

COCOpf: An Algorithm Portfolio Framework

Petr Baudiš

Dept. of Cybernetics, Czech Technical University, Technická 2, 166 27 Praha, Czech Republic

pasky@ucw.cz

Abstract. *Algorithm portfolios represent a strategy of composing multiple heuristic algorithms, each suited to a different class of problems, within a single general solver that will choose the best suited algorithm for each input. This approach recently gained popularity especially for solving combinatoric problems, but optimization applications are still emerging. The COCO platform [6] [5] of the BBOB workshop series is the current standard way to measure performance of continuous black-box optimization algorithms.*

As an extension to the COCO platform, we present the Python-based COCOpf framework that allows composing portfolios of optimization algorithms and running experiments with different selection strategies. In our framework, we focus on black-box algorithm portfolio and online adaptive selection. As a demonstration, we measure the performance of stock SciPy [8] optimization algorithms and the popular CMA algorithm [4] alone and in a portfolio with two simple selection strategies. We confirm that even a naive selection strategy can provide improved performance across problem classes.

Keywords

Algorithm portfolio, continuous black-box optimization, hyperheuristics, software engineering.

1. Introduction

The problem of Algorithm Selection is not new [16], but has only recently gained traction in particular in the field of combinatoric problem solvers. [9] There are multiple ways to approach the issue, in our work we adopt the abstraction of algorithm portfolios. [3] The basic idea is that we have a set of *heuristic algorithms* at our disposal (each suitable for a different class of problem instances), and for each problem instance on input, we apply a *selection strategy* to pick an algorithm from this portfolio and apply it on the problem. We can perform the selection once or along a fixed schedule (*offline selection*) based on features of the problem instance, or in multiple rounds first exploring the suitability of portfolio members, then allocating time to them based on their previous performance. In successive rounds,

algorithms can be either resumed from their previous state or restarted; we take the approach of resuming them, reserving the restart schedule to be an internal matter of each algorithm.¹

Applying algorithm portfolios on continuous black-box optimization is still a fresh area of research.² The main results so far lie either in population methods, combining a variety of genetic algorithms together, e.g. the AMALGAM-SO algorithm [17], or in offline methods based chiefly on exploratory landscape analysis [11].

In our work, we aim to allow the more traditional optimization methods to enter the mix as we regard the algorithms as *black-box*, i.e. we completely avoid modifying their inner working and we simply call them to retrieve a single result. This allows state-of-art methods to be combined in a single portfolio easily and offers easy implementation of whatever selection strategy desired as the existing algorithm modules can be simply reused without modification. Furthermore, we focus on online adaptive selection that selects algorithms based on their performance so far and does not require possibly expensive feature extraction and training.

The currently accepted de-facto standard for benchmarking optimization methods is the *COMparing Continuous Optimisers* COCO platform [6] [5] that was originally developed for the BBOB workshop series. It provides the infrastructure, glue code for both running experiments and preparing high quality figures, a set of common reference results and the code for a set of benchmark functions. Therefore, we chose to add algorithm portfolio support to this platform. The glue code is available in multiple languages, we extended the Python implementation. We are in part motivated by the convenient availability of optimization algorithms distributed along the popular SciPy library [8].

We have two options when testing algorithm selection methods in practice: either using only previously published data on performance of individual algorithms, or re-running optimization algorithms from scratch under the guide of a particular selection method. The former approach is certainly an easier option while benchmarking, nevertheless we

¹That is, sometimes resumption will simply continue the algorithm to a next iteration, on reaching a local optimum the algorithm may attempt to escape it and failing that, a complete restart may occur.

²Especially if we do not consider applying related methods to individual operator selection within genetic algorithms used for optimization.

opted for the latter model so that we can also readily apply all code on practical function optimization aside of mere benchmarking. Apart of that, the available performance data does not include information about algorithm iterations, listing only individual function evaluations. That makes it unclear when is a good time to switch. Lastly, the re-running model allows migration of intermediate solutions between algorithms as a simple extension.

In Section 2, we introduce the COCOpf software, explaining its architecture, general API and a few interesting implementation points. In Section 3, we establish a reference algorithm portfolio, benchmarking its individual members and two algorithm selection methods implemented in the framework. We present the benchmark results in Section 4 and conclude with a results summary and outline of a few directions of future work in Section 5.

2. COCOpf Framework

Our framework has been implemented as a Python module that can be imported by user applications. The framework provides the means to run the experiment on a given set of benchmark scenarios, maintains a population of algorithm instances and provides a complex wrapper to individual algorithms that uses an existing per-iteration callback mechanism to suspend and resume execution of the algorithm after a single iteration. Algorithm wrappers for the SciPy optimization methods and the reference implementation of CMA are available out-of-the box as well as two example algorithm selection strategies that are benchmarked below. The framework is publicly available as free software under the permissible MIT licence at <https://github.com/pasky/cocopf>.

