



HAL
open science

Improving Nevergrad's Algorithm Selection Wizard NGOpt through Automated Algorithm Configuration

Risto Trajanov, Ana Nikolikj, Gjorgjina Cenikj, Fabien Teytaud, Mathurin Videau, Olivier Teytaud, Tome Eftimov, Manuel López-Ibáñez, Carola Doerr

► **To cite this version:**

Risto Trajanov, Ana Nikolikj, Gjorgjina Cenikj, Fabien Teytaud, Mathurin Videau, et al.. Improving Nevergrad's Algorithm Selection Wizard NGOpt through Automated Algorithm Configuration. 17th Proceedings of Parallel Problem Solving from Nature - (PPSN) 2022, Sep 2022, Dortmund, Germany. pp.18-31, 10.1007/978-3-031-14714-2_2 . hal-03740762

HAL Id: hal-03740762







<https://hal.sorbonne-universite.fr/hal-03740762>

Submitted on 29 Jul 2022

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

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

Improving Nevergrad’s Algorithm Selection Wizard NGOpt through Automated Algorithm Configuration

Risto Trajanov¹, Ana Nikolikj^{2,3}, Gjorgjina Cenikj^{2,3}, Fabien Teytaud⁴,
Mathurin Videau⁵, Olivier Teytaud⁵, Tome Eftimov²
, Manuel López-Ibáñez⁶, and Carola Doerr⁷

¹ Faculty of Computer Science and Engineering, Ss. Cyril and Methodius Skopje, North Macedonia

² Jožef Stefan Institute, Ljubljana, Slovenia

³ Jožef Stefan International Postgraduate School, Ljubljana, Slovenia

⁴ Université Littoral Côte d’Opale

⁵ Facebook AI Research

⁶ ITIS Software, Universidad de Málaga, Spain

⁷ Sorbonne Université, CNRS, LIP6, Paris, France

Abstract. Algorithm selection wizards are effective and versatile tools that automatically select an optimization algorithm given high-level information about the problem and available computational resources, such as number and type of decision variables, maximal number of evaluations, possibility to parallelize evaluations, etc. State-of-the-art algorithm selection wizards are complex and difficult to improve. We propose in this work the use of automated configuration methods for improving their performance by finding better configurations of the algorithms that compose them. In particular, we use elitist iterated racing (irace) to find CMA configurations for specific artificial benchmarks that replace the hand-crafted CMA configurations currently used in the NGOpt wizard provided by the Nevergrad platform. We discuss in detail the setup of irace for the purpose of generating configurations that work well over the diverse set of problem instances within each benchmark. Our approach improves the performance of the NGOpt wizard, even on benchmark suites that were not part of the tuning by irace.

Keywords: Algorithm Configuration · Algorithm Selection · Black-box Optimization · Evolutionary Computation

1 Introduction

In the context of black-box optimization, the use of a portfolio of optimization algorithms [23], from which an algorithm is selected depending on the features of the particular problem instance to be solved, is becoming increasingly popular. The algorithm that encapsulates the selection rules and the portfolio of

algorithms is sometimes referred as a “wizard” [20]. Algorithm selection wizards provide versatile, robust and convenient tools, particularly for black-box optimization, where problems that arise in practice show a large variety of different requirements, both with respect to problem models, but also with respect to the (computational) resources that are available to solve them.

Building a competitive algorithm selection wizard is not an easy task, because it not only requires devising the rules for selecting an algorithm for a given problem instance, but also configuring the parameters of the algorithms to be selected, which is a difficult task by itself [3]. Although there are examples of wizards that were build automatically for SAT problems [29], many algorithm selection wizards are still hand-crafted. An example of a successful hand-crafted algorithm selection wizard is NGOpt [20], provided by the optimization platform Nevergrad [22].

NGOpt was designed by carefully studying the performance of tens of optimizers on a wide range of benchmark suites to design hand-crafted rules that aim to select the best optimizer for particular problem features. Through an iterative improvement process, new versions of NGOpt are designed by refining the hand-crafted rules and replacing optimizers by others that lead to a better performance of the new NGOpt version. The result is a complex algorithm selection wizard that outperforms many well-known stand-alone optimizers on a wide range of benchmark suites.

