



**HAL**  
open science

# The Impact of Initial Designs on the Performance of MATSuMoTo on the Noiseless BBOB-2015 Testbed: A Preliminary Study

Dimo Brockhoff, Bernd Bischl, Tobias Wagner

► **To cite this version:**

Dimo Brockhoff, Bernd Bischl, Tobias Wagner. The Impact of Initial Designs on the Performance of MATSuMoTo on the Noiseless BBOB-2015 Testbed: A Preliminary Study. GECCO '15 Companion, Jul 2015, Madrid, Spain. 10.1145/2739482.2768470 . hal-01161504

**HAL Id: hal-01161504**

**<https://inria.hal.science/hal-01161504>**

Submitted on 8 Jun 2015

**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.

# The Impact of Initial Designs on the Performance of MATSuMoTo on the Noiseless BBOB-2015 Testbed: A Preliminary Study

Dimo Brockhoff  
INRIA Lille - Nord Europe  
DOLPHIN research team  
Villeneuve d'Ascq, France  
firstname.lastname@inria.fr

Bernd Bischl  
LMU München  
Institut für Statistik  
Munich, Germany  
firstname.lastname@  
stat.uni-muenchen.de

Tobias Wagner  
TU Dortmund University  
Institute of Machining  
Technology (ISF)  
Dortmund, Germany  
lastname@isf.de

## ABSTRACT

Most surrogate-assisted algorithms for expensive optimization follow the same framework: After an initial design phase in which the true objective function is evaluated for a few search points, an iterative process builds a surrogate model of the expensive function and, based on the current model, a so-called infill criterion suggests one or more points to be evaluated on the true problem. The evaluations are used to successively update and refine the model. Implementing surrogate-assisted algorithms requires several design choices to be made. It is practically relevant to understand their impact on the algorithms' performance. Here, we start to look at the initial design phase and experimentally investigate the performance of the freely available MATLAB Surrogate Model Toolbox (MATSuMoTo) with regard to the initial design. The results are preliminary in the sense that not all possible choices are investigated, but we can make first well-founded statements about whether Latin Hypercube or uniform random sampling should be preferred and about the effect of the size of the initial design on the performance of MATSuMoTo on the 24 noiseless test functions of the BBOB-2015 test suite.

## Categories and Subject Descriptors

G.1.6 [Numerical Analysis]: Optimization—*global optimization, unconstrained optimization*; F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems

## Keywords

Benchmarking, Black-box optimization, Expensive problems

## 1. INTRODUCTION

Expensive optimization problems in which the evaluation of the objective function might take minutes or hours of (parallelized) optimization time occur frequently in practical applications, e.g., in the case of simulation-based engineering design or in biological and medical applications. Algorithms, especially designed for this expensive optimization scenario where only a few hundreds to thousands of function evaluations are affordable, often follow a common framework. In an initial design phase, a few representative solutions are sampled in the search space and evaluated on the true, expensive objective function. Then, an iterative process builds a surrogate- or meta-model of the objective function, suggests a batch of new solutions to be evaluated on the expensive objective function based on the model and a user-defined infill criterion, and uses the new evaluated solutions to again update the model. The actual implementation of an algorithm for an expensive optimization problem, requires a choice among several building blocks for each part of this iterative process. Many surrogate models and several infill criteria have been proposed in the literature [16].

Also for the initial design, several approaches have been suggested. Most often, Latin Hypercube sampling (LHS, [12]) is employed, but also uniform random or quasi-random sampling methods such as Halton or Sobol sequences [6, 15] can be applied. Over the last 15 years, many studies have been conducted in order to find the optimal design for fitting surrogate models. Koehler and Owen [11] have provided an overview of theoretical results and visual examples. Comparative studies [12, 4] can only conclude that LHS is superior to random sampling in some cases.

With respect to size of the initial design, the recommendation  $10 \cdot \text{DIM}$  was once established [10] with DIM being the search space dimension. More recent studies have shown that also much smaller initial designs can produce good results given that the modeled response has not too many local optima or is not too irregular [2, 9].

To measure and visualize performance, several benchmarking suites have been proposed of which the Black-box Optimization Benchmarking suite (BBOB) is one of the most sophisticated ones [7]. The underlying Coco software allows to automatically collect data during the optimization runs and to produce comparison plots and result tables with a single python script. Due to the simplicity of algorithm benchmarking, we will employ the BBOB framework here.

With respect to the algorithm implementations, we rely on the MATLAB Surrogate Model Toolbox (MATSuMoTo) which has recently been provided for free to the research community [13] and which allows to easily plug in and test several initial designs, surrogate models, and infill criteria. Note that the default version of the MATSuMoTo library has been already benchmarked on the BBOB testbed [3] without investigating the impact of its many parameters and algorithm design choices.

### Our Contribution.

In this paper, we provide a first preliminary experimental investigation of the impact of the initial design on the performance of MATSuMoTo on the BBOB testbed. In particular, we try to answer the following research questions:

- Q1: What is the effect of having larger initial designs, such as two times or ten times the default value of  $2 \cdot (\text{DIM} + 1)$  function evaluations?  
 Q2: What is the effect of replacing the LHS with simple (uniform pseudo-)random sampling?

