

# BBOB-Benchmarking a Simple Estimation of Distribution Algorithm with Cauchy Distribution

Petr Pošík

Czech Technical University in Prague, Faculty of Electrical Engineering, Dept. of Cybernetics  
Technická 2, 166 27 Prague 6  
posik@labe.felk.cvut.cz

## ABSTRACT

The restarted estimation of distribution algorithm (EDA) with Cauchy distribution as the probabilistic model is tested on the BBOB 2009 testbed. These tests prove that when using the Cauchy distribution and suitably chosen variance enlargement factor, the algorithm is usable for broad range of fitness landscapes, which is not the case for EDA with Gaussian distribution which converges prematurely. The results of the algorithm are of mixed quality and its scaling is at least quadratic.

## 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

## General Terms

Algorithms, Experimentation, Performance, Reliability

## Keywords

Benchmarking, Black-box optimization, Local search, Estimation-of-distribution algorithms, Evolutionary computation, Cauchy distribution

## 1. INTRODUCTION

Estimation of distribution algorithms (EDAs) [11] are a class of evolutionary algorithms (EAs) that do not use the crossover and mutation operators to create the offspring population. Instead, they build a probabilistic model describing the distribution of promising individuals and create offspring by sampling from the model. In real-valued spaces, such an algorithm can have a simple structure which is depicted in Fig. 1.

If the Gaussian distribution is employed as the model of promising individuals ([10], [3], [16], [1]), and the parameters

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'09, July 8–12, 2009, Montréal Québec, Canada.  
Copyright 2009 ACM 978-1-60558-505-5/09/07 ...\$5.00.

- 
1. Initialize the parameters  $\mu^0 = (\mu_1^0, \dots, \mu_D^0)$ ,  $\sigma^0 = (\sigma_1^0, \dots, \sigma_D^0)$ , and  $R^0 \in \mathcal{R}^{D \times D}$ ,  $R$  is orthonormal,  $D$  is the dimensionality of the search space. Generation counter  $t = 0$ .
  2. Sample  $N$  offspring from the search distribution (use  $R^t$  as a rotation matrix containing the base vectors,  $\sigma^t$  as relative scaling factors of individual components, and  $\mu^t$  as the distribution center).
  3. Evaluate the individuals.
  4. Select the  $\tau N$  best solutions (truncation selection).
  5. Estimate new  $\mu^{t+1}$ ,  $\sigma^{t+1}$ , and  $R^{t+1}$  using the selected individuals.
  6. Enlarge the  $\sigma^{t+1}$  by a constant factor  $k$  (global step size).
  7. Advance generation counter:  $t = t + 1$ .
  8. If termination condition is not met, go to step 2.
- 

Figure 1: Simple EDA used in this article.

of the distribution,  $\mu$ ,  $\sigma$ , and  $R$ , are learned by maximum likelihood (ML) estimation, the algorithm is very prone to premature convergence (i.e. the population converges on the slope of the fitness function) as recognized by many authors (see e.g. [3], [17], [12]). In [7], it was shown also theoretically that the distance traversed by a simple Gaussian EDA with truncation selection is bounded, and [5] showed similar results for tournament selection.

Many techniques that fight the premature convergence were developed, usually by means of artificially enlarging the ML estimate of variance of the learned distribution. In [17] it is suggested to use standard deviation greater than 1 when sampling the base Gaussian distribution (before it is multiplied by  $R \times \text{diag}(\sigma)$ ), e.g. to use  $\mathcal{G}(0, 1.5)$ , which amounts to use the variance enlargement factor  $k = 1.5$ . Adaptive variance scaling (AVS), i.e. enlarging the variance when better solutions were found and shrinking the variance in case of no improvement, was used along with various techniques to trigger the AVS only on the slope of the fitness function in [6] and [2]. The algorithm in Fig. 1, that suggests enlarging the population variance by a constant factor each generation, was studied in [18] where the minimal values of the ‘amplification coefficient’ were determined by a search in 1D case.

In [13], the theoretical model of the algorithm behavior in 1D was used to derive the minimal and maximal admissible values for  $k$ . However, in [15] it was shown experimentally that a constant multiplier does not ensure the desired properties of the algorithm when increasing the dimensionality of the search space.

