

Benchmarking Natural Evolution Strategies with Adaptation Sampling on the Noiseless and Noisy Black-box Optimization Testbeds

Tom Schaul

Courant Institute of Mathematical Sciences, New York University
Broadway 715, New York, USA
schaul@cims.nyu.edu

ABSTRACT

Natural Evolution Strategies (NES) are a recent member of the class of real-valued optimization algorithms that are based on adapting search distributions. Exponential NES (xNES) are the most common instantiation of NES, and particularly appropriate for the BBOB 2012 benchmarks, given that many are non-separable, and their relatively small problem dimensions. The technique of *adaptation sampling*, which adapts learning rates online further improves the algorithm’s performance. This report provides the most extensive empirical results on that combination (xNES-as) to date, on both the noise-free and noisy BBOB testbeds.

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

Keywords

Evolution Strategies, Natural Gradient, Benchmarking

1. INTRODUCTION

Evolution strategies (ES), in contrast to traditional evolutionary algorithms, aim at repeating the type of mutation that led to those good individuals. We can characterize those mutations by an explicitly parameterized *search distribution* from which new candidate samples are drawn, akin to estimation of distribution algorithms (EDA). Covariance matrix adaptation ES (CMA-ES [10]) innovated the field by introducing a parameterization that includes the full covariance matrix, allowing them to solve highly non-separable problems.

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’12 Companion, July 7–11, 2012, Philadelphia, PA, USA.
Copyright 2012 ACM 978-1-4503-1178-6/12/07 ...\$10.00.

A more recent variant, *natural evolution strategies* (NES [17, 6, 15, 16]) aims at a higher level of generality, providing a procedure to update the search distribution’s parameters for any type of distribution, by ascending the gradient towards higher expected fitness. Further, it has been shown [12, 11] that following the *natural gradient* to adapt the search distribution is highly beneficial, because it appropriately normalizes the update step with respect to its uncertainty and makes the algorithm scale-invariant.

Exponential NES (xNES), the most common instantiation of NES, used a search distribution parameterized by a mean vector and a full covariance matrix, and is thus most similar to CMA-ES (in fact, the precise relation is described in [4] and [5]). Given the relatively small problem dimensions of the BBOB benchmarks, and the fact that many are non-separable, it is also among the most appropriate NES variants for the task.

In this report, we retain the original formulation of xNES (including all parameter settings, except for an added stopping criterion), augmented with a technique for the online adaptation of its learning rate called *adaptation sampling*, which is designed to speed up convergence. We describe the empirical performance on all 54 benchmark functions (both noise-free and noisy) of the BBOB 2012 workshop.

2. NATURAL EVOLUTION STRATEGIES

Natural evolution strategies (NES) maintain a search distribution π and adapt the distribution parameters θ by following the *natural gradient* [1] of expected fitness J , that is, maximizing

$$J(\theta) = \mathbb{E}_\theta[f(\mathbf{z})] = \int f(\mathbf{z}) \pi(\mathbf{z} | \theta) d\mathbf{z}$$

Just like their close relative CMA-ES [10], NES algorithms are invariant under monotone transformations of the fitness function and linear transformations of the search space. Each iteration the algorithm produces n samples $\mathbf{z}_i \sim \pi(\mathbf{z} | \theta)$, $i \in \{1, \dots, n\}$, i.i.d. from its search distribution, which is parameterized by θ . The gradient w.r.t. the parameters θ can be rewritten (see [17]) as

$$\nabla_\theta J(\theta) = \nabla_\theta \int f(\mathbf{z}) \pi(\mathbf{z} | \theta) d\mathbf{z} = \mathbb{E}_\theta [f(\mathbf{z}) \nabla_\theta \log \pi(\mathbf{z} | \theta)]$$

from which we obtain a Monte Carlo estimate

$$\nabla_\theta J(\theta) \approx \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) \nabla_\theta \log \pi(\mathbf{z}_i | \theta)$$

of the search gradient. The key step then consists in replacing this gradient by the natural gradient defined as $\mathbf{F}^{-1}\nabla_\theta J(\theta)$ where $\mathbf{F} = \mathbb{E}[\nabla_\theta \log \pi(\mathbf{z}|\theta) \nabla_\theta \log \pi(\mathbf{z}|\theta)^\top]$ is the Fisher information matrix. The search distribution is iteratively updated using natural gradient ascent

$$\theta \leftarrow \theta + \eta \mathbf{F}^{-1} \nabla_\theta J(\theta)$$

with learning rate parameter η .

2.1 Exponential NES

While the NES formulation is applicable to arbitrary parameterizable search distributions [17, 11], the most common variant employs multinormal search distributions. For that case, two helpful techniques were introduced in [6], namely an exponential parameterization of the covariance matrix, which guarantees positive-definiteness, and a novel method for changing the coordinate system into a “natural” one, which makes the algorithm computationally efficient. The resulting algorithm, NES with a multivariate Gaussian search distribution and using both these techniques is called *xNES*, and the pseudocode is given in Algorithm 1.

