López-Ibáñez, Holger Hoos, and Thomas Stützle, ‘An experimental study of adaptive capping in irace’, in *LION*, eds., Roberto Battiti, Dmitri E. Kvasov, and Yaroslav D. Sergeyev, pp. 235–250, (2017).
- [10] Carola Doerr, Hao Wang, Furong Ye, Sander van Rijn, and Thomas Bäck, ‘IOHprofiler: A Benchmarking and Profiling Tool for Iterative Optimization Heuristics’, *arXiv e-prints:1810.05281*, (October 2018).
- [11] Katharina Eggensperger, Marius Lindauer, and Frank Hutter, ‘Pitfalls and best practices in algorithm configuration’, *J. Artif. Intell. Res.*, **64**, 861–893, (2019).
- [12] Stefan Falkner, Aaron Klein, and Frank Hutter, ‘BOHB: Robust and efficient hyperparameter optimization at scale’, in *ICML*, pp. 1436–1445, (2018).
- [13] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter, ‘Efficient and robust automated machine learning’, in *Advances in neural information processing systems*, pp. 2962–2970, (2015).
- [14] Xin Guo, Bas van Stein, and Thomas Bäck, ‘A new approach towards the combined algorithm selection and hyper-parameter optimization problem’, (2020), to appear at SSCI 2020.
- [15] Nikolaus Hansen, ‘CMA-ES with Two-Point Step-Size Adaptation’, *arXiv:0805.0231 [cs]*, (May 2008). arXiv: 0805.0231.
- [16] Nikolaus Hansen, ‘Benchmarking a BI-population CMA-ES on the BBOB-2009 Function Testbed’, in *GECCO Companion*, pp. 2389–2396. ACM, (2009).
- [17] Nikolaus Hansen, Anne Auger, Dimo Brockhoff, Dejan Tuar, and Tea Tuar, ‘COCO: Performance Assessment’, *arXiv:1605.03560 [cs]*, (May 2016). arXiv: 1605.03560.
- [18] Nikolaus Hansen, Anne Auger, Raymond Ros, Steffen Finck, and Petr Pošík, ‘Comparing results of 31 algorithms from the black-box optimization benchmarking BBOB-2009’, in *GECCO*, pp. 1689–1696. ACM, (2010).
- [19] Nikolaus Hansen and Andreas Ostermeier, ‘Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation’, in *CEC*, pp. 312–317, (May 1996).
- [20] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown, ‘Sequential model-based optimization for general algorithm configuration’, in *LION*, pp. 507–523. Springer, (2011).
- [21] Anja Janković and Carola Doerr, ‘Adaptive landscape analysis’, in *GECCO Companion*, pp. 2032–2035. ACM, (2019).
- [22] Grahame A. Jastrebski and Dirk V. Arnold, ‘Improving Evolution Strategies through Active Covariance Matrix Adaptation’, in *CEC*, pp. 2814–2821, (2006).
- [23] Pascal Kerschke, Holger H. Hoos, Frank Neumann, and Heike Trautmann, ‘Automated algorithm selection: Survey and perspectives’, *Evolutionary Computation*, **27**(1), 3–45, (2019).
- [24] Pascal Kerschke and Heike Trautmann, ‘The r-package FLACCO for exploratory landscape analysis with applications to multi-objective optimization problems’, in *CEC*, pp. 5262–5269. IEEE, (2016).
- [25] Lars Kotthoff, Chris Thornton, Holger H Hoos, Frank Hutter, and Kevin Leyton-Brown, ‘Auto-weka 2.0: Automatic model selection and hyperparameter optimization in weka’, *The Journal of Machine Learning Research*, **18**(1), 826–830, (2017).
- [26] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar, ‘Hyperband: A novel bandit-based approach to hyperparameter optimization’, *arXiv preprint arXiv:1603.06560*, (2016).
- [27] *Parameter Setting in Evolutionary Algorithms*, eds., Fernando G. Lobo, Cláudio F. Lima, and Zbigniew Michalewicz, volume 54 of *Studies in Computational Intelligence*, Springer, 2007.