Both questions will be answered empirically by comparing different MATSuMoTo variants on the 24 noiseless BBOB functions of [5].

## 2. ALGORITHM PRESENTATION

As baseline algorithm, we use the MATSuMoTo library in its default setting (January 2014 from <https://ccse.lbl.gov/people/julianem/>). Note that during our experiments, a small bug affecting the restarting mechanism of MATSuMoTo has been fixed such that the source code differs slightly from the original MATSuMoTo library.<sup>1</sup> The default setting of the MATSuMoTo library—as already previously benchmarked on the BBOB testbed [3]—employs LHS as initial design for  $2 \cdot (\text{DIM} + 1)$  function evaluations, a cubic Radial Basis Functions surrogate model, and an infill criterion which either chooses the next point either as a (small) random perturbation around the model’s minimum (exploitation) or, with a certain probability, uniformly at random in the whole variable domain (exploration).

In order to investigate the impact of the initial design on the library’s performance, we run four different versions of the algorithm that only differ in the type and size of the initial design as indicated in Table 1. The table also introduces the abbreviations for the algorithms used in the following. The default version, LHD-Default corresponds to the algorithm benchmarked in [3] and employs an LHS of size  $2 \cdot (\text{DIM} + 1)$  for fitting the first surrogate model. The algorithms LHD-2xDefault and LHD-10xDefault differ from LHD-Default only in the size of the initial design which is two times and ten times larger than the default, respectively. Finally, RAND-2xDefault samples the initial points in the bounded search space uniformly at random for the same number of function evaluations as LHD-2xDefault. All algorithm variants are run in total for 50·DIM function evaluations and otherwise use the default settings of the library. The upper and lower bounds for the initial sampling within MATSuMoTo are set to -5 and 5, respectively, according to the definition of the BBOB test functions which guarantees that the global optimum lies within this range.

<sup>1</sup>The source code for reproducing our results can be found on the Coco/BBOB web page under <http://coco.gforge.inria.fr/doku.php?id=bbob-2015-results>.

algorithm name	initial design	length of initial design phase
LHD-Default	LHS	default
LHD-2xDefault	LHS	2x default
RAND-2xDefault	random	2x default
LHD-10xDefault	LHS	10x default

**Table 1: Overview of the algorithms investigated. LHS: Latin hypercube design, default length of initial design phase is  $2 \cdot (\text{DIM} + 1)$ .**

	2-D	3-D	5-D	10-D	20-D
LHD-Default	0.0044	0.0040	0.16	0.59	1.9
LHD-2xDefault	0.0048	0.0013	0.20	0.54	1.8
RAND-2xDefault	0.0068	0.031	0.22	0.56	1.8
LHD-10xDefault	0.0021	0.013	0.14	0.41	1.4

**Table 2: Running times per function evaluation on function  $f_8$  (in seconds).**

## 3. CPU TIMING

In order to evaluate the CPU time, we have run all four variants of MATSuMoTo on BBOB function  $f_8$  with restarts for at least 30 seconds *and* for at least 50·DIM function evaluations. The code was run on a Mac Intel(R) Core(TM) i7-2635QM CPU @ 2GHz with 1 processor, 4 cores, and 8GB of RAM. The time per function evaluation in seconds for all algorithms and dimensions are given in Table 2. We observe that the runtime per function evaluation does not differ much between the algorithm variants, only the large design results in slightly shorter times. With increasing dimension, however, the computation time per function evaluation increases drastically. Nevertheless, with less than two seconds in 20-D, the computation time of the algorithm itself is neglectable in an expensive optimization scenario.

## 4. EXPERIMENTAL RESULTS

Results from experiments according to [7] on the benchmark functions given in [5, 8] are presented in Figures 1, 2 and 3 and in Tables 3 and 4. The **expected running time (ERT)**, used in the figures and tables, depends on a given target value,  $f_t = f_{\text{opt}} + \Delta f$ , and is computed as sum over all evaluations executed in each run until  $f_t$  has been obtained divided by the number of trials that actually reached  $f_t$  [7, 14]. **Statistical significance** is tested with the rank-sum test for a given target  $\Delta f_t$  using, for each trial, either the number of function evaluations to reach  $\Delta f_t$  (inverted and multiplied by -1), or, if the target was not reached, the best  $\Delta f$ -value and the smallest number of overall function evaluations for any unsuccessful trial under consideration.

Besides the mentioned LHD-Default, LHD-2xDefault, LHD-10xDefault, and RAND-2xDefault, the pure random search of [1] and SMAC [9], an improved version of the Efficient Global Optimization (EGO) algorithm, have been used as a baseline algorithms. In the following, we investigate the figures and tables, provided by the BBOB framework, with the two main research question introduced in Sec. 1 in mind.

### 4.1 Impact of the Size of the Initial Design

Comparing the algorithms, which only differ in the size of the initial design (LHD-Default, LHD-2xDefault, and LHD-10xDefault), we can state two main observations.

First, during the initial design phase, all three algorithms follow, as expected, roughly the same trend in the empirical cumulative distribution plots, which resembles the performance of the pure random search. The actual improvements start with the proposal of new search points according to the model. Increasing the initial design by a factor of ten results in a maximum performance loss at this time of about a factor of 3 in 5-D and of about a factor of 6 in 20-D.