Algorithm 1: Exponential NES (xNES)

input: f , μ_{init} , η_σ , η_B , u_k

μ	\leftarrow	μ_{init}	
initialize	σ	\leftarrow	1
	\mathbf{B}	\leftarrow	\mathbb{I}

repeat

for $k = 1 \dots n$ do
draw sample $\mathbf{s}_k \sim \mathcal{N}(0, \mathbb{I})$
$\mathbf{z}_k \leftarrow \mu + \sigma \mathbf{B}^\top \mathbf{s}_k$
evaluate the fitness $f(\mathbf{z}_k)$
end

sort $\{(\mathbf{s}_k, \mathbf{z}_k)\}$ with respect to $f(\mathbf{z}_k)$
and assign utilities u_k to each sample

compute gradients

$\nabla_\delta J \leftarrow \sum_{k=1}^n u_k \cdot \mathbf{s}_k$
$\nabla_M J \leftarrow \sum_{k=1}^n u_k \cdot (\mathbf{s}_k \mathbf{s}_k^\top - \mathbb{I})$
$\nabla_\sigma J \leftarrow \text{tr}(\nabla_M J)/d$
$\nabla_B J \leftarrow \nabla_M J - \nabla_\sigma J \cdot \mathbb{I}$

update parameters

$\mu \leftarrow \mu + \sigma \mathbf{B} \cdot \nabla_\delta J$
$\sigma \leftarrow \sigma \cdot \exp(\eta_\sigma/2 \cdot \nabla_\sigma J)$
$\mathbf{B} \leftarrow \mathbf{B} \cdot \exp(\eta_B/2 \cdot \nabla_B J)$

until stopping criterion is met

2.2 Adaptation Sampling

First introduced in [11] (chapter 2, section 4.4), *adaptation sampling* is a new meta-learning technique [14] that can adapt hyper-parameters online, in an economical way that is grounded on a measure statistical improvement.

Here, we apply it to the learning rate of the global step-size η_σ . The idea is to consider whether a larger learning-rate $\eta'_\sigma = \frac{3}{2}\eta_\sigma$ would have been more likely to generate the good samples in the current batch. For this we determine the (hypothetical) search distribution that would have resulted from such a larger update $\pi(\cdot|\theta')$. Then we compute

importance weights

$$w'_k = \frac{\pi(\mathbf{z}_k|\theta')}{\pi(\mathbf{z}_k|\theta)}$$

for each of the n samples \mathbf{z}_k in our current population, generated from the actual search distribution $\pi(\cdot|\theta)$. We then conduct a *weighted* Mann-Whitney test [11] (appendix A) to determine if the set $\{\text{rank}(\mathbf{z}_k)\}$ is inferior to its reweighted counterpart $\{w'_k \cdot \text{rank}(\mathbf{z}_k)\}$ (corresponding to the larger learning rate), with statistical significance ρ . If so, we increase the learning rate by a factor of $1 + c'$, up to at most $\eta_\sigma = 1$ (where $c' = 0.1$). Otherwise it decays to its initial value:

$$\eta_\sigma \leftarrow (1 - c') \cdot \eta_\sigma + c' \cdot \eta_{\sigma,\text{init}}$$

The procedure is summarized in algorithm 2 (for details and derivations, see [11]). The combination of xNES with adaptation sampling is dubbed *xNES-as*.

One interpretation of why adaptation sampling is helpful is that half-way into the search, (after a local attractor has been found, e.g., towards the end of the valley on the Rosenbrock benchmarks f_8 or f_9), the convergence speed can be boosted by an increased learning rate. For such situations, an online adaptation of hyper-parameters is inherently well-suited.

Algorithm 2: Adaptation sampling

input : $\eta_{\sigma,t}, \eta_{\sigma,\text{init}}$, θ_t , θ_{t-1} , $\{(\mathbf{z}_k, f(\mathbf{z}_k))\}$, c' , ρ

output: $\eta_{\sigma,t+1}$

compute hypothetical θ' , given θ_{t-1} and using $3/2\eta_{\sigma,t}$

for $k = 1 \dots n$ do
$w'_k = \frac{\pi(\mathbf{z}_k \theta')}{\pi(\mathbf{z}_k \theta)}$
end
$S \leftarrow \{\text{rank}(\mathbf{z}_k)\}$
$S' \leftarrow \{w'_k \cdot \text{rank}(\mathbf{z}_k)\}$
if weighted-Mann-Whitney(S, S') $< \rho$ then
return $(1 - c') \cdot \eta_\sigma + c' \cdot \eta_{\sigma,\text{init}}$
else
return $\min((1 + c') \cdot \eta_\sigma, 1)$
end