In design of such a framework, a choice presents itself — to either build a generic main program which allows an experiment to choose a certain configuration and plug in custom callbacks, or to prepare a set of “lego bricks” that replace common parts of the main program but letting each experiment provide a separate main program. We opted for the latter; it reduces encapsulation and requires slightly more repetitive code, but it allows a much greater degree of flexibility encouraging diverse experiments.

2.1. Experiment Wrapper

The **Experiment** class wraps up the common code that sets up the COCO benchmark, initializing the logs, sets of dimensions and function instances etc. It also provides an iterator that in turn returns individual function instances to be minimized, and a pretty-printing progress report function. After an experiment is finished, the standard BBOB postprocessing tools may be used to examine and plot the performance data. An example of usage is shown in Fig. 1.

```
e = Experiment(maxfev=10000, shortname="EG50",
              comments="Eps-greedy, eps=0.5")
for i in e.finstances():
    minimization(i)
    e.f.finalizerun()
    e.f.report(i)
```

Fig. 1. A minimalistic main body of a Python script using the Experiment class. Function *minimization* is to be provided by the user.

2.2. Minimization Wrappers

All algorithms are executed via the **MinimizeMethod** class. Its primary purpose is to produce a callable Python object from a user-provided string name of algorithm. By default, it calls the *scipy.optimize.minimize* API with unchanged method name and wraps it in the *scipy.optimize.basinhopping* API. However, it is expected the user will create a subclass adding custom optimizers; a demo subclass adding the CMA optimization is included.³

We were presented with a software engineering problem of how to suspend, resume and cancel the algorithms between iterations without requiring any modification to the algorithm code, while all we have is a callback executed after each iteration. Our answer is the **MinimizeStepping** class which creates a separate thread for each algorithm instance and uses the standard Python *Queue* mechanism to pause execution of each thread within the callback code based on instructions from the main thread. At any time, only a single thread is running. To cancel execution (when time is up or a solution is found) we pass a special message through the queue that raises an exception within the callback, eventually cancelling the whole algorithm and thread orderly. An “mlog” facility is provided to log function evaluation status at the end of each algorithm iteration.

2.3. Portfolio Population

The **Population** class maintains a population of algorithms drawn from the algorithm portfolio; each population member corresponds to a tuple of method, solution coordinates, current value and number of iterations invested. Usually, the size of population is static, but dynamically adding new members is also supported.

In a typical case, the population size is the same as the portfolio size, meaning that only a single solution is tracked for each algorithm. However, in principle this may not be the case. In fact, we can also make a population from solutions all using a single algorithm — this special case corresponds to the problem of finding an efficient restart or resume schedule for a particular algorithm.

³The custom optimizers may also legitimately want to use the Basin Hopping algorithm (see Sec. 3.1), but SciPy allows using custom minimization routines with the *basinhopping* API only starting with version 0.14.

3. Benchmarking Simple Portfolios

To showcase COCOpf, we composed a reference portfolio from readily available open source Python optimizers and implemented two selection strategies.

3.1. Reference Algorithm Portfolio

As a testbed algorithm portfolio, we have chosen the six stock minimizers provided by SciPy v0.13 [8] that are available for direct use:⁴

- **Nelder-Mead** uses the Simplex algorithm. [12] [19]
- **Powell** implements a tweaked version of the Powell’s conjugate direction method. [14]
- **CG** is a nonlinear conjugate gradient method using the Fletcher-Reeves method. [13, p. 120–122]
- **BFGS** uses the quasi-Newton method of Broyden, Fletcher, Goldfarb and Shanno [13, p.136] with numerically approximated derivations.
- **L-BFGS-B** is a limited-memory variant of BFGS with box constraints. [1] [20]
- **SLSQP** implements Sequential Least Squares Programming [10] with inequalities as the box constraints.

All of these SciPy minimizers perform local minimization; to achieve global optimization, we wrap them in the “Basin Hopping” restart strategy [18] (also provided by SciPy), which is conceptually similar to Simulated Annealing with a fixed temperature. The strategy performs the hopping step 100 times (the default) before restarting from scratch. On restart, the starting position is determined uniformly randomly in the $[-5, 5]^k$ space.

To improve the portfolio performance on more difficult functions, we also included the popular CMA algorithm [4] (in the “production“ version of its official reference Python implementation), i.e. a genetic algorithm that follows the Covariance Matrix Adaptation evolution strategy.

In accordance with our black-box approach, we kept all the algorithms at their default settings and did not tune any of their parameters. Where applicable, we set the box constraints to $[-6, 6]^k$ bounds.⁵

Each algorithm provides a callback method that is invoked after a single iteration of the algorithm; we suspend the algorithm after each callback invocation.