Second, a smaller initial design seems to be beneficial—for sure in an any-time scenario, but also for the final budget of 50-DIM function evaluations. In accordance with [2], the exception is the class of multi-modal problems ( $f_{15}$ – $f_{19}$ ) where LHD-10xDefault takes the lead in the empirical cumulative distribution function plots of Fig. 2 and Fig. 3 from about 30–40-DIM function evaluations. The results are not significant, however, and the effect is mainly caused by a reduced variance in the expected running time for the more difficult targets. For the smaller budget of function evaluations, the trend of a small initial design being favorable is supported by the performance of SMAC which tries to dispense with the initial design [9]. It is outperforming all other compared algorithms in the beginning of the optimization.

Considering question Q1 from above, we can conclude that smaller initial designs seem beneficial on the BBOB testbed and that, among the chosen settings, the default initial design size works best. We expect, however, that even smaller initial designs might further improve the performance.

## 4.2 Impact of the Type of Initial Design

Regarding question Q2 from above, we can state that there is no significant difference between LHS and random sampling when the initial design size of MATSuMoTo is set to twice the default value. When looking at the Empirical Cumulative Distribution Functions, aggregating over all 24 test functions, the graphs for LHD-2xDefault and RAND-2xDefault are almost superimposed. Even if the interaction with the initial design size cannot be evaluated based on the reported experiments, initial experiments, not shown here, showed the same effect for other sizes of the initial design.

## 5. CONCLUSIONS

Preliminary experiments on the impact of MATSuMoTo's initial design on its performance on the noiseless BBOB testbed revealed that (i) a longer initial design phase seems detrimental for the any-time performance and that (ii) replacing the default LHS with uniform (pseudo-)random sampling does not change much for the algorithm's performance. The results are the first of this kind on the BBOB test suite and preliminary in the sense that the MATSuMoTo also offers *symmetric* LHS and an two-level factorial designs. Based on our results, however, these options will have less effect compared to the design size.

**Acknowledgments:** The authors acknowledge support by the French National Research Agency (ANR) under grant ANR-12-MONU-0009 (NumBBO).

## 6. REFERENCES

- [1] A. Auger and R. Ros. Benchmarking the pure random search on the BBOB-2009 testbed. In F. Rothlauf, editor, *GECCO (Companion)*, pages 2479–2484. ACM, 2009.
- [2] T. Bartz-Beielstein and M. Preuß. Considerations of budget allocation for sequential parameter optimization (SPO). In L. Paquete et al., editors, *Proc. Workshop on Empirical Methods for the Analysis of Algorithms (EMAA 2006)*, pages 35–40, 2006.
- [3] D. Brockhoff. Comparison of the MATSuMoTo Library for Expensive Optimization on the Noiseless Black-Box Optimization Benchmarking Testbed. In *Congress on Evolutionary Computation (CEC 2015)*. IEEE Press, 2015. accepted.
- [4] D. Bursztyn and D. Steinberg. Comparison of designs for computer experiments. *Statistical Planning and Inference*, 136(3):1103–1119, 2006.
- [5] S. Finck, N. Hansen, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking 2009: Presentation of the noiseless functions. Technical Report 2009/20, Research Center PPE, 2009. Updated February 2010.
- [6] J. H. Halton. Algorithm 247: Radical-inverse Quasi-random Point Sequence. *Communications of the ACM*, 7(12):701–702, Dec. 1964.
- [7] N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization benchmarking 2012: Experimental setup. Technical report, INRIA, 2012.
- [8] N. Hansen, S. Finck, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Technical Report RR-6829, INRIA, 2009. Updated February 2010.
- [9] F. Hutter, H. Hoos, and K. Leyton-Brown. An Evaluation of Sequential Model-Based Optimization for Expensive Blackbox Functions. In *GECCO workshop on Black-Box Optimization Benchmarking (BBOB'2013)*, pages 1209–1216. ACM Press, 2013.
- [10] D. R. Jones, M. Schonlau, and W.-J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [11] J. R. Koehler and A. B. Owen. Computer experiments. In S. Ghosh and C. R. Rao, editors, *Handbook of Statistics*, volume 13, pages 261–308. Elsevier, New York, NY, 1996.
- [12] M. D. McKay, R. J. Beckman, and W. J. Conover. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics*, 21:239–245, 1979.
- [13] J. Müller. MATSuMoTo: The MATLAB Surrogate Model Toolbox For Computationally Expensive Black-Box Global Optimization Problems. *ArXiv e-prints*, Apr. 2014.
- [14] K. Price. Differential evolution vs. the functions of the second ICEO. In *Proceedings of the IEEE International Congress on Evolutionary Computation*, pages 153–157, 1997.
- [15] I. M. Sobol. On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967.
- [16] F. A. C. Viana, T. W. Simpson, V. Balabanov, and V. Toropov. Metamodeling in multidisciplinary design optimization: How far have we really come? *AIAA Journal*, 52:670–690, 2014.