Table 1: Default parameter values for xNES (including the utility function and adaptation sampling) as a function of problem dimension d .

parameter	default value
n	$4 + \lfloor 3 \log(d) \rfloor$
$\eta_\sigma = \eta_B$	$\frac{3(3 + \log(d))}{5d\sqrt{d}}$
u_k	$\frac{\max(0, \log(\frac{n}{2} + 1) - \log(k))}{\sum_{j=1}^n \max(0, \log(\frac{n}{2} + 1) - \log(j))} - \frac{1}{n}$
ρ	$\frac{1}{2} - \frac{1}{3(d+1)}$
c'	$\frac{1}{10}$

3. EXPERIMENTAL SETTINGS

We use identical default hyper-parameter values for all benchmarks (both noisy and noise-free functions), which are taken from [6, 11]. Table 1 summarizes all the hyper-parameters used.

In addition, we make use of the provided target fitness f_{opt} to trigger *independent* algorithm restarts¹, using a simple ad-hoc procedure: If the log-progress during the past 1000d evaluations is too small, i.e., if

$$\log_{10} \left| \frac{f_{\text{opt}} - f_t}{f_{\text{opt}} - f_{t-1000d}} \right| < (r+2)^2 \cdot m^{3/2} \cdot [\log_{10} |f_{\text{opt}} - f_t| + 8]$$

where m is the remaining budget of evaluations divided by 1000d, f_t is the best fitness encountered until evaluation t and r is the number of restarts so far. The total budget is $10^5 d^{3/2}$ evaluations.

Implementations of this and other NES algorithm variants are available in Python through the PyBrain machine learning library [13], as well as in other languages at www.idsia.ch/~tom/nest.html.

4. CPU TIMING

A timing experiment was performed to determine the CPU-time per function evaluation, and how it depends on the problem dimension. For each dimension, the algorithm was restarted with a maximum budget of $10000/d$ evaluations, until at least 30 seconds had passed.

Our xNES-as implementation (in Python, based on the PyBrain [13] library), running on an Intel Xeon with 2.67GHz, required an average time of 1.1, 0.9, 0.7, 0.7, 0.9, 2.7 milliseconds per function evaluation for dimensions 2, 5, 10, 20, 40, 80 respectively (the function evaluations themselves take about 0.1ms).

5. RESULTS

Results of xNES-as on the noiseless testbed (from experiments according to [7] on the benchmark functions given in [2, 8]) are presented in Figures 1, 3 and 5, and in Table 2.

Similarly, results of xNES-as on the testbed of noisy functions (from experiments according to [7] on the benchmark functions given in [3, 9]) are presented in Figures 2, 4 and 5, and in Table 3.

6. DISCUSSION

The top rows in Figures 3 and 4 give a good overview picture, showing that across all benchmarks taken together, xNES-as performs almost as well as the best and better than most of the BBOB 2009 contestants. Beyond this high-level perspective, the results speak for themselves, of course, we will just highlight a few observations.

According to Tables 2 and 3, the only conditions where xNES-as significantly outperforms *all* algorithms from the BBOB2009 competition are on function f_{115} and in the early phase on f_{18} , f_{118} and f_{119} . We observe the worst performance on multimodal functions like f_3 , f_4 and f_{15} that other algorithms tackle very easily. Comparing different types of

¹It turns out that this use of f_{opt} is technically not permitted by the BBOB guidelines, so strictly speaking a different restart strategy should be employed, for example the one described in [11].

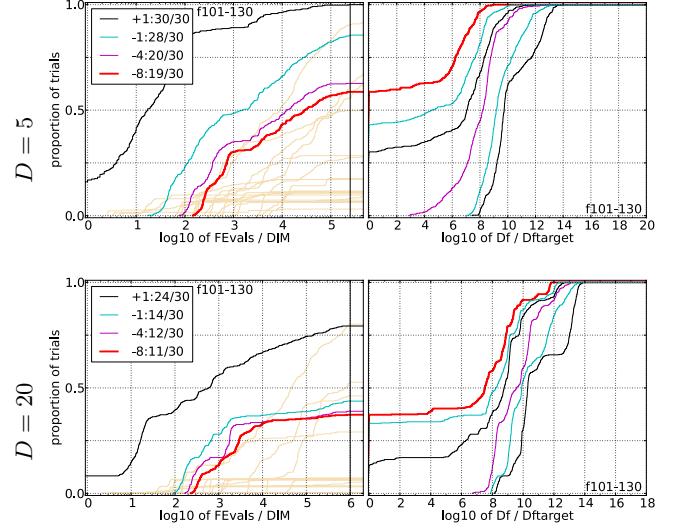


Figure 4: Empirical cumulative distribution functions (ECDFs) of the 30 noisy benchmark functions. Plotted is the fraction of trials versus running time (left subplots) or versus Δf (right subplots) (see Figure 3 for details).

noise, xNES-as appears to be least sensitive to Cauchy noise and most sensitive to uniform noise (see Figure 2).