In this work, we attempt to investigate the potential of improving NGOpt via automated algorithm design techniques. Rebuilding NGOpt from scratch using automated methods is a daunting task that would waste the knowledge already encoded in it and the human effort already invested in its creation. Instead, we examine how automatic algorithm configuration (AC) methods may be used in a judicious manner to improve NGOpt by focusing on improving one of its most critical underlying optimizers.

AC methods such as irace [16], aim to find a configuration, i.e., a setting of the parameters of an algorithm, that gives good expected performance over a large space of problem instances by evaluating configurations on a *training* subset from such space. AC methods, including irace, are typically designed to handle categorical and numerical parameter spaces and stochastic performance metrics, not only due to the inherent stochasticity of randomized optimization algorithms but also because the training instances represent a random sample of the problems of interest. Traditionally, AC methods have been used to tune the parameters of specific algorithms [3], however, more recently, they have been used to automatically design algorithms from flexible frameworks [1, 12, 18, 19] and to build algorithm portfolios [28]. For more background on AC, we refer the interested reader to two recent surveys [8, 25].

As mentioned above, we do not wish to re-build nor replace NGOpt but, instead, iteratively improve it. To do so, we first answer the question of whether the hand-crafted NGOpt may be further improved by replacing some of its components by algorithm configurations obtained from the application of an AC method, in particular, irace. When applied to a complex algorithm selection

wizard such as NGOpt and a diverse black-box optimization scenario, such as the one tackled by NGOpt, the correct setup of an AC method is crucial. We discuss in detail our approach here, which we believe is applicable to other similar scenarios. Our results show that our proposed approach clearly improves the hand-crafted NGOpt, even on benchmark suites not considered in the tuning process by irace.

This paper is structured as follows. Section 2 gives an introduction to the concept of algorithm selection wizard and describes, in particular, the one used in our experiments, NGOpt. Section 3 presents the experimental setup and the benchmark suites used for tuning the target algorithm. In Section 4 we present the integration of the tuned algorithms into the new algorithm selection wizard. The evaluation results are presented in Section 5. Section 6 concludes the paper.

2 Preliminaries

Algorithm Selection Wizards. Modern algorithm selection wizards make use of a number of different criteria to choose which algorithm(s) are executed on a given problem instance, and for how long. Algorithm selection wizards take into account a priori information about the problem and the resources that are available to solve it. Using this information, the algorithm selection wizards recommend one or several algorithms to be run, along with a protocol to assign computational budget to each of these.

NGOpt is built atop of several dozens of state-of-the-art algorithms that have been added to Nevergrad over the last years. However, NGOpt does not only select which algorithms to execute on which problem instances, but it also combines algorithms in several ways, e.g., by enriching classic approaches with local surrogate models. Since all algorithms submitted to Nevergrad are periodically run on all fitting benchmark problems available within the environment, a very large amount of performance data is available. In light of this rich data set, and in light of the tremendous performance gains obtained through automated algorithm designs [2, 12, 21] and selection rules [11], it is therefore surprising that both the decision rules of NGOpt and the configuration of the algorithms available in Nevergrad are hand-picked.

CMA. Among the algorithms that compose NGOpt, a key component is CMA, an instance of the family of covariance matrix adaptation evolution strategies (CMA-ES [7]). In Nevergrad’s CMA implementation, the following parameters are explicitly exposed, making them straightforward candidates for the configuration task. The `scale` parameter controls the scale of the search of the algorithm in the search domain: a small value leads to start close to the center. The `elitist` parameter is a Boolean variable that can have values ‘True’ or ‘False’ and it represents a switch for elitist mode, i.e., preserving the best individual even if it includes a parent. The `diagonal` parameter is another Boolean setting that controls the use of diagonal CMA [24]. Finally, the population size is another crucial parameter that is usually set according to the dimension of the problem. Instead of setting its value directly, we decided to create

a higher-level integer parameter `popsiz_factor` and let the population size be $\lfloor 4 + \text{popsiz_factor} \cdot \log d \rfloor$, where d is the problem dimension. The parameter search space of CMA is shown in Table 3.

Automated Algorithm Configuration. Given a description of the parameter space of a target algorithm, a class of problems of interest, and a computational budget (typically measured in target algorithm runs), the goal of an AC method is to find a parameter configuration of the target algorithm that is expected to perform well on the problem class of interest. Due to the stochasticity of randomized algorithms and the fact that only a limited sample of (*training*) problem instances can be evaluated in practice, specialized methods for AC have been developed in recent years [8, 9, 25] that try to avoid inherent human biases and limitations in manual parameter tuning. We have selected here the *elitist iterated racing* algorithm implemented by the irace package [16], since it has shown good performance for configuring black-box optimization algorithms in similar scenarios [13, 14, 15, 17].