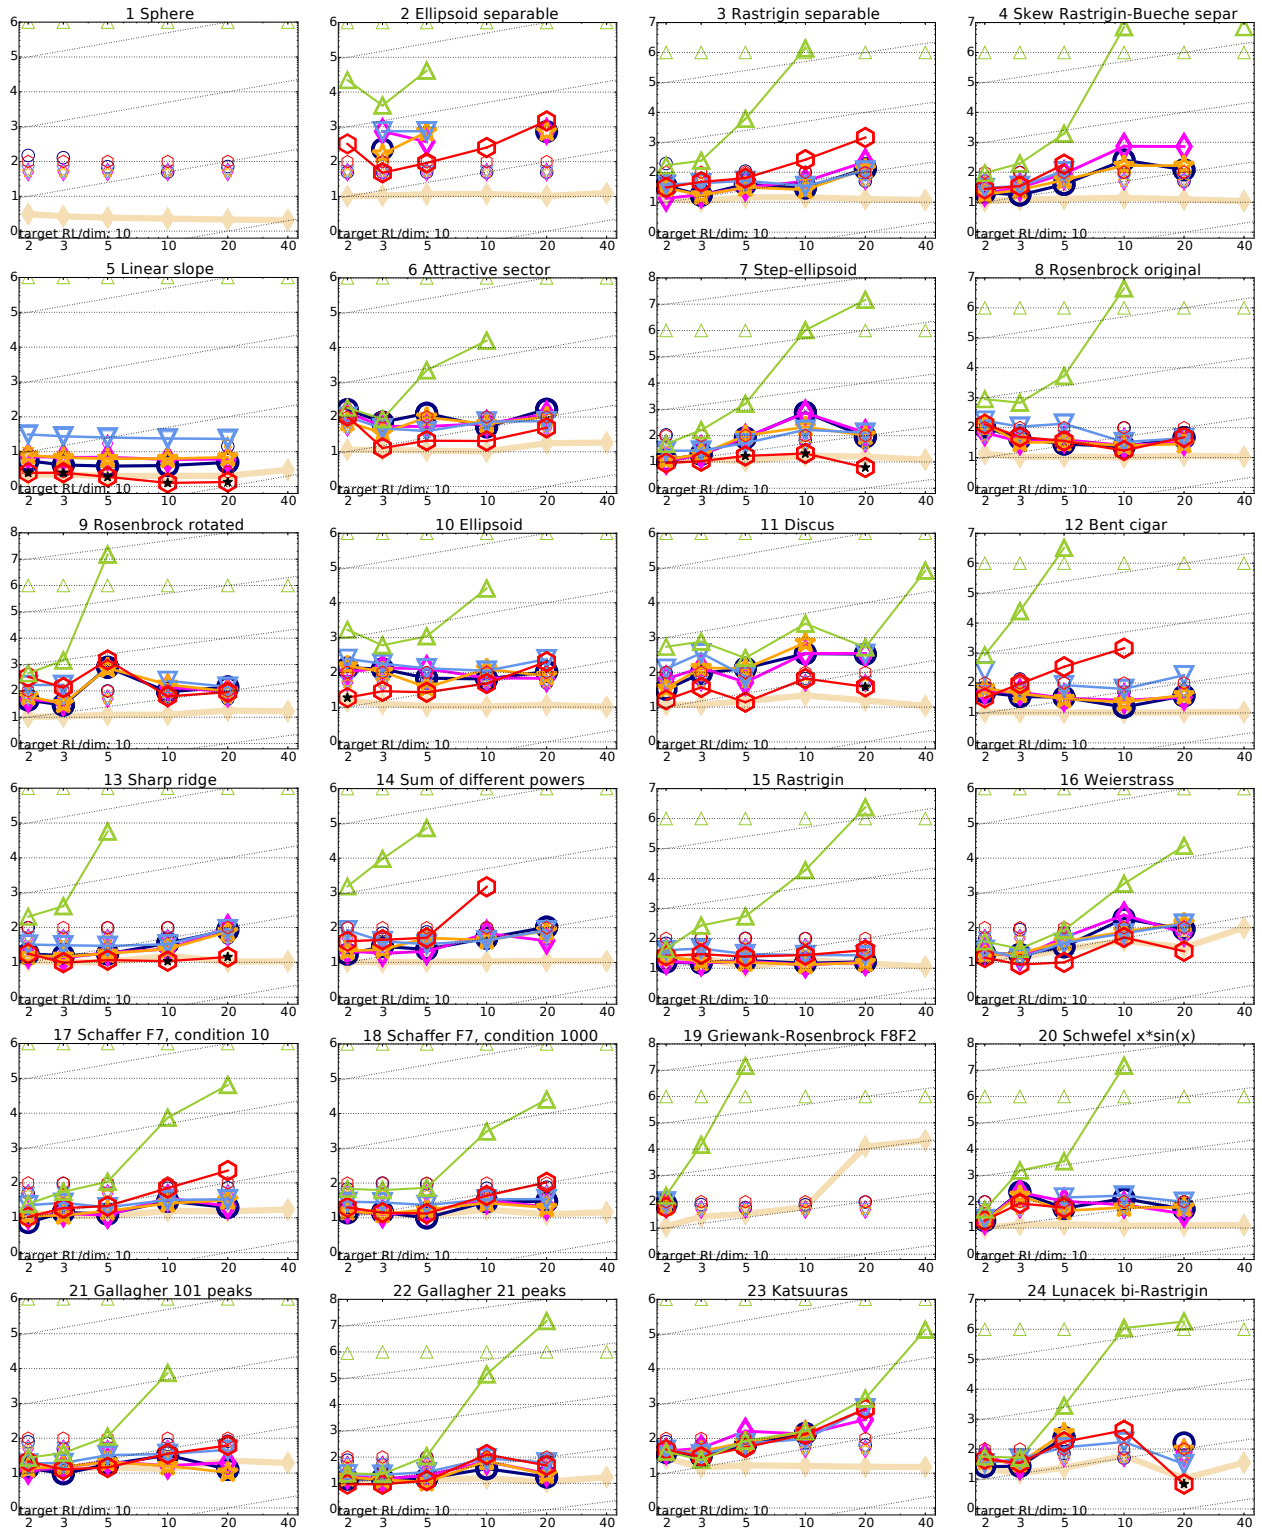


Figure 1: Expected running time (ERT in number of  $f$ -evaluations as  $\log_{10}$  value) divided by dimension versus dimension. The target function value is chosen such that the bestGECCO2009 artificial algorithm just failed to achieve an ERT of  $10 \times \text{DIM}$ . Different symbols correspond to different algorithms given in the legend of  $f_1$  and  $f_{24}$ . Light symbols give the maximum number of function evaluations from the longest trial divided by dimension. Black stars indicate a statistically better result compared to all other algorithms with  $p < 0.01$  and Bonferroni correction number of dimensions (six). Legend:  $\circ$ :LHD-Default,  $\diamond$ :LHD-2xDefault,  $\star$ :RAND-2xDefault,  $\nabla$ :LHD-10xDefault,  $\square$ :SMAC,  $\triangle$ :RANDOMSEARCH