From Figure 5 we observe a good loss ratio across the board on all benchmarks, with the best ones on moderate functions, ill-conditioned functions, and for all levels of noise. These results hold equally in dimensions 5 and 20. On the other hand, the algorithm is less competitive on (noisy or noise-free) multimodal benchmarks, which we expect to be directly related to its small default population size.

Acknowledgements

The author wants to thank the organizers of the BBOB workshop for providing such a well-designed benchmark setup, and especially such high-quality post-processing utilities.

This work was funded in part through AFR postdoc grant number 2915104, of the National Research Fund Luxembourg.

7. REFERENCES

- [1] S. I. Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10:251–276, 1998.
- [2] 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.
- [3] S. Finck, N. Hansen, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking 2010: Presentation of the noisy functions. Technical Report 2009/21, Research Center PPE, 2010.
- [4] N. Fukushima, Y. Nagata, S. Kobayashi, and I. Ono. Proposal of distance-weighted exponential natural evolution strategies. In *2011 IEEE Congress of Evolutionary Computation*, pages 164–171. IEEE, 2011.

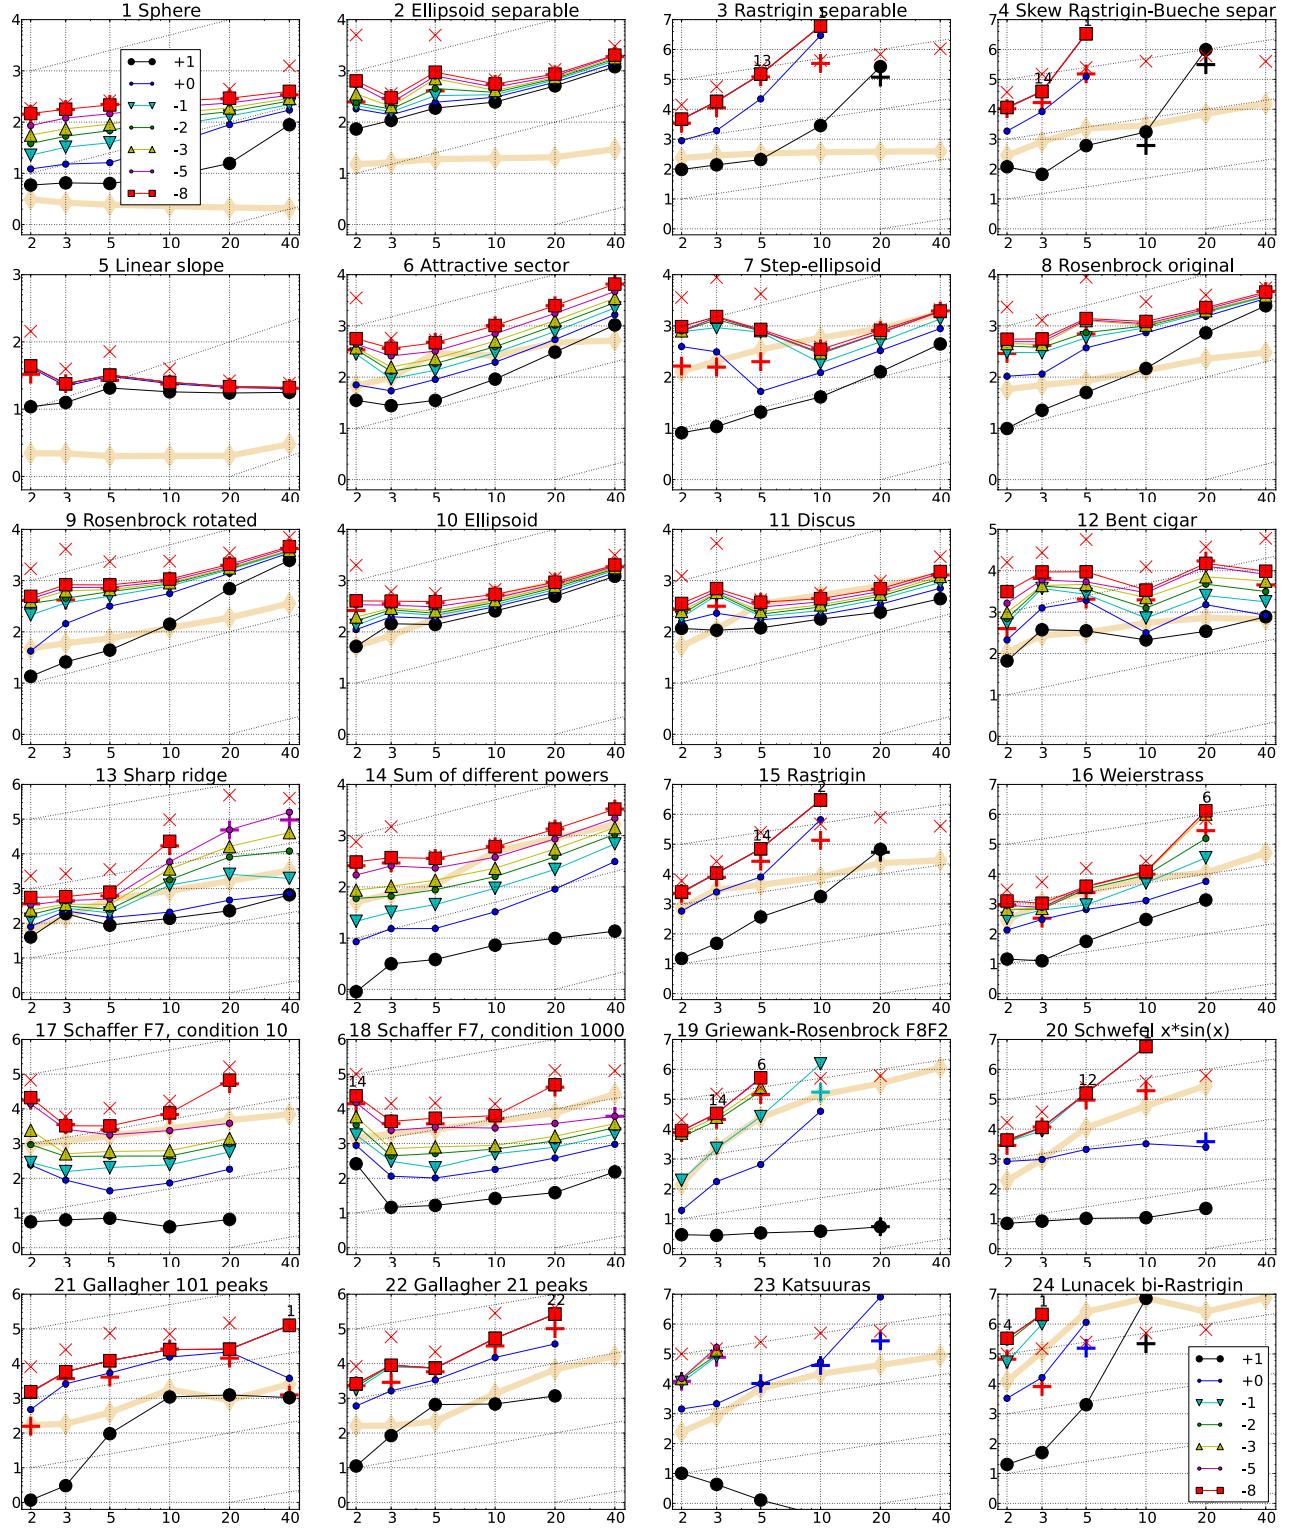


Figure 1: Expected number of f -evaluations (ERT, with lines, see legend) to reach $f_{\text{opt}} + \Delta f$, median number of f -evaluations to reach the most difficult target that was reached at least once (+) and maximum number of f -evaluations in any trial (x), all divided by dimension and plotted as \log_{10} values versus dimension. Shown are $\Delta f = 10^{\{1,0,-1,-2,-3,-5,-8\}}$. Numbers above ERT-symbols indicate the number of successful trials. The light thick line with diamonds indicates the respective best result from BBOB-2009 for $\Delta f = 10^{-8}$. Horizontal lines mean linear scaling, slanted grid lines depict quadratic scaling.