3 Experimental Setup

To show the influence of the parameter tuning of the CMA included in NGOpt [22], we have performed several parameter tuning experiments using irace. The setup and results of this process are described in this section.

3.1 Setup of Irace for Tuning CMA

A training *instance* is defined by a benchmark function, a dimension, a rotation and the budget available to CMA. We also define “blocks” of instances: all instances within a block are equal except for the benchmark function and there are as many instances within a block as benchmark functions. We setup irace so that, within each race, the first elimination test (*FirstTest*) happens after seeing 5 blocks and subsequent elimination tests (*EachTest*) happen after each block. Moreover, configurations are evaluated by irace on blocks in order of increasing budget first and increasing dimension second, such that we can quickly discard poor-performing configurations on small budgets and only good configurations are evaluated on large ones [26]. The performance criterion optimized by irace is the objective value of the point recommended by CMA after it has exhausted its budget. Since Nevergrad validates performance according to the mean loss (as explained later), the elimination test used by irace is set to t-test. Finally, we set a maximum of 10 000 individual runs of CMA as the termination criterion of each irace run. By parallelizing each irace run across 4 CPUs, the runtime of a single run of irace was around 8 hours.

3.2 Benchmark Suites Used for Tuning

Nevergrad contains a very large number of benchmark *suites*. Each suite is a collection of benchmark problems that share some common properties such as

Table 1. Artificial benchmarks used as training instances by irace. The third column gives the name of the CMA variant obtained by irace on that particular benchmark (for the parameter values, see Table 4). The detailed code can be found in [22].

Benchmark name	Context	Optimized CMA
YABBOB	dimension $\in [2, 50]$, budget $\in [50, 12800]$	CMAstd
YASMALLBBOB	Budget < 50	CMAsmall
YATUNINGBBOB	Budget < 50 and dim ≤ 15	CMAtuning
YAPARABBOB	Num-workers=100	CMApara
YABOUNDEDBBOB	Box-constrained, budget ≤ 300 , dim ≤ 40	CMAbounded

similar domains or constraints, similar computational budgets, etc. A large number of different academic and real-world problems are available. We selected five suites (YABBOB, YASMALLBBOB, YATUNINGBBOB, YAPARABBOB, and YABOUNDEDBBOB) as training instances for irace, all of them derived from the BBOB benchmark suite [5] of the COCO benchmarking environment [6]. See Table 1 for their various characteristics. We ran irace once for each benchmark suite separately, by setting up their problem instances as described above. All runs of irace used the parameter space of CMA shown in Table 3. For each benchmark suite, we selected one CMA configuration from the ones returned by irace by doing 10 validation runs using the best configurations obtained by irace and the default CMA configuration. Each run consisted of running each CMA configuration on the whole instance space. For each instance we declared a winner configuration that yielded the smallest result. Using majority vote then we determined a winner from the 10 validation runs. The optimized configurations selected are named CMAstd for YABBOB, CMAsmall for YASMALLBBOB, CMAtuning for YATUNINGBBOB, CMApara for YAPARABBOB, and CMAbounded for YABOUNDEDBBOB (Table 1). Their parameter values are shown in Table 4.

Table 2. Artificial and real-world benchmarks used for testing. The detailed code can be found in [22].

Type	Name	Context
Artificial	YABIGBBOB	Budget 40000 to 320000
	YABOXBBOB	Box constrained
	YAHDBBOB	Dimension 100 to 3000
	YAWIDEBBOB	Different settings (multi-objective, noisy, discrete...)
	Deceptive	Hard benchmark, far from YABBOB
Real-world	SeqMLTuning	Hyperparameter tuning for SciKit models
	SimpleTSP	Traveling Salesman Problem, black-box
	ComplexTSP	Traveling Salesman Problem, black-box, nonlinear terms
	UnitCommitment	Unit Commitment for power systems
	Photonics	Simulators of nano-scale structural engineering
	GP 007	Gym environments used in [27]. Game simulator for the 007 game

Algorithm 1 Pseudocode for MetaCMA.

```

1: procedure METACMA(budget, dimension, num_workers, fully_bounded)
2:   if fully_bounded then return CMAbounded    ▷ All variables have bounded
   domain.
3:   if budget < 50 then
4:     if dimension ≤ 15 then return CMA tuning else return CMA small
5:   if num_workers > 20 then return CMA para else return CMA std

```

Table 3. Search space for the tuning of CMA by irace.