- [28] Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Leslie Pérez Cáceres, Mauro Birattari, and Thomas Stützle, ‘The irace package: Iterated racing for automatic algorithm configuration’, *Operations Research Perspectives*, **3**, 43 – 58, (2016).
- [29] Manuel López-Ibáñez, Jérémie Dubois-Lacoste, Thomas Stützle, and Mauro Birattari, ‘The irace package, iterated race for automatic algorithm configuration’, Technical Report TR/IRIDIA/2011-004, IRIDIA, Université Libre de Bruxelles, Belgium, (2011).
- [30] Manuel López-Ibáñez and Leslie Pérez Cáceres. The irace package: Iterated race for automatic algorithm configuration. <http://iridia.ulb.ac.be/irace/>.
- [31] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph, ‘Exploratory landscape analysis’, in *GECCO*, pp. 829–836. ACM, (2011).
- [32] Alejandro Píad-Morffis, Suilan Estévez-Velarde, Antonio Bolufé-Röhler, James Montgomery, and Stephen Chen, ‘Evolution strategies with threshold convergence’, in *CEC*, pp. 2097–2104, (May 2015).
- [33] Zbynek Pitra, Jakub Repický, and Martin Holena, ‘Landscape analysis of gaussian process surrogates for the covariance matrix adaptation evolution strategy’, in *GECCO*, pp. 691–699. ACM, (2019).
- [34] Quentin Renau, Johann Dreö, Carola Doerr, and Benjamin Doerr, ‘Expressiveness and robustness of landscape features’, in *GECCO Companion*, pp. 2048–2051. ACM, (2019).
- [35] Chris Thornton, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown, ‘Auto-weka: Combined selection and hyperparameter optimization of classification algorithms’, in *ACM SIGKDD*, pp. 847–855. ACM, (2013).
- [36] Sander van Rijn. Modular CMA-ES framework from [38], v0.3.0. <https://github.com/sjvrijn/ModEA>. Available also as pypi package at <https://pypi.org/project/ModEA/0.3.0/>, 2018.
- [37] Sander van Rijn, Carola Doerr, and Thomas Bäck, ‘Towards an adaptive cma-es configurator’, in *PPSN*, volume 11101 of *Lecture Notes in Computer Science*, pp. 54–65. Springer, (2018).
- [38] Sander van Rijn, Hao Wang, Matthijs van Leeuwen, and Thomas Bäck, ‘Evolving the structure of Evolution Strategies’, in *SSCI*, pp. 1–8, (2016).
- [39] Sander van Rijn, Hao Wang, Bas van Stein, and Thomas Bäck, ‘Algorithm Configuration Data Mining for CMA Evolution Strategies’, in *GECCO*, pp. 737–744. ACM, (2017).
- [40] Diederick Vermetten, Sander van Rijn, Thomas Bäck, and Carola Doerr, ‘Github repository for online selection of CMA-ES variants’. https://github.com/Dvermetten/Online_CMA-ES_Selection.
- [41] Diederick Vermetten, Sander van Rijn, Thomas Bäck, and Carola Doerr, ‘Online selection of CMA-ES variants’, in *GECCO*, pp. 951–959. ACM, (2019).
- [42] Diederick Vermetten, Hao Wang, Thomas Bäck, and Carola Doerr, ‘Github repository with the data used in this paper’. https://github.com/Dvermetten/CMA_ES_hyperparam.
- [43] Hao Wang, Michael Emmerich, and Thomas Bäck, ‘Mirrored Orthogonal Sampling with Pairwise Selection in Evolution Strategies’, in *SAC*, pp. 154–156. ACM, (2014).
- [44] Hao Wang, Michael Emmerich, and Thomas Bäck, ‘Cooling strategies for the moment-generating function in bayesian global optimization’, in *CEC*, pp. 1–8, (2018).
- [45] Hao Wang, Bas van Stein, Michael Emmerich, and Thomas Bäck, ‘A new acquisition function for bayesian optimization based on the moment-generating function’, in *SMC*, pp. 507–512. IEEE, (2017).
- [46] David H. Wolpert and William G. Macready, ‘No free lunch theorems for optimization’, *IEEE Trans. Evolutionary Computation*, **1**(1), 67–82, (1997).