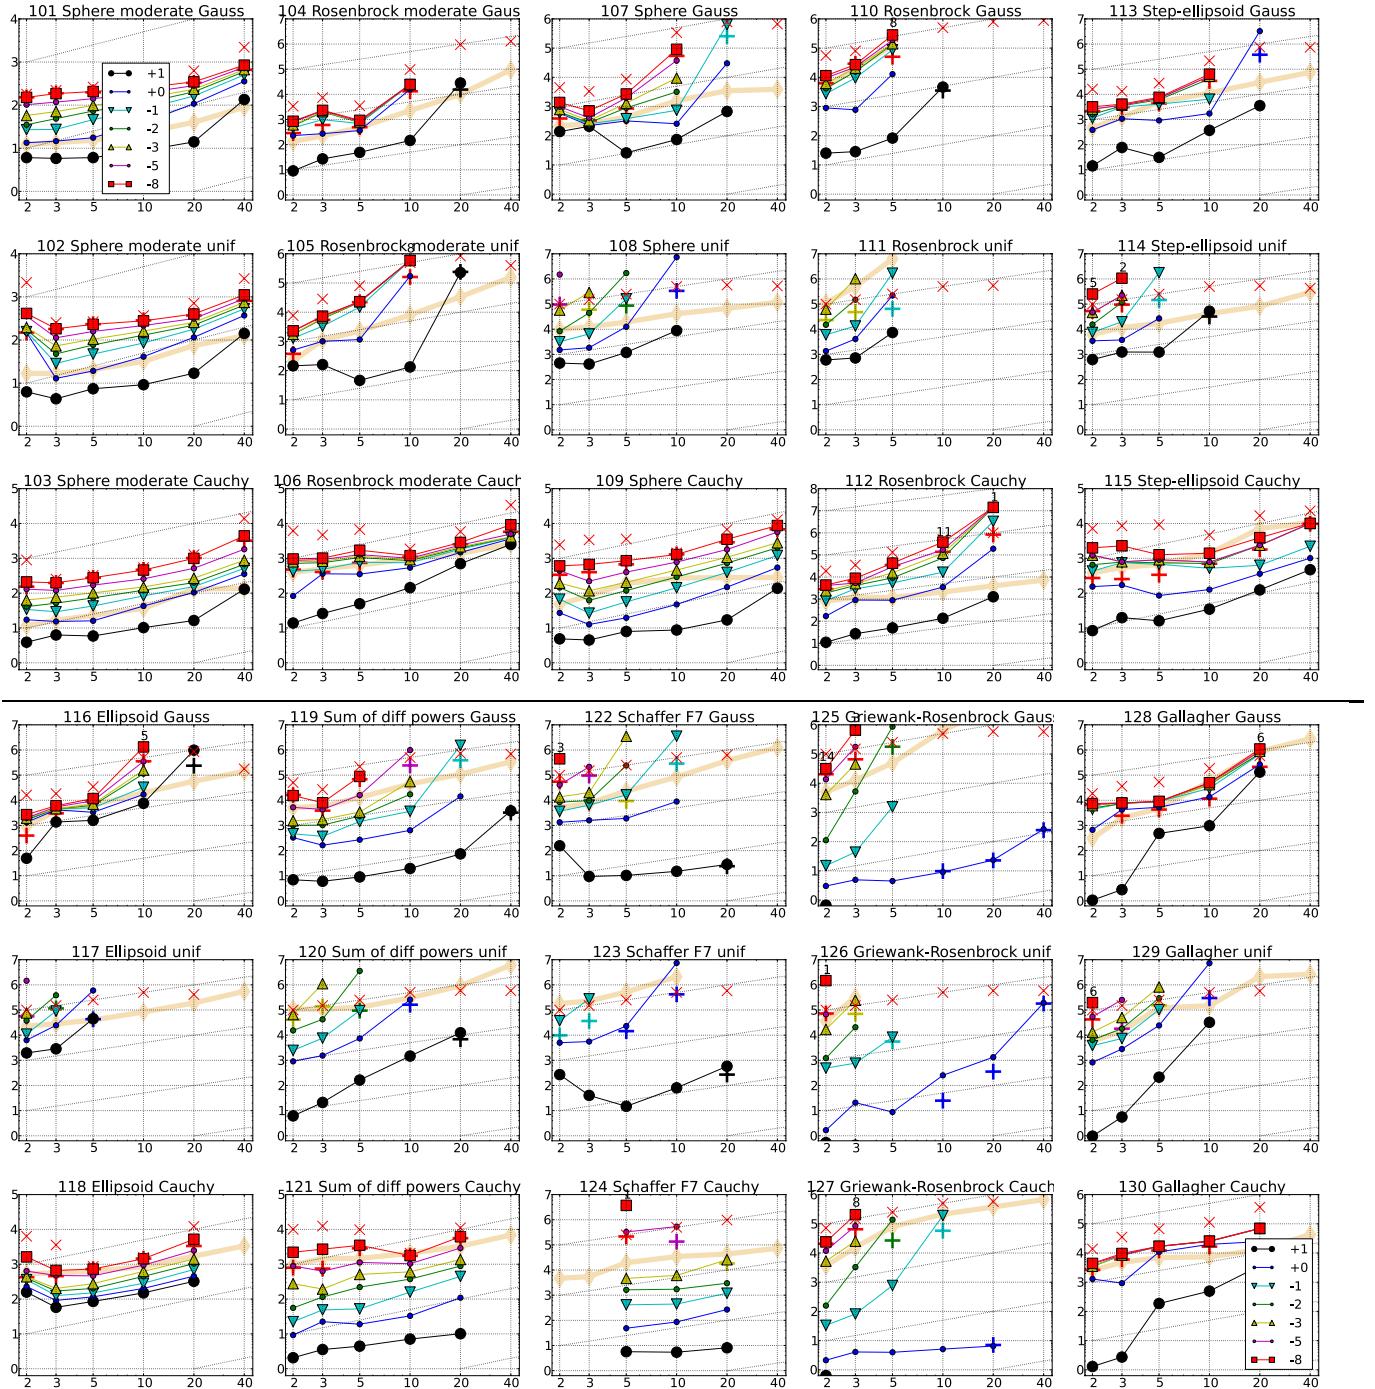


Figure 2: Expected number of f -evaluations (ERT, with lines, see legend) to reach $f_{\text{opt}} + \Delta f$, median number of f -evaluations to reach the most difficult target that was reached at least once (+) and maximum number of f -evaluations in any trial (\times), all divided by dimension and plotted as \log_{10} values versus dimension. Shown are $\Delta f = 10^{\{1,0,-1,-2,-3,-5,-8\}}$. Numbers above ERT-symbols indicate the number of successful trials. The light thick line with diamonds indicates the respective best result from BBOB-2009 for $\Delta f = 10^{-8}$. Horizontal lines mean linear scaling, slanted grid lines depict quadratic scaling.