Parameter	Default value	Domain
scale	1	$(0, 1) \subset \mathbb{R}$
popsize-factor	3	$[1, 9] \subset \mathbb{N}$
elitist	False	{True, False}
diagonal	False	{True, False}

4 MetaCMA and its Integration in Nevergard

The CMA configurations found by irace were further combined into a new algorithm selection wizard. For this purpose, we propose MetaCMA (Algorithm 1), a deterministic ensemble for noiseless single-objective continuous optimization that, based on deterministic rules, switches between different tuned CMA configurations. For problems where all variables have bounded domain, MetaCMA selects CMAbounded. Otherwise, for low budget (function evaluations) available, the MetaCMA model selects either the CMA tuning configuration (for low dimension) or the CMA small configuration (otherwise). If number of workers is more than 20, then the MetaCMA uses the CMA para. If neither of the above-mentioned rules is met, the MetaCMA will switch to CMA std, which was tuned on YABBOB.

Given that NGOpt considers many more cases than MetaCMA, we have to integrate it inside NGOpt so that we have both the improved performance of our MetaCMA and the generality of NGOpt. NGTuned corresponds to NGOpt with all instances of CMA in NGOpt replaced by MetaCMA.

5 Experimental Results

To evaluate the performance of the MetaCMA model and its integration in NGOpt (NGTuned), we compared it with the previous variant of NGOpt and other state-of-the-art algorithms in three different scenarios. The first scenario consists of the benchmark suites that were used by irace during tuning. All algorithms have been rerun independently of the tuning process. The second scenario involves artificial benchmark suites whose problem instances were not

Table 4. Best parameters found by irace for each benchmark suite.

Parameter	CMAsstd (Yabbob)	CMAsmall (Yasmallbbob)	CMAtuning (Yatuningbbob)	CMAPara (Yaparabbob)	CMABounded (Yaboundedbbob)
scale	0.3607	0.4151	0.4847	0.8905	1.5884
popsize-factor	3	9	1	8	1
elitist	False	False	True	True	True
diagonal	False	False	False	True	True




Fig. 1. Results on YABBOB: suite used in the tuning, moderate dimension and budget.




Fig. 2. Results on YASMALLBBOB: benchmark suite used in the tuning, low budget.




Fig. 3. Results on YATUNINGBBOB: benchmark suite used in the tuning, counterpart of YABBOB with low budget and low dimension.

used in the tuning process. The third scenario involves real-world problems not presented in the tuning process of irace. Details about these benchmark suites are given in Table 1.

The selection of algorithms involved in the evaluation of each benchmark suite differ because each suite has its own set of state-of-the-art algorithms.




Fig. 4. Results on YAPARABBOB: benchmark suite used in the tuning, parallel case.




Fig. 5. YABOUNDEDBBOB: suite used in the tuning, bounded domain.

5.1 Evaluation on Artificial Benchmark Suites Presented in the Tuning Process of Irace

Figures 1, 2, 3, 4, and 5 present results on benchmark suites used for running irace. On the left side of each figure is a robustness heatmap, where the rows and columns represent different algorithms, and the values of the heatmap reflect the number of times the algorithm in each row outperforms the algorithm in each column. We use mean quantiles over multiple settings (compared to other optimization methods) in an anytime manner, meaning that all budgets are taken into account, not only the maximum one. This is robust to arbitrary scalings of objective functions. The algorithms’ global score, defined as their mean scores (the number of times it performs better than the other algorithms), is used for ranking the rows and columns. The six best-performing algorithms are listed as rows and sorted by their mean score, while the columns contain all of the algorithms involved in the algorithm portfolio. The labels in the columns contain the mean score of each algorithm. For instance, in Figure 1, the label for NGTuned is “1/18:76.1% +- 0.4”, meaning that it was ranked first, out of 18 evaluated algorithms, and its mean score is 76.1%, which is the mean number of times it performed better than the rest of the algorithms, on all problems and all budgets in the benchmark suite. On the right side of each figure, we present the algorithms’ performance with different budgets. The x-axis represents the budget, while the y-axis represents the rescaled losses (normalized objective value, averaged over all settings), for a given budget, and linearly normalized in the range [0, 1]. Apart from the algorithm name, the label of each algorithm in the legend contains two values in brackets. The first one denotes the mean loss achieved by the algorithm