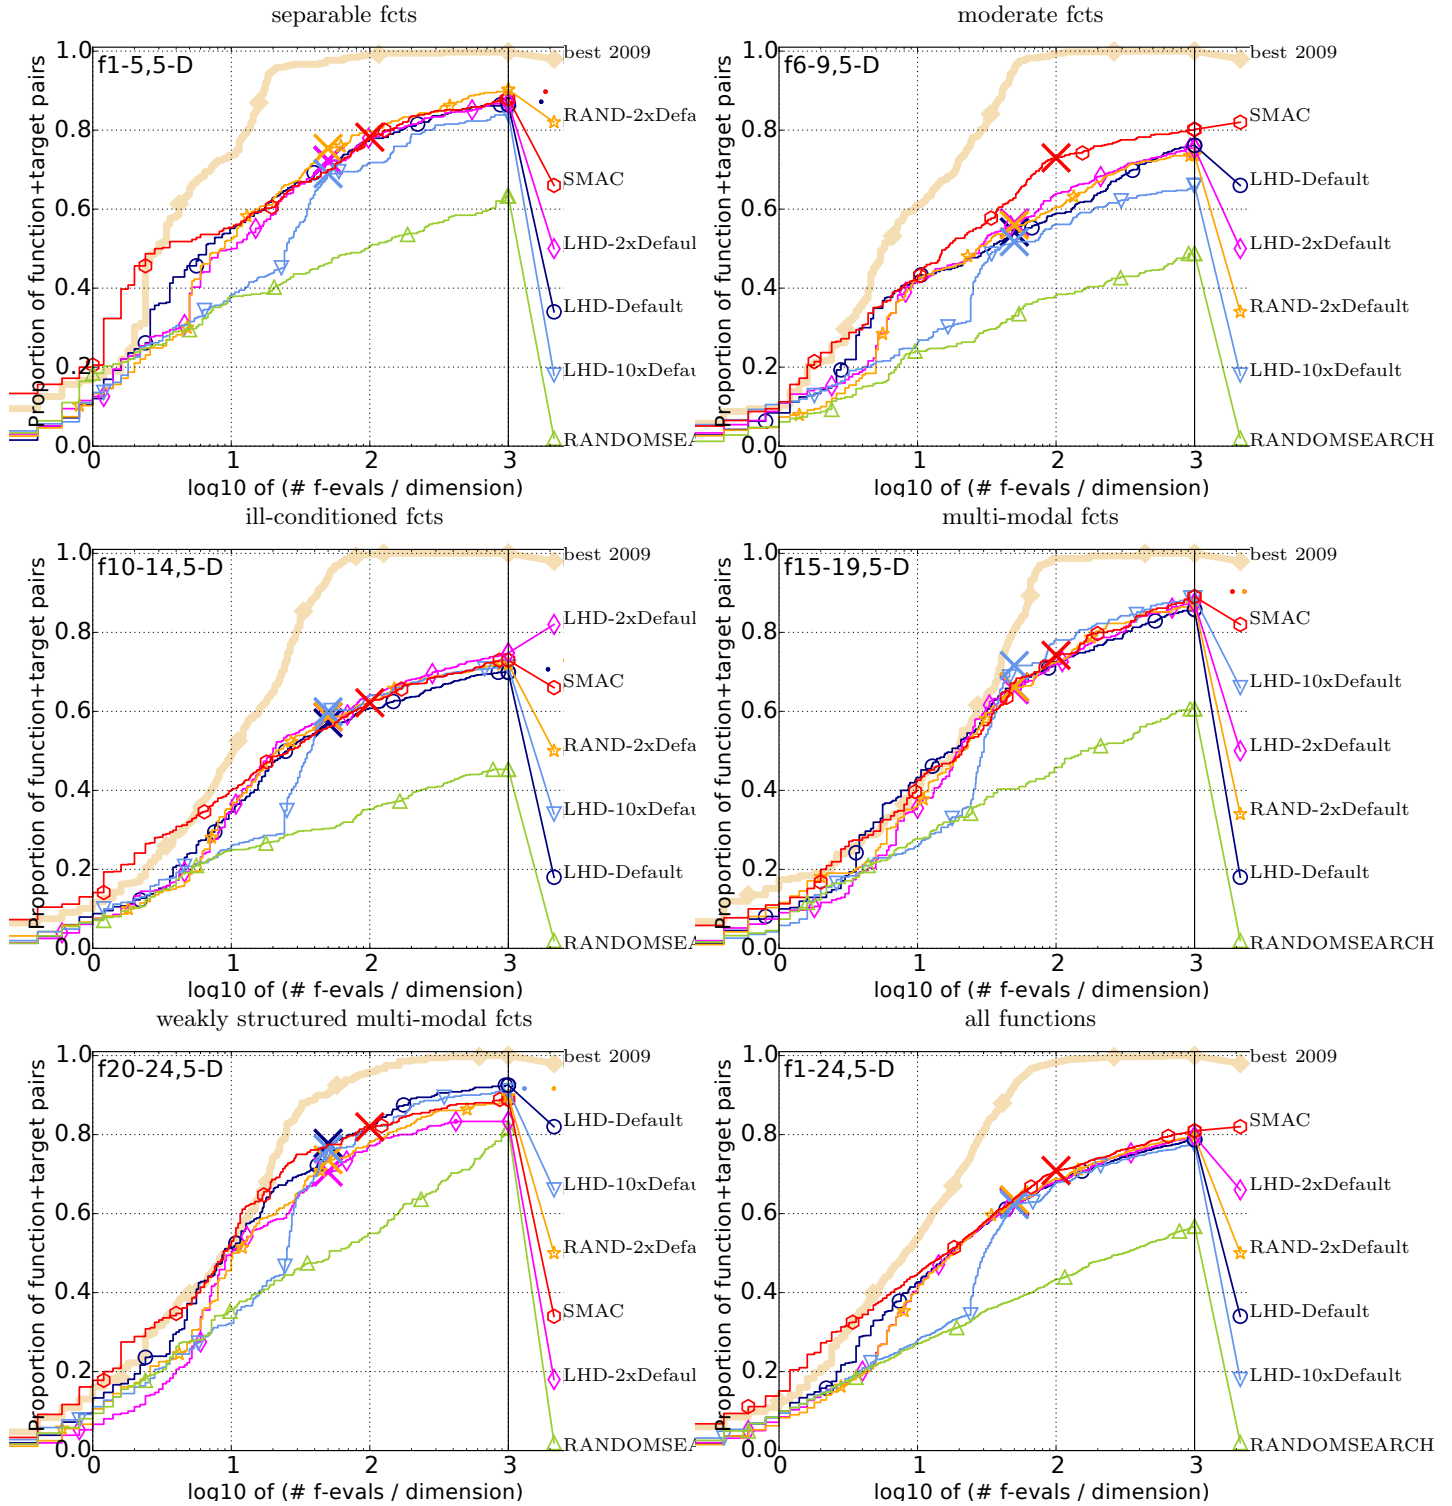


Figure 2: Bootstrapped empirical cumulative distribution of the number of objective function evaluations divided by dimension (FEvals/DIM) for all functions and subgroups in 5-D. The targets are chosen from  $10^{[-8..2]}$  such that the bestGECCO2009 artificial algorithm just not reached them within a given budget of  $k \times \text{DIM}$ , with  $k \in \{0.5, 1.2, 3, 10, 50\}$ . The “best 2009” line corresponds to the best ERT observed during BBOB 2009 for each selected target.