⁴Other minimizers are also available in SciPy, but they are either unsuitable for black-box minimization or do not provide a callback mechanism to allow pausing execution between individual iterations.

⁵The optimum is guaranteed to be in $[-5, 5]^k$, but some minimizers may behave erratically when the optimum is near to or at the bounds.

The reference portfolio is not perfectly balanced and certainly would benefit from algorithms with more diverse performance. Our motivation for going with these is that the implementations are free software, trivial to obtain and therefore representing a very low barrier of entry. Also, a SciPy-based portfolio gives valuable feedback to SciPy users regarding the best algorithm to use on a certain function, and provides a gateway to introduce an algorithm portfolio selection strategy to SciPy itself in the future.

3.2. Algorithm Selection Strategies

While the core of our followup work involves researching various algorithm selection strategies, in the context of presentation of the COCOpf framework, we decided to demonstrate just the performance of two simple simplest methods whose source code is also a part of the framework:

- The **UNIF** strategy performs a number of rounds where in every round, a uniformly randomly selected algorithm is run for a single iteration.
- The **EG50** strategy performs a number of rounds where in every round, an *epsilon-greedy policy* with $\varepsilon = 0.5$ selects the algorithm to run for one iteration. That is, the algorithm with the currently best solution is run with $p = 0.5$ and a randomly chosen algorithm is run otherwise.

4. Results

Results from experiments according to [5] on the benchmark functions given in [2, 7] are presented in Figures 2, 3 and 4.⁶ The **expected running time (ERT)** used in the figures depends on a given target function value, $f_t = f_{\text{opt}} + \Delta f$, and is computed over all relevant trials as the number of function evaluations executed during each trial while the best function value did not reach f_t , summed over all trials and divided by the number of trials that actually reached f_t [5, 15].

If we review the results in context of choosing a best (black-box function) SciPy optimizer, we can see that there is no single right choice of a SciPy optimizer that works on all function instances and computational budgets. We can observe that **Powell** generally works best for separable functions while **BFGS** and **SLSQP** are good choices for many functions, especially in high dimensions. For non-separable functions, **CMA** sooner or later overtakes the SciPy optimizers but often with a slow start (which is relevant in case a function evaluation is very expensive) and is not competitive with the SciPy optimizers on separable functions.

⁶The full results dataset is available at <http://pasky.or.cz/sci/cocopf-scipy>.