for the maximum budget (after the losses for each problem have been normalized in the range $[0, 1]$), while the second one denotes the mean loss for the second largest budget. Figure 1 features the results on the YABBOB benchmark suite. The heatmap (left) shows that NGTuned outperforms the other algorithms, followed by MetaCMA, while NGOpt is in the 7th position. Figure 2 features the results on the YASMALBBOB benchmark suite. The best results are obtained by NGTuned, followed by MetaCMA, while NGOpt is ranked in the 6th position. In addition, with regard to different budgets, MetaCMA outperforms NGOpt by obtaining lower losses. In Figure 3, the results from the evaluation on YATUNINGBBOB benchmark suite are presented. It follows the trend of NGTuned being the best performing algorithm. Comparing NGTuned and NGOpt according to the convergence it seems that till 2×10^1 , both achieved similar loss, but later on for budgets above 2×10^1 , NGOpt can provide better losses. The comparison between the left and the right plot (keeping the rescaling in mind) suggests that NGTuned has a better rank than NGOpt, but that the rescaled loss is better for NGOpt. The results for YAPARABBOB benchmark suite are presented in Figure 4. In terms of mean scores obtained across all budgets, MetaCMA is the best algorithm, followed by NGTuned. In terms of the loss achieved for different budgets, NGTuned and MetaCMA provide similar results, with the biggest difference being in the largest budget, where NGTuned achieves a lower loss. Due to rescalings and the difference criteria (ranks on the left, normalized loss on the right) the visual comparison between algorithms can differ. Both outperform NGOpt. On the YABOUNDEDBBOB benchmark suite (Figure 5), NGOpt achieves better performance than NGTuned. Here, several algorithms from [10] have been tested: the low budget makes them relevant. The best results are achieved by NLOPT_LN_BOBYQA [4], closely followed by MetaCMA, with a difference of 0.4% in the average quantiles referring to the number of times they outperform other methods. From the line plot on the right of the figure, we can see that MetaCMA consistently provides the lowest normalized loss, for different budgets.

5.2 Evaluation on Benchmark Suites Not Used for Tuning

Artificial Benchmarks. Figures 6, 7, 9, 8, and 10 present results on benchmark suites related to YABBOB though with different configurations. The problem instances used were not used in the tuning process.

Figure 6 shows the results on the YABIGBBOB benchmark suite, which consists of the same problem instances as the YABBOB benchmark suite with longer budgets. Here, NGOpt achieves the best performance, while NGTuned is second ranked. From the line plot on the right of Figure 6, we can see that NGOpt and NGTuned are better than the other algorithms over all budgets. The results for the YABOXBBOB benchmark suite (box-constrained and low budget) are presented in Figure 7. In this case, all of the proposed algorithms appear as top ranked, with MetaCMA being the best performing one. As for budget up to around 10^3 , MetaCMA performs the best, and is later outperformed by NGTuned, which uses the MetaModel on top of MetaCMA. The YAHDBBOB bench-




Fig. 6. Results on YABIGBBOB: longer budgets than for BBOB. Not used in the irace trainings, but same functions as YABBOB.




Fig. 7. Results on YABOBBOB: box-constrained and low budget: not used in the irace training but has similarities.




Fig. 8. YAHDBBOB: high-dim. counterpart of YABBOB, not used for training.

mark suite is a high-dimensional counter part of YABBOB. On it, NGTuned is selected as the best algorithm in Figure 8, followed by NGOpt with 0.7% difference referring to the % of times they outperform other algorithms. The right subplot shows that these algorithms together with MetaCMA achieved the best losses, across all budgets.