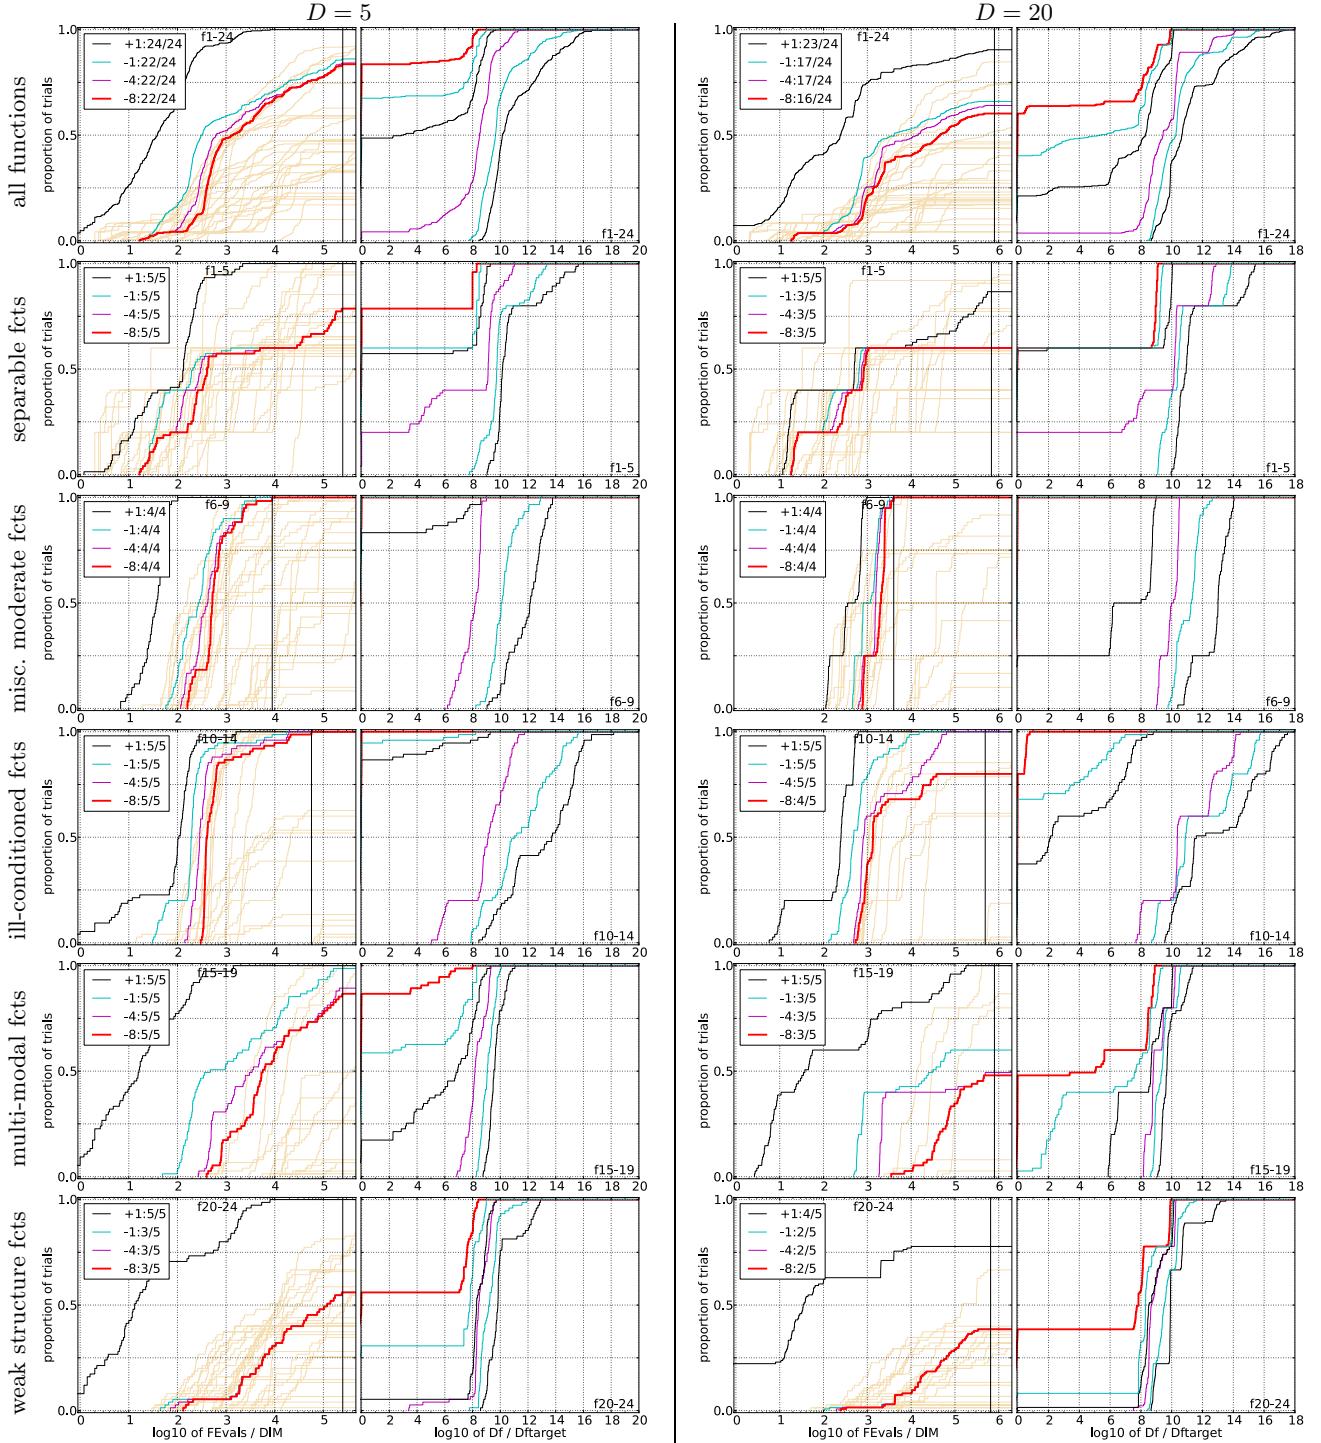


Figure 3: Empirical cumulative distribution functions (ECDFs), plotting the fraction of trials with an outcome not larger than the respective value on the x -axis. Left subplots: ECDF of number of function evaluations (FFEvals) divided by search space dimension D , to fall below $f_{\text{opt}} + \Delta f$ with $\Delta f = 10^k$, where k is the first value in the legend. Right subplots: ECDF of the best achieved Δf divided by 10^{-8} for running times of $D, 10D, 100D, \dots$ function evaluations (from right to left cycling black-cyan-magenta). The thick red line represents the most difficult target value $f_{\text{opt}} + 10^{-8}$. Legends indicate the number of functions that were solved in at least one trial. Light brown lines in the background show ECDFs for $\Delta f = 10^{-8}$ of all algorithms benchmarked during BBOB-2009.