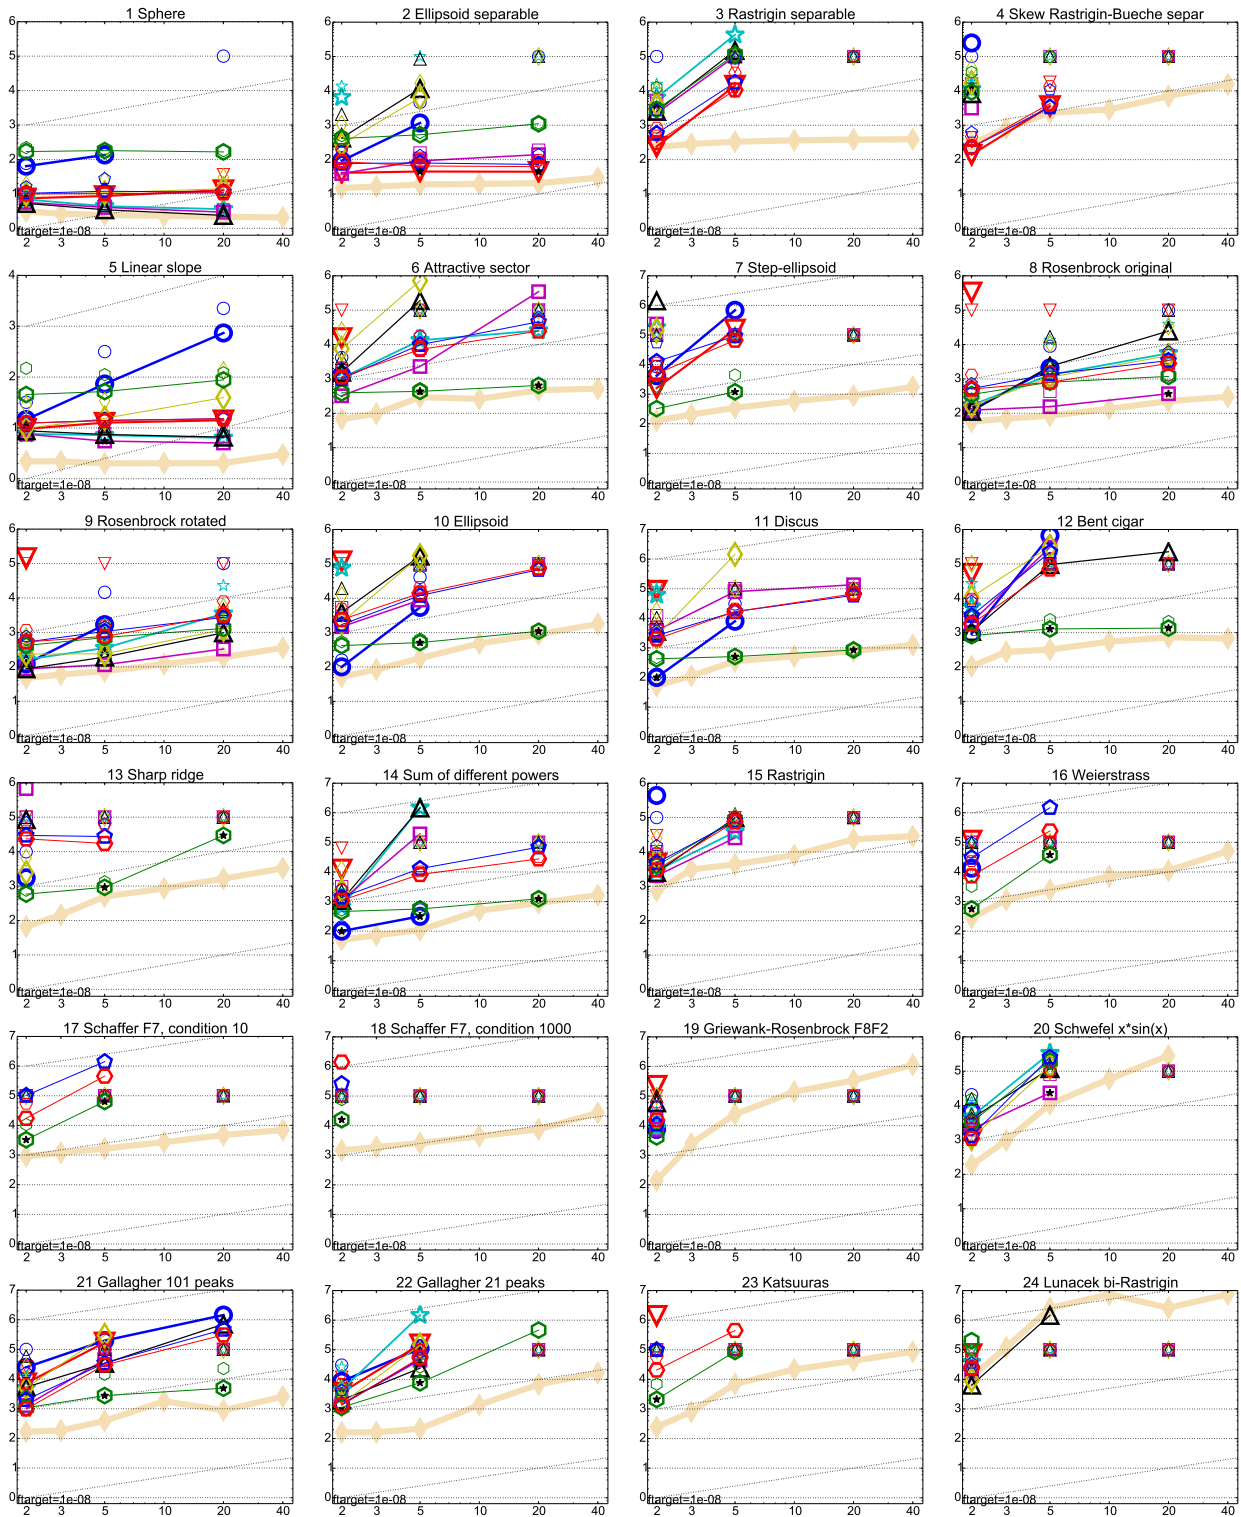


Fig.2. Expected running time divided by dimension versus dimension. Expected running time (ERT in number of f -evaluations) divided by dimension for target function value 10^{-8} as \log_{10} values versus dimension. Light symbols give the maximum number of function evaluations from the longest trial divided by dimension. Horizontal lines give linear scaling, slanted dotted lines give quadratic scaling. Black stars indicate statistically better result compared to all other algorithms with $p < 0.01$ and Bonferroni correction number of dimensions (six). Legend: \circ :Nelder-Mead, ∇ :Powell, \star :CG, \square :BFGS, \triangle :L-BFGS-B, \diamond :SLSQP, \odot :CMA, \odot :UNIF, \odot :EG50