Figure 9 shows results on the YAWIDEBBOB benchmark suite, which is considered to be specially difficult and covers many different test problem instances (continuous, discrete, noisy, noise-free, parallel or not, multi-objective or not). On YAWIDEBBOB, NGTuned is selected as the best algorithm, followed by NGOpt, with a difference of 1.7% in the average score related to the number of times they outperform the other algorithms. MetaCMA is ranked as third. From the line plot we can see that these algorithms achieved lowest losses, however the behavior is not stable over different budgets. We need to point out here that




Fig. 9. Results on YAWIDEBBOB: this benchmark matters particularly as it covers many different test cases, continuous, discrete, noisy, noise-free, parallel or not, multi-objective or not. It is designed for validating algorithm selection wizards. The different problems have different budgets, so that it is not surprising that the curve is not decreasing.




Fig. 10. Deceptive functions: this benchmark [22] is quite orthogonal to the YA*BBOB benchmarks, and very hard (infinitely many local minima arbitrarily close to the optimum / infinite condition number / infinitely tiny path to the optimum) .

NGTuned is not using CMA for multi-objective problems, so that some parts of YAWIDEBBOB are not modified by our work. The results indicate that the automated configuration that is performed for such benchmark suites also helps to improve the performance of the NGOpt on benchmark suites that contain a mix of different problem instances. The results for the Deceptive functions benchmark suite are presented in Figure 10, where NGTuned is the winning algorithm in terms of the number of times it performs better than the other algorithms and the mean loss achieved for different budgets over all settings. This benchmark suite is quite orthogonal to the YABBOB variants and very hard to optimize, since the deceptive functions contain infinitely many local minima arbitrarily close to the optimum or have an infinite condition number or have an infinite tiny path to the optimum.

Real-world Problems. We also tested our algorithms on some of the real-world benchmark suites presented in Nevergrad. Due to length constraints, the figures displaying these results are available only in the full version of this paper. On the SimpleTsp benchmark suite, NGTuned outperforms NGOpt in terms of the average frequency of winning over other algorithms and mean loss achieved for different budgets over all settings. On the SeqMITuning benchmark NG-

Tuned appears among the best-performing algorithms with rank six. It outperforms NGOpt and MetaCMA which do not appear in the set of best performing algorithms. On the ComplexTSP problem, NGTuned is found to be the best performing algorithm. It also surpasses NGOpt. With regard to mean performance per budget, they are very similar and consistently better than the other algorithms. On the Unit Commitment benchmark suite, NGTuned is the third ranked algorithm. It outperforms NGOpt but not enough for competing with the best variants of DE. According to the convergence analysis, it seems that NGTuned and NGOpt have similar behavior. On the Photonics problem, NGOpt and NGTuned appear between the top ranked algorithms showing very similar performance. On the benchmark suite (GP), which is a challenge in genetic programming, NGOpt and NGTuned show almost identical but very bad performance overall: it is mentioned in [22] that NGOpt has counterparts dedicated to neural reinforcement learning (MixDeterministic RL and ProgNoisyRL) that perform better: we confirm this. On the 007 benchmark suite, NGTuned is the best performing algorithm. Regarding the mean loss achieved for different budgets, NGTuned shows similar performance to NGOpt.

6 Conclusion

We have shown in this paper how automated algorithm configuration methods like irace may be used to improve an algorithm selection wizard like NGOpt. NGTuned was better than NGOpt in most benchmark suites, included those not used for tuning our MetaCMA. Our results improve the algorithm selection wizard not only for the benchmark suites used during the tuning process but also for a wider set of benchmark suites, including real-world test cases. Given the complexity of NGOpt, we have an impact on few cases, so that the overall impact on NGOpt across all Nevergrad suites remains moderate: in some cases, NGTuned is equivalent to NGOpt, and in some cases (in which CMA is crucial) it performs clearly better. In some cases, the differences between the two algorithms are small. If we consider only clear differences, we have gaps for YABBOB, 007, YASMALLBBOB, YAPARABBOB, YABOXBBOB, SimpleTSP, and Deceptive: in all these cases, NGTuned performs better than NGOpt.

Our tuning process still contains a few manual steps that require an expert, such as deciding which benchmark suites to use for tuning, as well as the integration of the tuned CMA configurations into NGOpt. Our ultimate goal is to automatize all steps of the process, and to also tune other base components as well as the decision rules that underlie the NGOpt algorithm selection process.

Acknowledgments. M. López-Ibáñez is a “Beatriz Galindo” Senior Distinguished Researcher (BEAGAL 18/00053) funded by the Spanish Ministry of Science and Innovation (MICINN). C. Doerr is supported by the Paris Ile-de-France Region (AlgoSelect) and by the INS2I institute of CNRS (RandSearch). T. Eftimov, A. Nikolikj, and G. Cenikj is supported by the Slovenian Research Agency: research core fundings No. P2-0098 and project No. N2-0239. G. Cenikj is also supported by the Ad Futura grant for postgraduate study.