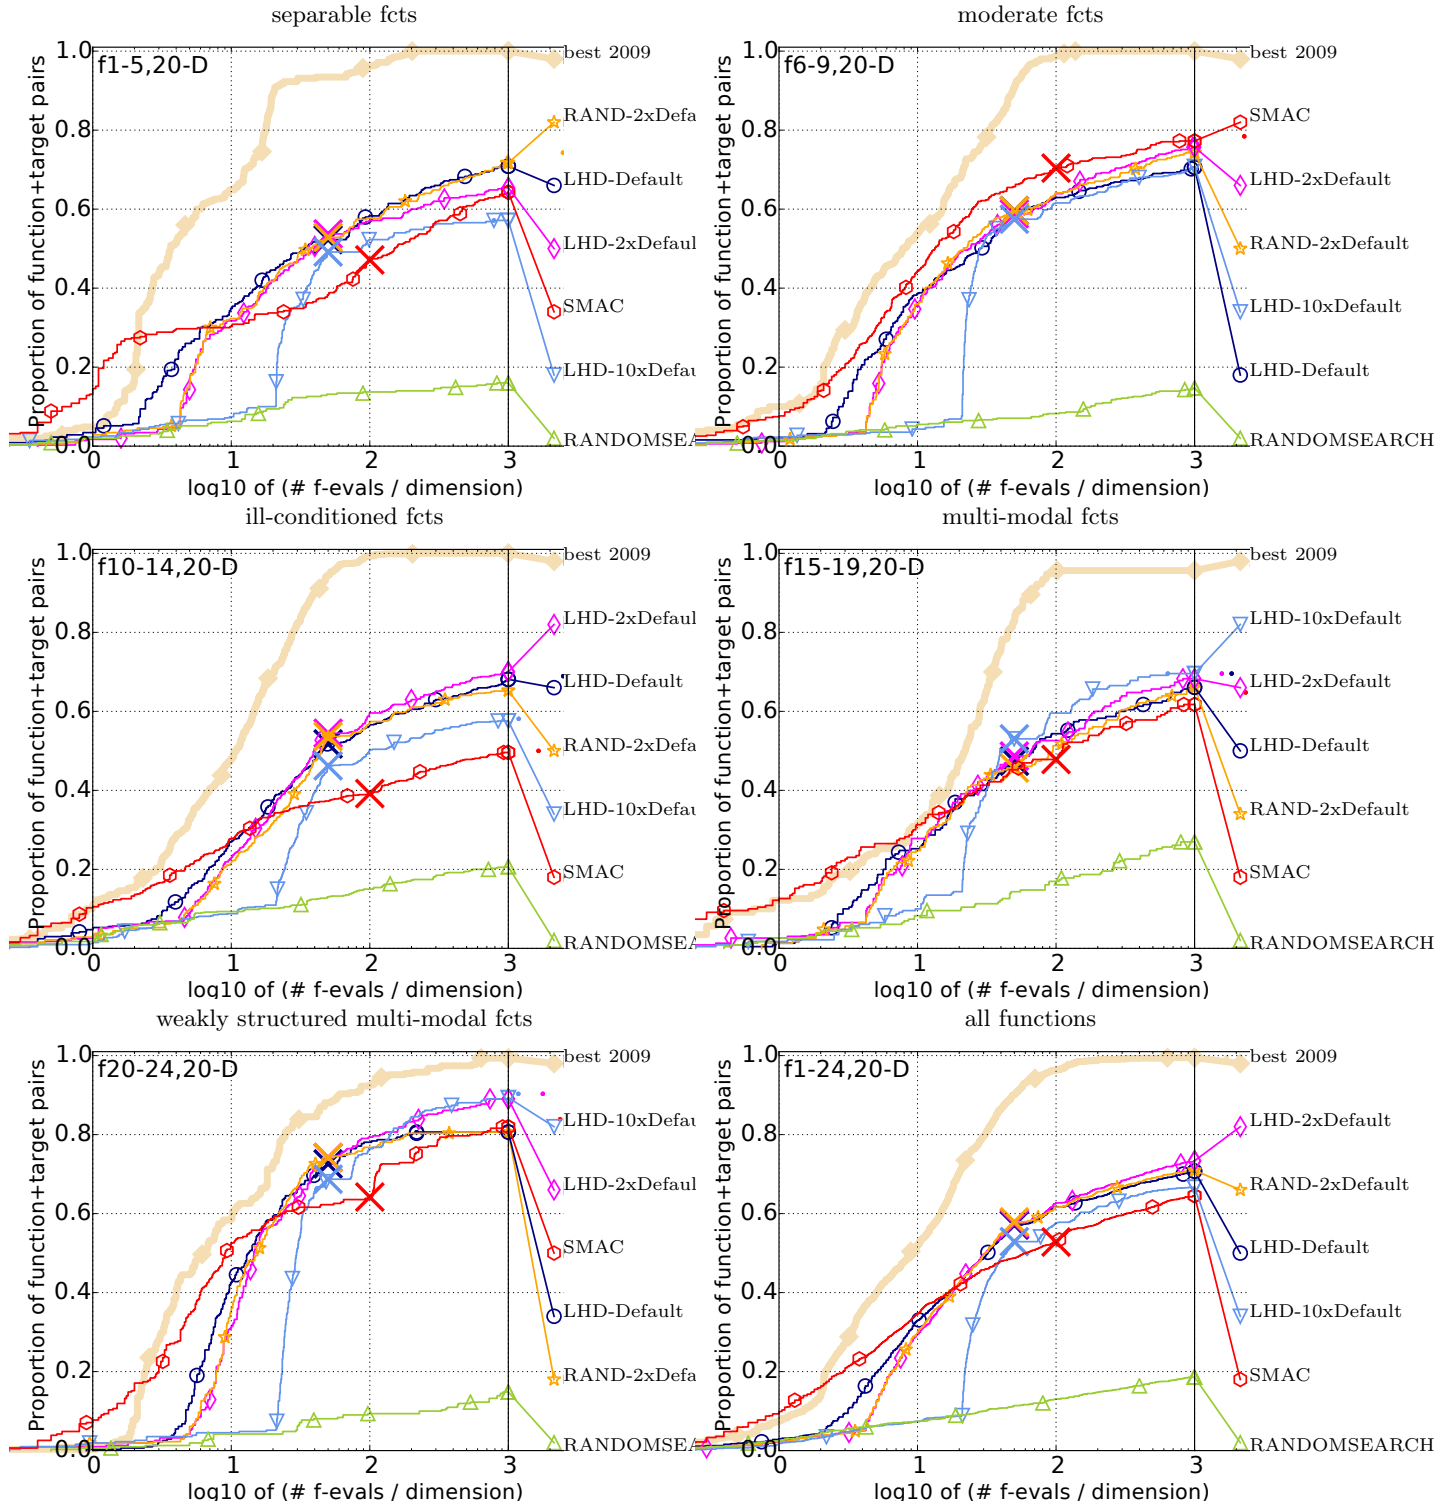


Figure 3: Bootstrapped empirical cumulative distribution of the number of objective function evaluations divided by dimension (FEvals/DIM) for all functions and subgroups in 20-D. The targets are chosen from  $10^{[-8..2]}$  such that the bestGECCO2009 artificial algorithm just not reached them within a given budget of  $k \times \text{DIM}$ , with  $k \in \{0.5, 1.2, 3, 10, 50\}$ . The “best 2009” line corresponds to the best ERT observed during BBOB 2009 for each selected target.





#FEs/D											#FEs/D													
		0.5	1.2	3	10	50	#succ			0.5	1.2	3	10	50	#succ			0.5	1.2	3	10	50	#succ	
#FEs/D	f1	6.3e+1:24	4.0e+1:42	1.0e-8:43	1.0e-8:43	1.0e-8:43	15/15			f13	1.6e+3:28	1.0e+3:64	6.3e+2:79	4.0e+1:211	2.5e+0:1724	15/15			0.5	1.2	3	10	50	#succ
LHD-Default		2.5(0.5)	2.0(0.2)	$\infty$	$\infty$	$\infty$	0/15	LHD-Default		2.6(0.6)	1.8(0.3)	2.6(1)	8.0(4)	$\infty$	1000	0/15	LHD-Default		2.5e+1:15	1.6e+1:42	1.0e+1:75	1.6e+0:219	6.3e+4:1106	15/15
LHD-2xDefault		3.9(0.2)	2.5(0.3)	$\infty$	$\infty$	$\infty$	0/15	LHD-2xDefault		3.3(1)	2.2(0.2)	2.8(1)	9.4(13)	$\infty$	1000	0/15	LHD-2xDefault		5.7(2)	3.1(1)	2.9(1)	9.5(9)	$\infty$	1000
RAND-2xDefault		1.1(0.1)	1.9(0.8)	8.5(4)	72(34)	$\infty$	0/15	RAND-2xDefault		8.1(2)	3.7(1)	3.1(1)	7.1(4)	$\infty$	1000	0/15	RAND-2xDefault		8.8(2)	4.1(1)	3.3(1)	7.3(5)	$\infty$	1000
LHD-10xDefault		17(0.1)	10(0.1)	$\infty$	$\infty$	$\infty$	0/15	LHD-10xDefault		15(0.2)	6.0(0.2)	6.3(0.3)	8.5(9)	$\infty$	1000	0/15	LHD-10xDefault		25(7)	11(0.5)	6.9(1)	7.7(12)	$\infty$	1000