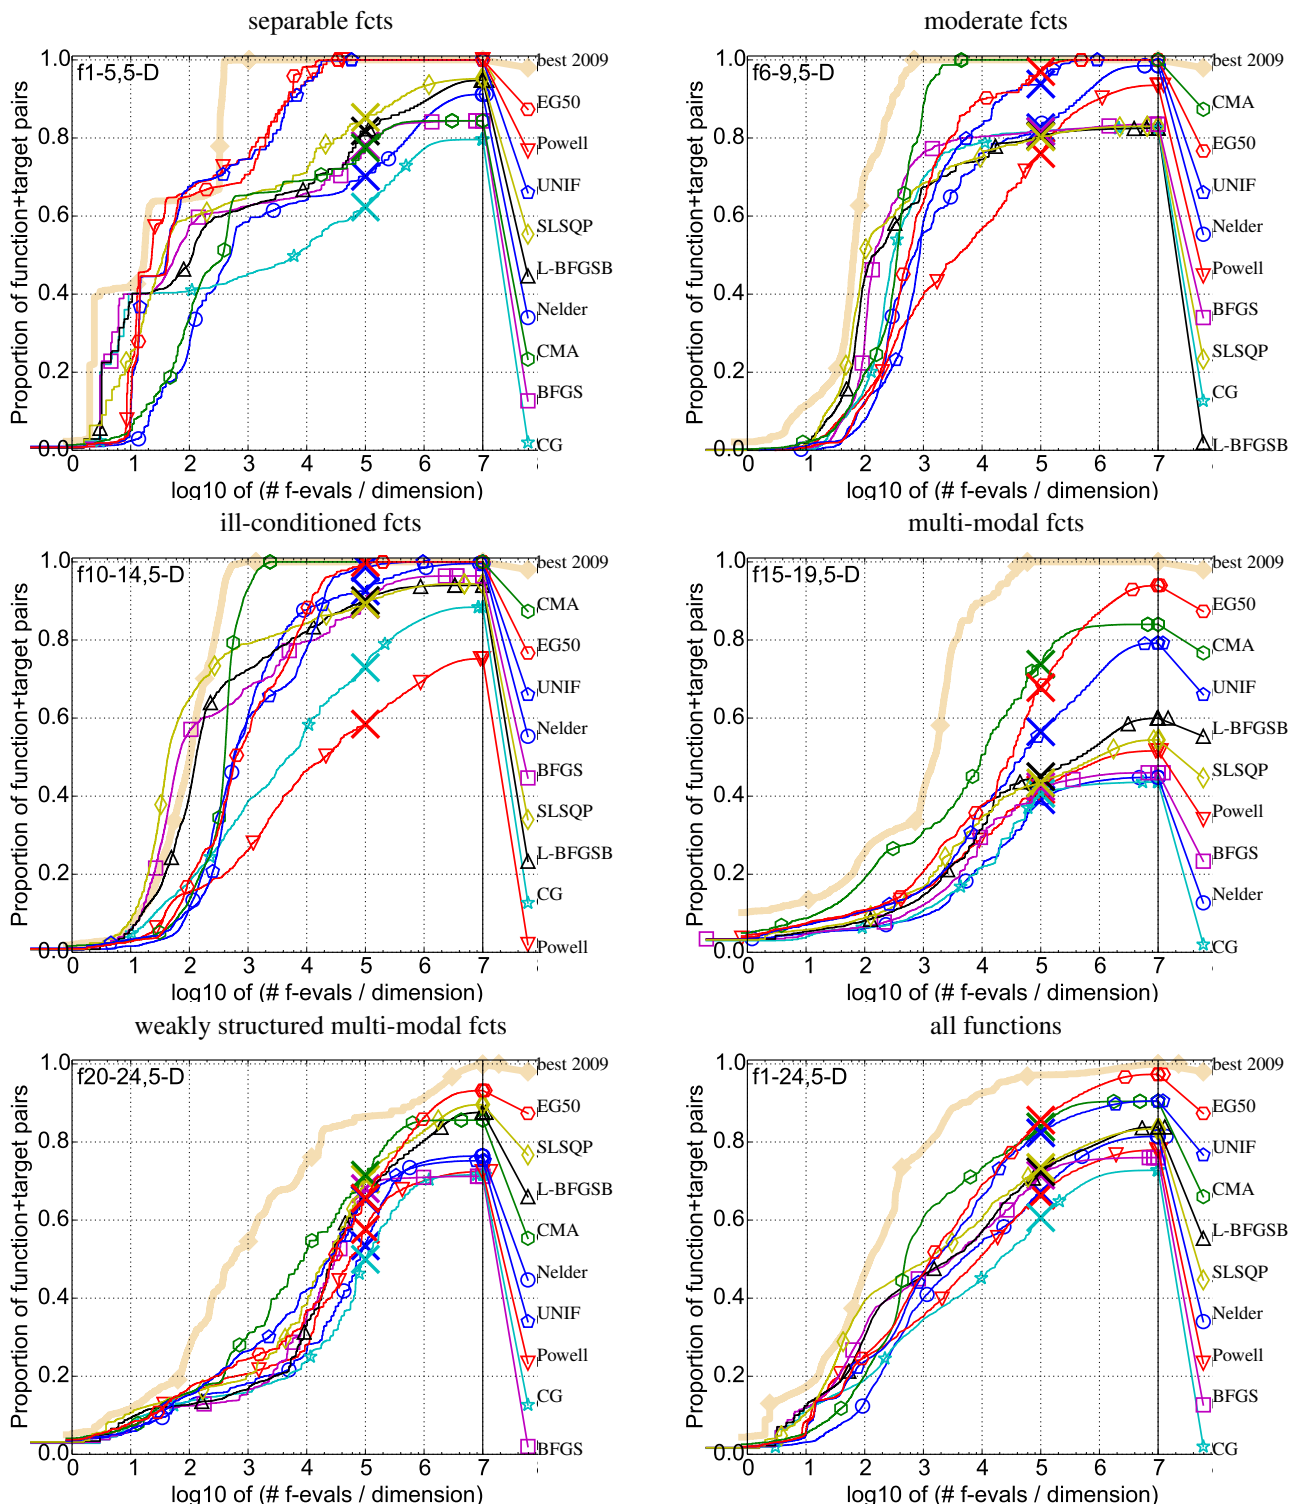


Fig.3. Bootstrapped empirical cumulative distribution of the number of objective function evaluations divided by dimension (FEvals/D) for 50 targets in $10^{[-8..2]}$ for all functions and subgroups in 5-D. The “best 2009” line corresponds to the best ERT observed during BBOB 2009 for each single target.