References

1. Aziz-Alaoui, A., Doerr, C., Dréo, J.: Towards large scale automated algorithm design by integrating modular benchmarking frameworks. In: Chicano, F., Krawiec, K. (eds.) *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO Companion 2021*, pp. 1365–1374, ACM Press, New York, NY (2021), doi: [10.1145/3449726.3463155](https://doi.org/10.1145/3449726.3463155)
2. Bezerra, L.C.T., López-Ibáñez, M., Stützle, T.: Automatically designing state-of-the-art multi- and many-objective evolutionary algorithms. *Evolutionary Computation* **28**(2), 195–226 (2020), doi: [10.1162/evco.a.00263](https://doi.org/10.1162/evco.a.00263)
3. Birattari, M.: *Tuning Metaheuristics: A Machine Learning Perspective*, Studies in Computational Intelligence, vol. 197. Springer, Berlin/Heidelberg (2009), doi: [10.1007/978-3-642-00483-4](https://doi.org/10.1007/978-3-642-00483-4)
4. Cartis, C., Fiala, J., Marteau, B., Roberts, L.: Improving the flexibility and robustness of model-based derivative-free optimization solvers (2018)
5. Hansen, N., Auger, A., Finck, S., Ros, R.: Real-parameter black-box optimization benchmarking 2009: Experimental setup. Tech. Rep. RR-6828, INRIA, France (2009), URL <https://hal.inria.fr/inria-00362633/document>
6. Hansen, N., Auger, A., Ros, R., Mersmann, O., Tušar, T., Brockhoff, D.: COCO: A platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* **36**(1), 1–31 (2020), doi: [10.1080/10556788.2020.1808977](https://doi.org/10.1080/10556788.2020.1808977)
7. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* **9**(2), 159–195 (2001), doi: [10.1162/106365601750190398](https://doi.org/10.1162/106365601750190398)
8. Hoos, H.H.: Automated algorithm configuration and parameter tuning. In: Hamadi, Y., Monfroy, E., Saubion, F. (eds.) *Autonomous Search*, pp. 37–71, Springer (2012), doi: [10.1007/978-3-642-21434-9_3](https://doi.org/10.1007/978-3-642-21434-9_3)
9. Huang, C., Li, Y., Yao, X.: A survey of automatic parameter tuning methods for metaheuristics. *IEEE Transactions on Evolutionary Computation* **24**(2), 201–216 (2020), doi: [10.1109/TEVC.2019.2921598](https://doi.org/10.1109/TEVC.2019.2921598)
10. Johnson, S.G.: The nlopt nonlinear-optimization package (1994), URL <http://github.com/stevengj/nlopt>
11. Kerschke, P., Hoos, H.H., Neumann, F., Trautmann, H.: Automated algorithm selection: Survey and perspectives. *Evolutionary Computation* **27**(1), 3–45 (Mar 2019), doi: [10.1162/evco.a.00242](https://doi.org/10.1162/evco.a.00242)
12. KhudaBukhsh, A.R., Xu, L., Hoos, H.H., Leyton-Brown, K.: SATenstein: Automatically building local search SAT Solvers from Components. *Artificial Intelligence* **232**, 20–42 (2016), doi: [10.1016/j.artint.2015.11.002](https://doi.org/10.1016/j.artint.2015.11.002)
13. Liao, T., Molina, D., Stützle, T.: Performance evaluation of automatically tuned continuous optimizers on different benchmark sets. *Applied Soft Computing* **27**, 490–503 (2015)
14. Liao, T., Montes de Oca, M.A., Stützle, T.: Computational results for an automatically tuned CMA-ES with increasing population size on the CEC’05 benchmark set. *Soft Computing* **17**(6), 1031–1046 (2013), doi: [0.1007/s00500-012-0946-x](https://doi.org/10.1007/s00500-012-0946-x)
15. Liao, T., Stützle, T., Montes de Oca, M.A., Dorigo, M.: A unified ant colony optimization algorithm for continuous optimization. *European Journal of Operational Research* **234**(3), 597–609 (2014)
16. López-Ibáñez, M., Dubois-Lacoste, J., Pérez Cáceres, L., Stützle, T., Birattari, M.: The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* **3**, 43–58 (2016), doi: [10.1016/j.orp.2016.09.002](https://doi.org/10.1016/j.orp.2016.09.002)