SMAC		<b>0.80(0.3)</b> * <sup>4</sup>	<b>0.67(0.2)</b> * <sup>4</sup>	$\infty$	$\infty$	$\infty$	0/15	SMAC		<b>0.81(0.3)</b> * <sup>2</sup>	<b>0.66(0.1)</b> * <sup>4</sup>	<b>0.84(0.1)</b> * <sup>4</sup>	<b>4.4(0.3)</b> * <sup>4</sup>	$\infty$	2000	0/15	SMAC		2.0(0.8)* <sup>2</sup>	3.3(4)	19(26)	$\infty$	$\infty$	2000
RANDOMSEARCH		§92(1135)	3.7e4(4e4)	$\infty$	$\infty$	$\infty$	0/15	RANDOMSEARCH		358(81)	1.5e5(2e5)	$\infty$	$\infty$	$\infty$	2e7	0/15	RANDOMSEARCH		155(127)	1839(1525)	4.5e4(5e4)	$\infty$	$\infty$	2e7

Table 4: Expected running time (ERT in number of function evaluations) divided by the respective best ERT measured during BBOB-2009 in dimension 20. The ERT and in braces, as dispersion measure, the half difference between 90 and 10%-tile of bootstrapped run lengths appear for each algorithm and run-length based target, the corresponding best ERT (preceded by the target  $\Delta f$ -value in *italics*) in the first row. #succ is the number of trials that reached the target value of the last column. The median number of conducted function evaluations is additionally given in *italics*, if the target in the last column was never reached. Entries, succeeded by a star, are statistically significantly better (according to the rank-sum test) when compared to all other algorithms of the table, with  $p = 0.05$  or  $p = 10^{-k}$  when the number  $k$  following the star is larger than 1, with Bonferroni correction by the number of instances.