In this article, we assess another premature convergence fighting technique suggested in [14], namely using the Cauchy distribution instead of Gaussian. This modification allows us to use a constant multiplier  $k$  which works for both the slope-like local neighborhood and the valley-like neighborhood in the same time.

## 2. ALGORITHM DESCRIPTION

The model parameters are computed actually using ML estimates for Gaussian distribution (even though they are subsequently used for the Cauchy distribution—they cannot be considered maximum-likelihood anymore, but they can serve as a heuristic). The distribution center  $\mu$  is computed as the average of the selected data points, the rotation matrix  $R$  and the standard deviations  $\sigma$  are obtained by eigendecomposition of the covariance matrix of the selected data points.

The sampling process can be described as follows. First, for each  $i$ -th offspring a Cauchy distributed number  $r_i$  is sampled; it will be the radius, i.e. the distance of the offspring from the origin (for a negative radius, absolute value is used, even though it is not necessary). Then, the values of  $r_i$  are divided by the  $\frac{1+\tau}{2}$ -quantile of the Cauchy distribution ( $\tau = 0.3$  is the used selection proportion of the truncation selection). It is just a form of standardizing the distribution used in [14]. Now, for each offspring a random vector is sampled, uniformly distributed on the unit sphere, which is then multiplied by the radius  $r_i$ . This operation results in an isotropically distributed set of offspring vectors  $z_i$ . Based on  $z_i$ , the final offspring  $x_i$  are generated by

$$x_i = \mu + k \cdot R \times \text{diag}(\sigma) \times z_i, \quad (1)$$

where  $k$  is the constant multiplier.

### 2.1 Parameter Settings

The algorithm has three parameters which determine its behaviour:

1. the selection proportion ( $\tau = 0.3$  is used throughout this article),
2. the population size  $N$  (depends on the search space dimensionality), and
3. the variance enlarging factor  $k$  (also depends on the search space dimensionality).

By running a small experiment with the algorithm on the Rosenbrock's function it was observed, that the values of  $k$  and  $N$  shown in Fig. 1 give acceptable performance.

Based on these measurements, the following models for  $k$  and  $N$  were obtained:

$$N = 10^{1.05} \cdot D^{1.36} \quad (2)$$

$$k = \begin{cases} (0.3 + 0.25 \cdot D) \cdot \sqrt{D} & \text{if } D < 5, \\ (1.45 + 0.013 \cdot D) \cdot \sqrt{D} & \text{if } D \geq 5. \end{cases} \quad (3)$$

D	2	3	5	10	20
k	0.7	1.2	1.5	1.6	1.7
N	30	50	80	250	700

**Table 1: Suggested values of parameters  $k$  and  $N$  for various values of  $D$  and for selection proportion  $\tau = 0.3$ .**

In each iteration, a check is performed if the model standard deviations did not fall below an acceptable limit. The limit was set to  $10^{-10}$ , and if smaller values of  $\sigma_i$  are detected, they are artificially enlarged to the least acceptable value.

### 2.2 Box Constraints? No

Several experiments with using box constraints in the form  $\langle -5, 5 \rangle^D$  or  $\langle -6, 6 \rangle^D$  were carried out. Subjectively, the best (the most regular) results were produced when the algorithm was run completely unconstrained, i.e. no constraints are used for the results presented in the next section.

### 2.3 The Crafting Effort

The algorithm has no additional parameters so that no further tuning is necessary; the same setup is used for all benchmark functions—the crafting effort CrE = 0.

### 2.4 Invariance Properties

Cauchy EDA is *invariant with respect to translation and rotation* as can be seen in the next section on the graphs for unrotated–rotated versions of Ellipsoid (2 and 10), Rastrigin (3 and 15), or Rosenbrock (8 and 9) functions. The algorithm uses fitness values only in the rank-based (truncation) selection, thus it is also *invariant against order-preserving transformations of the fitness function*.

## 3. EXPERIMENTAL PROCEDURE

The standard experimental procedure of BBOB was adopted: the algorithm was run on 24 test functions, 5 instances each, 3 runs on each instance. Each run was finished

- after finding a solution with fitness difference  $\Delta f \leq 10^{-8}$ , or
- after performing more than  $5 \cdot 10^4 \times D$  function evaluations.