17. López-Ibáñez, M., Liao, T., Stützle, T.: On the anytime behavior of IPOP-CMA-ES. In: Coello Coello, C.A., et al. (eds.) *Parallel Problem Solving from Nature, PPSN XII, Lecture Notes in Computer Science*, vol. 7491, pp. 357–366, Springer, Heidelberg (2012), doi: [10.1007/978-3-642-32937-1_36](https://doi.org/10.1007/978-3-642-32937-1_36)
18. López-Ibáñez, M., Stützle, T.: The automatic design of multi-objective ant colony optimization algorithms. *IEEE Transactions on Evolutionary Computation* **16**(6), 861–875 (2012), doi: [10.1109/TEVC.2011.2182651](https://doi.org/10.1109/TEVC.2011.2182651)
19. Mascia, F., López-Ibáñez, M., Dubois-Lacoste, J., Stützle, T.: Grammar-based generation of stochastic local search heuristics through automatic algorithm configuration tools. *Computers & Operations Research* **51**, 190–199 (2014), doi: [10.1016/j.cor.2014.05.020](https://doi.org/10.1016/j.cor.2014.05.020)
20. Meunier, L., Rakotoarison, H., Wong, P., Rozière, B., Rapin, J., Teytaud, O., Moreau, A., Doerr, C.: Black-box optimization revisited: Improving algorithm selection wizards through massive benchmarking. *IEEE Trans. Evol. Comput.* **26**(3), 490–500 (2022), doi: [10.1109/TEVC.2021.3108185](https://doi.org/10.1109/TEVC.2021.3108185)
21. Pagnozzi, F., Stützle, T.: Automatic design of hybrid stochastic local search algorithms for permutation flowshop problems. *European Journal of Operational Research* **276**, 409–421 (2019), doi: [10.1016/j.ejor.2019.01.018](https://doi.org/10.1016/j.ejor.2019.01.018)
22. Rapin, J., Teytaud, O.: Nevergrad: A gradient-free optimization platform. <https://github.com/FacebookResearch/Nevergrad> (2018)
23. Rice, J.R.: The algorithm selection problem. *Advances in Computers* **15**, 65–118 (1976)
24. Ros, R., Hansen, N.: A simple modification in CMA-ES achieving linear time and space complexity. In: Rudolph, G., et al. (eds.) *Parallel Problem Solving from Nature, PPSN X, Lecture Notes in Computer Science*, vol. 5199, pp. 296–305, Springer, Heidelberg (2008)
25. Schede, E., Brandt, J., Tornede, A., Wever, M., Bengs, V., Hüllermeier, E., Tierney, K.: A survey of methods for automated algorithm configuration (2022), doi: [10.48550/ARXIV.2202.01651](https://doi.org/10.48550/ARXIV.2202.01651)
26. Styles, J., Hoos, H.H.: Ordered racing protocols for automatically configuring algorithms for scaling performance. In: Blum, C., Alba, E. (eds.) *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2013*, pp. 551–558, ACM Press, New York, NY (2013), ISBN 978-1-4503-1963-8, doi: [10.1145/2463372.2463438](https://doi.org/10.1145/2463372.2463438)
27. Videau, M., Leite, A., Teytaud, O., Schoenauer, M.: Multi-objective genetic programming for explainable reinforcement learning. In: Medvet, E., Pappa, G., Xue, B. (eds.) *Proceedings of the 25th European Conference on Genetic Programming, EuroGP 2022*, pp. 256–281, *Lecture Notes in Computer Science*, Springer Nature, Cham, Switzerland (2022)
28. Xu, L., Hoos, H.H., Leyton-Brown, K.: Hydra: Automatically configuring algorithms for portfolio-based selection. In: Fox, M., Poole, D. (eds.) *Proceedings of the AAAI Conference on Artificial Intelligence*, AAAI Press (2010)
29. Xu, L., Hutter, F., Hoos, H.H., Leyton-Brown, K.: SATzilla: portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research* **32**, 565–606 (Jun 2008), doi: [10.1613/jair.2490](https://doi.org/10.1613/jair.2490)