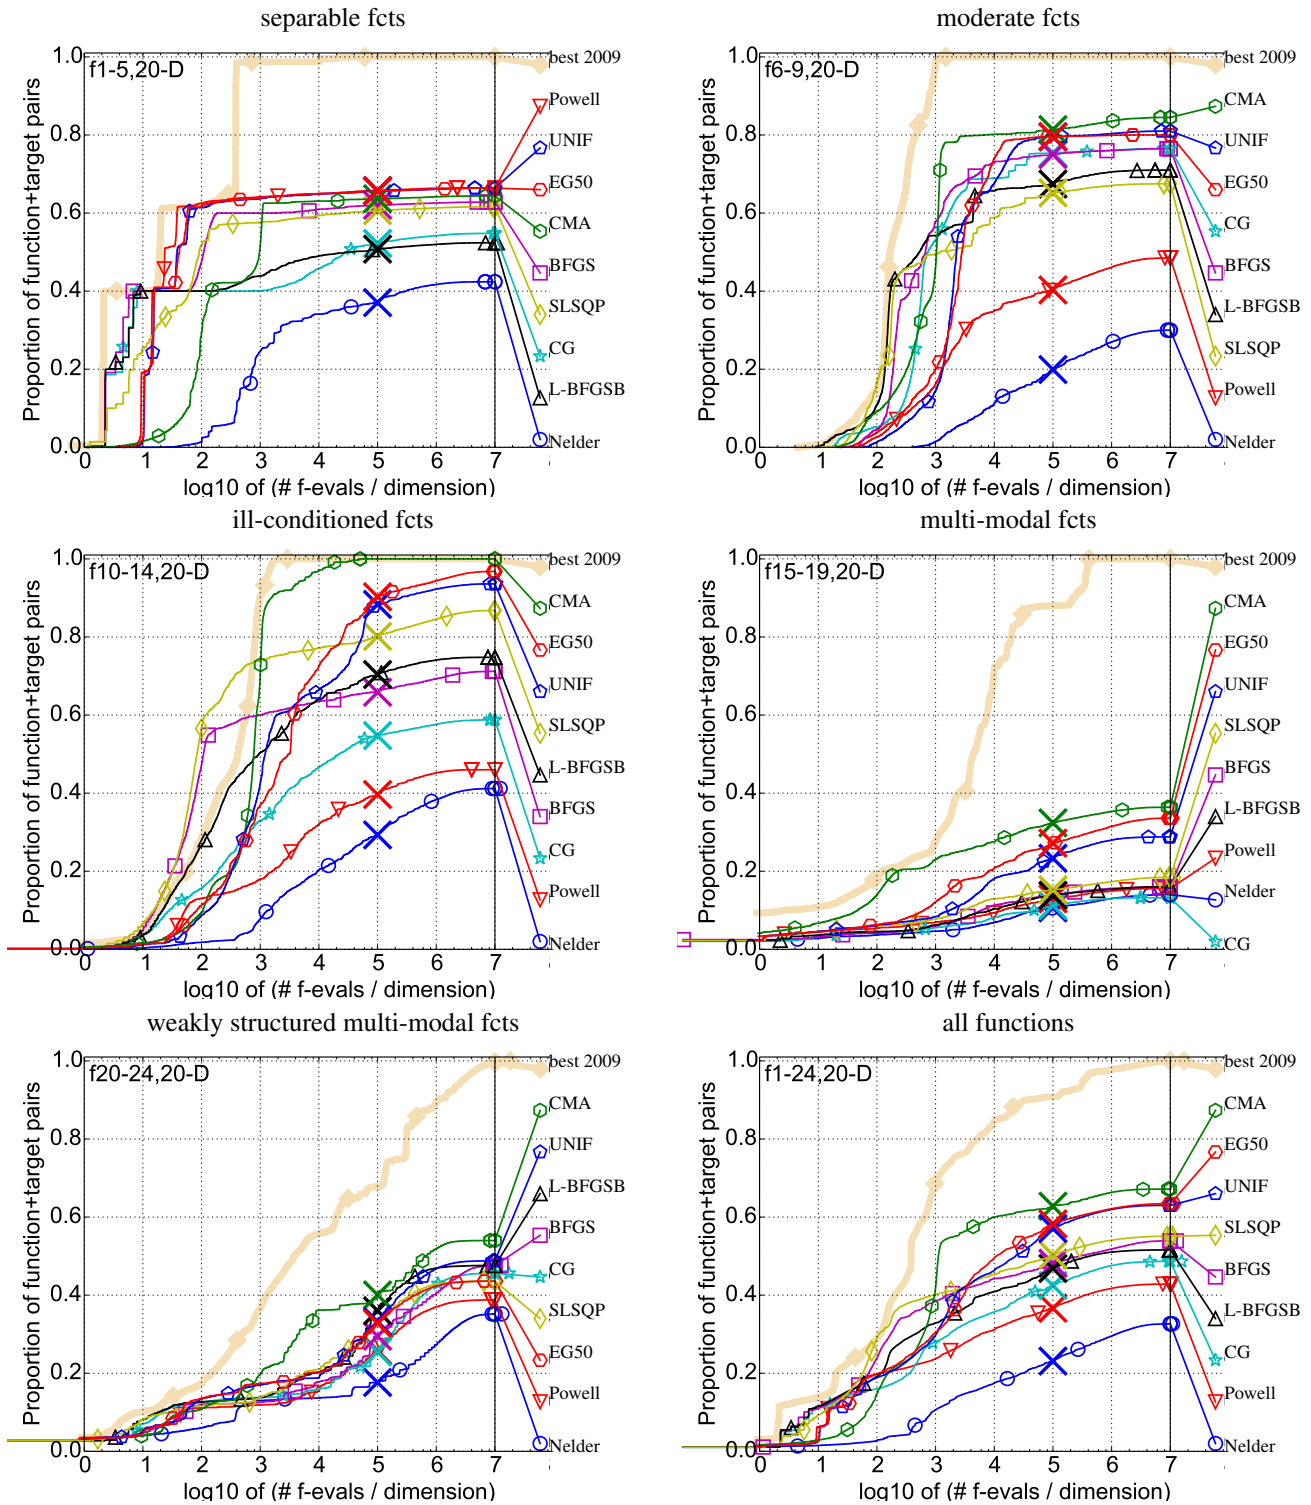


Fig.4. Bootstrapped empirical cumulative distribution of the number of objective function evaluations divided by dimension (FEvals/D) for 50 targets in $10^{[-8..2]}$ for all functions and subgroups in 20-D. The “best 2009” line corresponds to the best ERT observed during BBOB 2009 for each single target.

5. Discussion and Conclusion

We described an easy-to-use framework that allows users to plug in and experiment with custom algorithm selection strategies and then directly move to use these strategies on practical problems instead of benchmarks. The framework currently focuses on online black-box algorithm portfolios, being careful not to require any modification to stock and third-party optimization routines besides a simple callback mechanism that is commonly provided. The framework could be easily modified to deal with off-line scenarios too.

Our insights may provide guidance for users of the SciPy optimizers. Regarding algorithm portfolios, the results show that even a very naive algorithm selection strategy can be beneficial; while it cannot beat the best algorithm CMA in high-dimensional cases on average over all functions, it provides a more consistent performance than CMA, translating to best-of-all performance in low dimensions.

5.1. Future Work

The currently chosen reference portfolio is somewhat ad hoc, chosen in part to contribute value to the SciPy project, but not optimal for further general research into algorithm portfolios. In the long run, we hope to transition to a better balanced portfolio.