Each individual launch of the basic Cauchy EDA was interrupted (and the algorithm was restarted)

- after finding a solution with fitness difference  $\Delta f \leq 10^{-8}$ , or
- after performing more than the allowed number of function evaluations, or
- after the model converged too much, i.e. when  $\max_i |\sigma_i| < 10^{-8}$ .

## 4. RESULTS

Results from experiments according to [8] on the benchmark functions given in [4, 9] are presented in Figures 2 and 3 and in Table 2. The method solved at least 1 run (out of

15 runs) for 21 (2D), 16 (3D), 16 (5D), 15 (10D), and 10 (20D) out of 24 functions, respectively.

As it turned out, the effect of the restarting mechanism is not large since the vast majority of runs ended up with only 1 launch, i.e. the run either found the solution or was stopped by the maximum allowed number of evaluations. The third stopping criterion—stop if all  $\sigma$  are below certain threshold—was used only seldom.

## 5. CPU TIMING EXPERIMENT

The multistart algorithm was run with the maximal number of evaluations set to  $10^5$ , the basic algorithm was restarted for at least 30 seconds. The average time demands were  $5.1 \cdot 10^{-5}$ ,  $1.7 \cdot 10^{-5}$ ,  $9.4 \cdot 10^{-6}$ ,  $8.6 \cdot 10^{-6}$ , and  $1.1 \cdot 10^{-5}$  seconds per function evaluation for 2-, 3-, 5-, 10-, and 20-dimensional search space, respectively. The experiment was conducted on Intel Core 2 CPU, T5600, 1.83 GHz, 1 GB RAM with Windows XP SP3 in MATLAB R2007b.

## 6. CONCLUSIONS

The restarted EDA with Cauchy distribution turned out to be effective for unimodal and ill-conditioned functions. On the other hand, its performance is bad when solving multimodal functions or functions with weak structure. Since it needs large number of function evaluations, it is likely that the restarting mechanism has only minor effect since only a few restarts were actually performed.

## Acknowledgements

The author is supported by the Grant Agency of the Czech Republic with the grant no. 102/08/P094 entitled “Machine learning methods for solution construction in evolutionary algorithms”.

## 7. REFERENCES

- [1] C. W. Ahn, R. S. Ramakrishna, and D. E. Goldberg. Real-coded bayesian optimization algorithm: Bringing the strength of BOA into the continuous world. In *Proceedings of the Genetic and Evolutionary Computation Conference—GECCO 2004*, pages 840–851. Springer-Verlag, Berlin, 2004.
- [2] P. A. N. Bosman, J. Grahl, and F. Rothlauf. SDR: A better trigger for adaptive variance scaling in normal EDAs. In *GECCO '07: Proceedings of the 9th annual conference on Genetic and Evolutionary Computation*, pages 492–499, New York, NY, USA, 2007. ACM Press.
- [3] P. A. N. Bosman and D. Thierens. Expanding from discrete to continuous estimation of distribution algorithms: The IDEA. In *PPSN VI: Proceedings of the 6th International Conference on Parallel Problem Solving from Nature*, pages 767–776, London, UK, 2000. Springer-Verlag.
- [4] 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.
- [5] C. Gonzales, J. A. Lozano, and P. Larrañaga. Mathematical modelling of UMDAc algorithm with tournament selection. *International Journal of Approximate Reasoning*, 31(3):313–340, 2002.
- [6] J. Grahl, P. A. N. Bosman, and F. Rothlauf. The correlation-triggered adaptive variance scaling IDEA. In *Proceedings of the 8th annual conference on Genetic and Evolutionary Computation Conference – GECCO 2006*, pages 397–404, New York, NY, USA, 2006. ACM Press.
- [7] J. Grahl, S. Minner, and F. Rothlauf. Behaviour of UMDAc with truncation selection on monotonous functions. In *IEEE Congress on Evolutionary Computation, CEC 2005*, volume 3, pages 2553–2559, 2005.
- [8] N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization benchmarking 2009: Experimental setup. Technical Report RR-6828, INRIA, 2009.
- [9] 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.
- [10] P. Larrañaga, J. A. Lozano, and E. Bengoetxea. Estimation of distribution algorithms based on multivariate normal distributions and Gaussian networks. Technical Report KZZA-IK-1-01, Dept. of Computer Science and Artificial Intelligence, University of Basque Country, 2001.
- [11] P. Larrañaga and J. A. Lozano, editors. *Estimation of Distribution Algorithms*. GENA. Kluwer Academic Publishers, 2002.
- [12] J. Ocenasek, S. Kern, N. Hansen, and P. Koumoutsakos. A mixed bayesian optimization algorithm with variance adaptation. In X. Yao, editor, *Parallel Problem Solving from Nature - PPSN VIII*, pages 352–361. Springer-Verlag, Berlin, 2004.
- [13] P. Pošík. Gaussian EDA and truncation selection: Setting limits for sustainable progress. In *IEEE SMC International Conference on Distributed Human-Machine Systems, DHMS 2008*, Athens, Greece, 2008. IEEE.
- [14] P. Pošík. Preventing premature convergence in a simple EDA via global step size setting. In *Parallel Problem Solving from Nature Ů- PPSN X*, volume 5199 of *Lecture Notes in Computer Science*, pages 549–558. Springer, 2008.
- [15] P. Pošík. Truncation selection and Gaussian EDA: Bounds for sustainable progress in high-dimensional spaces. In *EvoWorkshops 2008*, volume 4974 of *LNCS*, pages 525–534. Springer, 2008.
- [16] S. Rudlof and M. Köppen. Stochastic hill climbing by vectors of normal distributions. In *First Online Workshop on Soft Computing*, Nagoya, Japan, 1996.
- [17] B. Yuan and M. Gallagher. On the importance of diversity maintenance in estimation of distribution algorithms. In H. G. Beyer and U. M. O’Reilly, editors, *Proceedings of the Genetic and Evolutionary Computation Conference GECCO-2005*, volume 1, pages 719–726, New York, NY, USA, 2005. ACM Press.
- [18] B. Yuan and M. Gallagher. A mathematical modelling technique for the analysis of the dynamics of a simple continuous EDA. In *IEEE Congress on Evolutionary Computation, CEC 2006*, pages 1585–1591, Vancouver, Canada, 2006. IEEE Press.

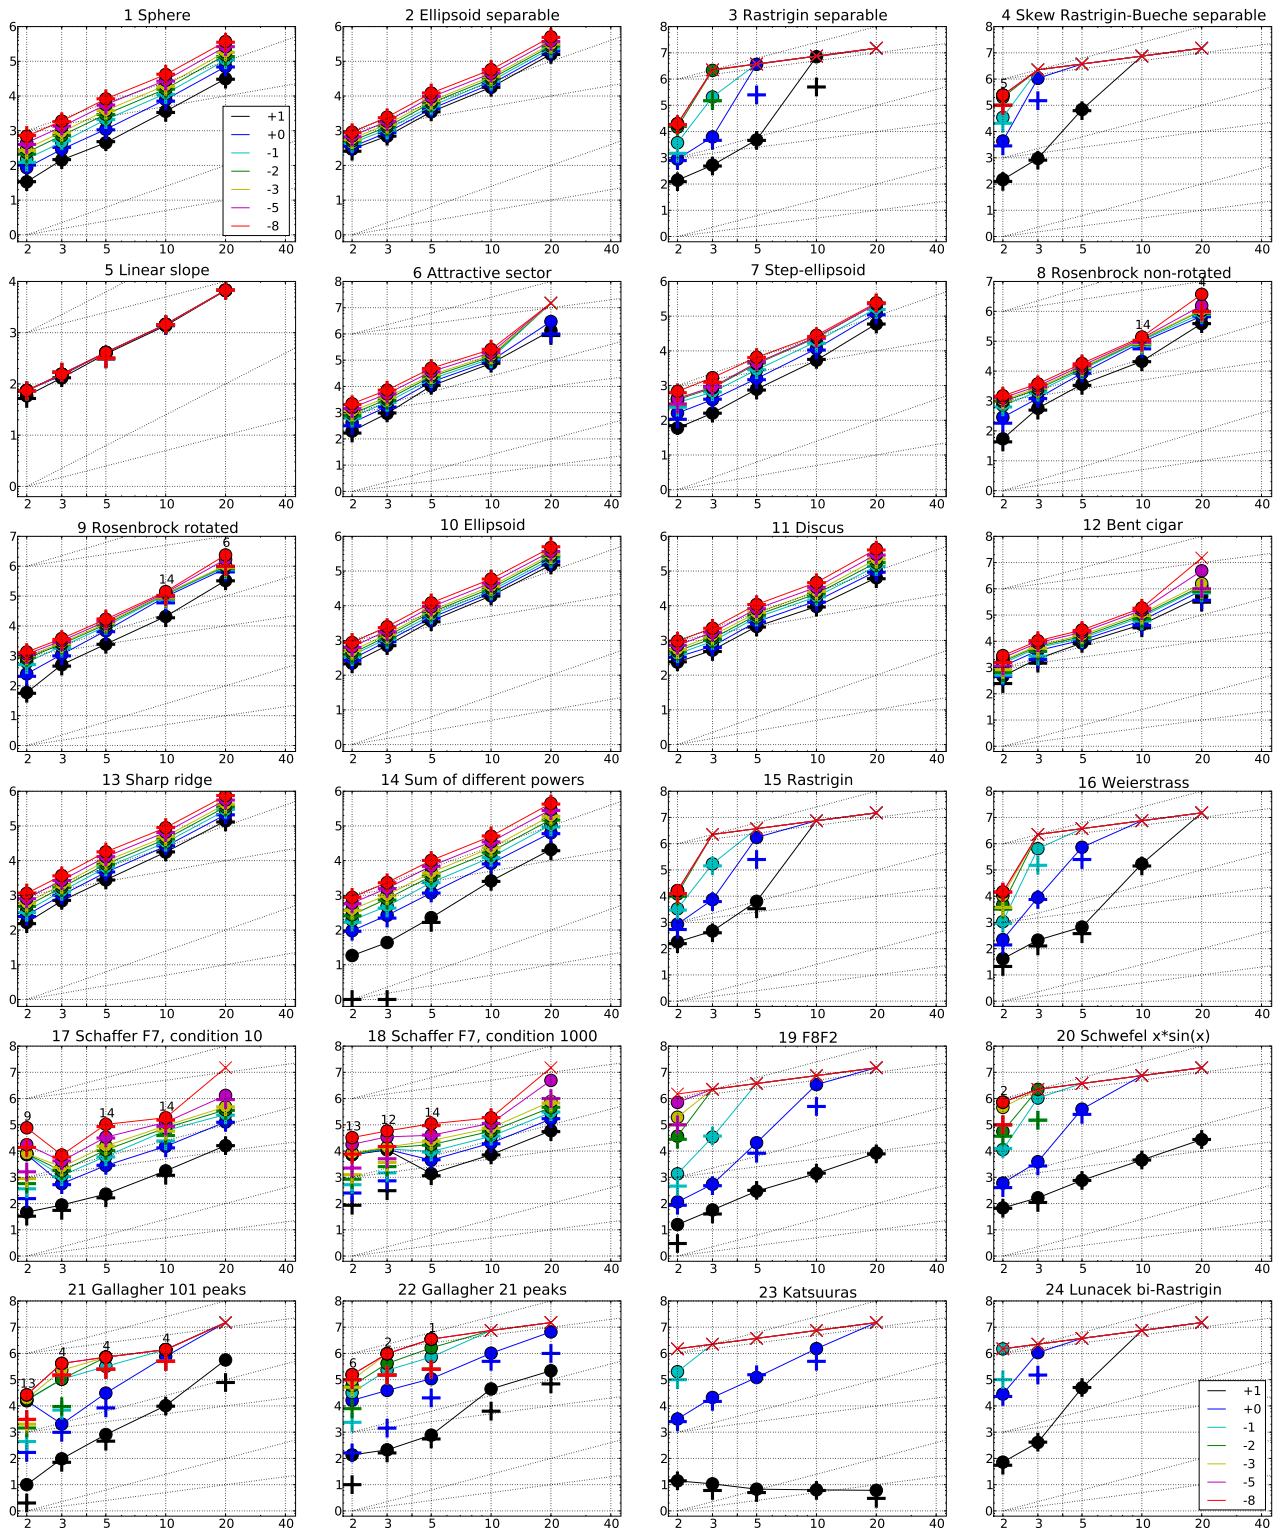


Figure 2: Expected Running Time (ERT,  $\bullet$ ) to reach  $f_{\text{opt}} + \Delta f$  and median number of function evaluations of successful trials (+), shown for  $\Delta f = 10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-8}$  (the exponent is given in the legend of  $f_1$  and  $f_{24}$ ) versus dimension in log-log presentation. The  $\text{ERT}(\Delta f)$  equals to  $\#FEs(\Delta f)$  divided by the number of successful trials, where a trial is successful if  $f_{\text{opt}} + \Delta f$  was surpassed during the trial. The  $\#FEs(\Delta f)$  are the total number of function evaluations while  $f_{\text{opt}} + \Delta f$  was not surpassed during the trial from all respective trials (successful and unsuccessful), and  $f_{\text{opt}}$  denotes the optimal function value. Crosses ( $\times$ ) indicate the total number of function evaluations  $\#FEs(-\infty)$ . Numbers above ERT-symbols indicate the number of successful trials. Annotated numbers on the ordinate are decimal logarithms. Additional grid lines show linear and quadratic scaling.