A common scheme of online algorithm selection schemes is that a *credit* is assigned to an algorithm in each round based on its performance (e.g. relative improvement, absolute value) and *accrued* over the rounds (e.g. averaging it or using an adaptation rule). Credit-based selection is so common that it should be directly supported by the framework; diverse pluggable strategies would then enable exploration of various interesting strategy combinations. A work-in-progress preview implementation is already available.

Actually *running* the individual algorithms is appropriate for optimizing non-benchmark functions and useful in the future when solutions may migrate within the portfolio population. However, a useful feature would be to allow testing selection strategies based on just the saved “mlog” data (see Sec. 2.2) from individual algorithm runs.

Acknowledgements

Research described in the paper was supervised by Dr. Petr Pošík and supported by the CTU grant SGS14/194/OHK3/3T/13 “Automatic adaptation of search algorithms”.

References

- [1] BYRD, R. H., LU, P., NOCEDAL, J., AND ZHU, C. A limited memory algorithm for bound constrained optimization. *SIAM Journal on*

- Scientific Computing* 16, 5 (1995), 1190–1208.
- [2] FINCK, S., HANSEN, N., ROS, R., AND AUGER, A. Real-parameter black-box optimization benchmarking 2009: Presentation of the noiseless functions. Tech. Rep. 2009/20, Research Center PPE, 2009. Updated February 2010.
- [3] GOMES, C. P., AND SELMAN, B. Algorithm portfolios. *Artificial Intelligence* 126, 1 (2001), 43–62.
- [4] HANSEN, N. The CMA evolution strategy: a comparing review. In *Towards a new evolutionary computation. Advances on estimation of distribution algorithms*, J. Lozano, P. Larranaga, I. Inza, and E. Bengoetxea, Eds. Springer, 2006, pp. 75–102.
- [5] HANSEN, N., AUGER, A., FINCK, S., AND ROS, R. Real-parameter black-box optimization benchmarking 2012: Experimental setup. Tech. rep., INRIA, 2012.
- [6] HANSEN, N., ET AL. Comparing continuous optimisers: Coco. <http://coco.gforge.inria.fr/>.
- [7] HANSEN, N., FINCK, S., ROS, R., AND AUGER, A. Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Tech. Rep. RR-6829, INRIA, 2009. Updated February 2010.
- [8] JONES, E., OLIPHANT, T., PETERSON, P., ET AL. SciPy: Open source scientific tools for Python, 2001–.
- [9] KOTTHOFF, L. Algorithm selection for combinatorial search problems: A survey. *AI Magazine* (2014).
- [10] KRAFT, D. *A software package for sequential quadratic programming*. DFVLR Obersaffenhofen, Germany, 1988.
- [11] MUÑOZ, M., KIRLEY, M., AND HALGAMUGE, S. The algorithm selection problem on the continuous optimization domain. In *Computational Intelligence in Intelligent Data Analysis*, C. Moewes and A. Nürnberger, Eds., vol. 445 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, 2013, pp. 75–89.
- [12] NELDER, J. A., AND MEAD, R. A simplex method for function minimization. *The Computer Journal* 7, 4 (1965), 308–313.
- [13] NOCEDAL, J., AND WRIGHT, S. *Numerical optimization*. Springer-Verlag, 2006.
- [14] POWELL, M. J. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The computer journal* 7, 2 (1964), 155–162.
- [15] PRICE, K. Differential evolution vs. the functions of the second ICEO. In *Proceedings of the IEEE International Congress on Evolutionary Computation* (1997), pp. 153–157.
- [16] RICE, J. R. The algorithm selection problem. *Advances in Computers* 15 (1976), 65–118.
- [17] VRUGT, J. A., ROBINSON, B. A., AND HYMAN, J. M. Self-adaptive multimethod search for global optimization in real-parameter spaces. *IEEE Trans. on Evolutionary Computation* 13, 2 (2009), 243–259.
- [18] WALES, D. J., AND DOYE, J. P. K. Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A* 101, 28 (1997), 5111–5116.
- [19] WRIGHT, M. H. Direct search methods: Once scorned, now respectable. *Pitman Research Notes in Math. Series* (1996), 191–208.
- [20] ZHU, C., BYRD, R. H., LU, P., AND NOCEDAL, J. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)* 23, 4 (1997), 550–560.

About Author...

Petr Baudiš has received his Bachelors degree and Masters degree in Theoretical Computer Science at the Charles University in Prague. His Master thesis “MCTS with Information Sharing” presented a state-of-art Computer Go program. Currently, he is a PhD student at the Czech Technical University with principal research interest in Algorithm Portfolios and applications of Monte Carlo Tree Search.