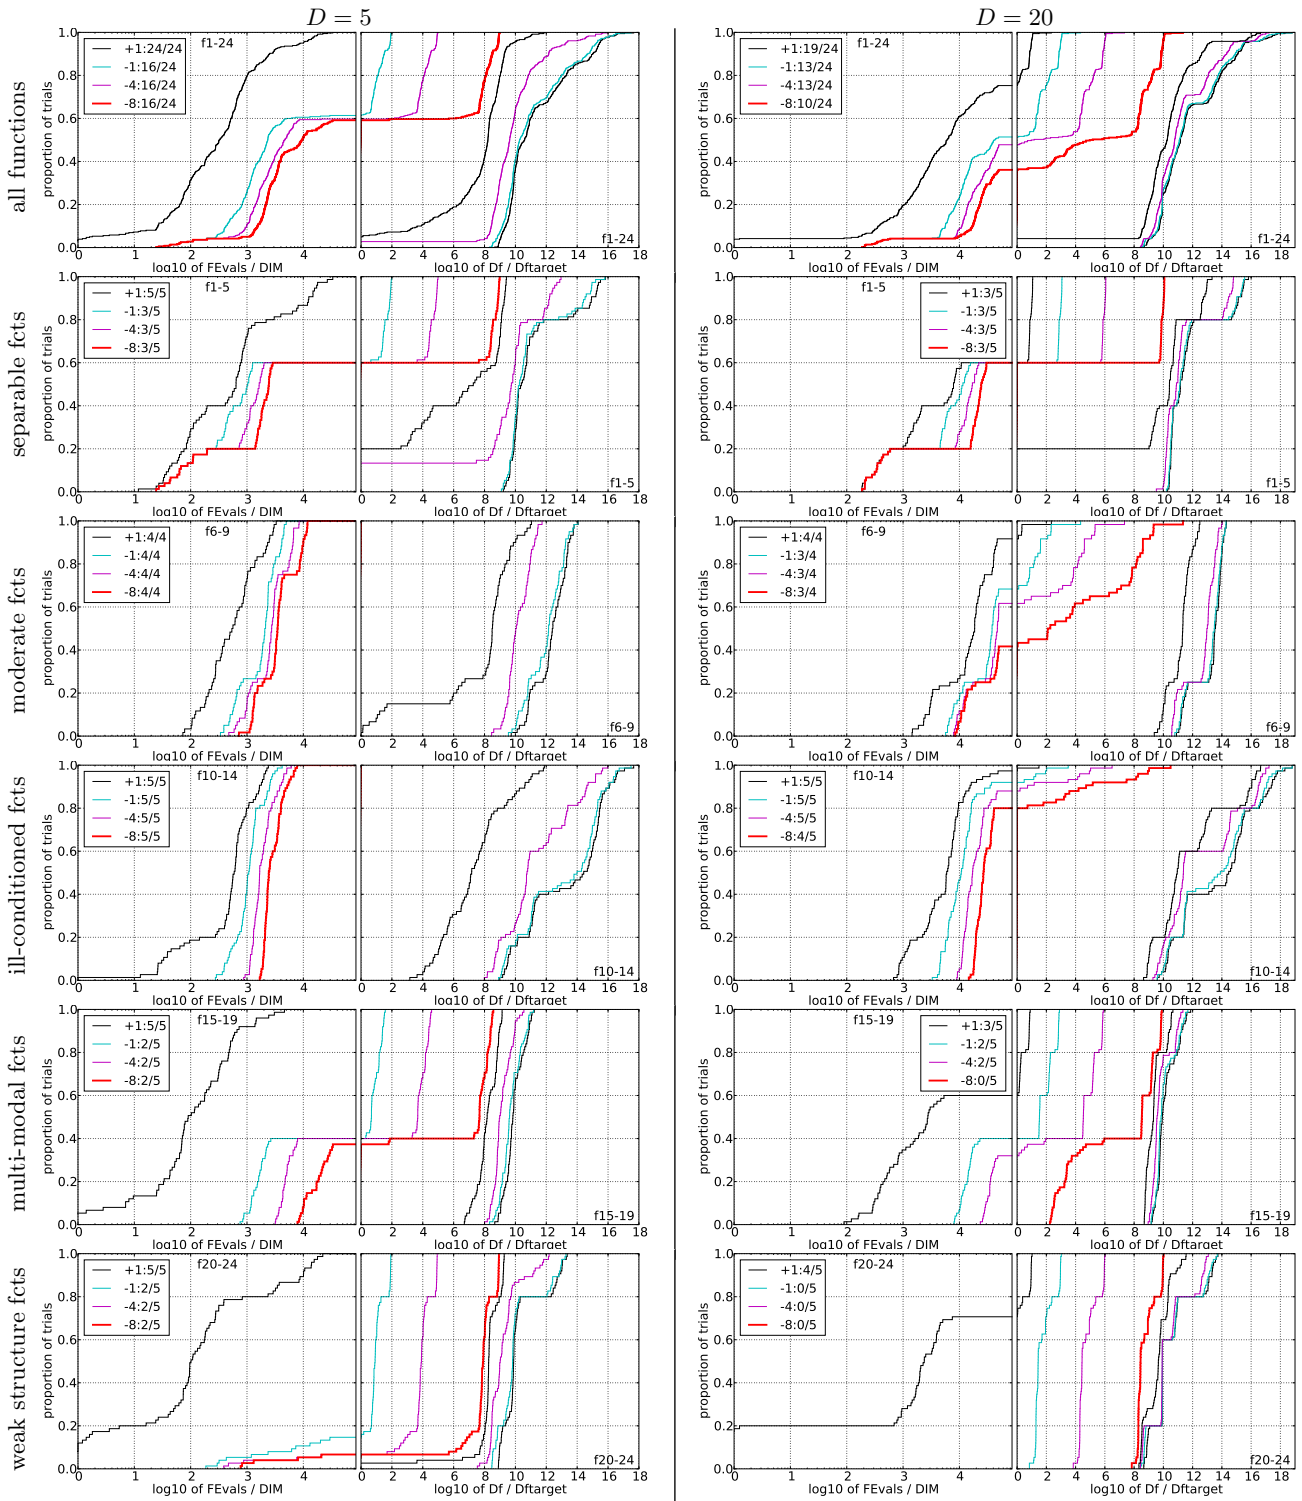


Figure 3: Empirical cumulative distribution functions (ECDFs), plotting the fraction of trials versus running time (left) or  $\Delta f$ . Left subplots: ECDF of the running time (number of function evaluations), divided by search space dimension  $D$ , to fall below  $f_{\text{opt}} + \Delta f$  with  $\Delta f = 10^k$ , where  $k$  is the first value in the legend. Right subplots: ECDF of the best achieved  $\Delta f$  divided by  $10^k$  (upper left lines in continuation of the left subplot), and best achieved  $\Delta f$  divided by  $10^{-8}$  for running times of  $D, 10D, 100D \dots$  function evaluations (from right to left cycling black-cyan-magenta). Top row: all results from all functions; second row: separable functions; third row: misc. moderate functions; fourth row: ill-conditioned functions; fifth row: multi-modal functions with adequate structure; last row: multi-modal functions with weak structure. The legends indicate the number of functions that were solved in at least one trial. FEvals denotes number of function evaluations,  $D$  and DIM denote search space dimension, and  $\Delta f$  and Df denote the difference to the optimal function value.