Δf	1e+1	1e+0	1e-1	1e-3	1e-5	1e-7	#succ
f₁	11 2.9(2)	12 6.6(3)	12 16(5)	12 37(8)	12 60(12)	12 78(17)	15/15
f₂	83 11(5)	87 14(9)	88 19(18)	90 39(62)	92 43(63)	94 49(92)	15/15
f₃	716 1.5(0.7)	1622 69(80)	1637 454(375)	1646 452(384)	1650 451(377)	1654 450(382)	15/15
f₄	809 3.8(5)	1633 381(434)	1688 9998(10786)	1817 9287(9745)	1886 8949(9769)	1903 8868(8768)	15/15
f₅	10 10(4)	10 15(9)	10 16(9)	10 16(8)	10 16(8)	10 16(8)	15/15
f₆	114 1.5(1)	214 2.1(0.6)	281 2.4(0.6)	580 2.0(0.2)	1038 1.5(0.2)	1332 1.6(0.2)	15/15
f₇	24 4.4(3)	324 0.81(0.3)	1171 3.2(4)	1572 2.6(3)	1572 2.6(3)	1597 2.6(3)	15/15
f₈	73 3.4(2)	273 6.8(3)	336 8.7(4)	391 16(13)	410 16(13)	422 16(12)	15/15
f₉	35 6.4(2)	127 13(4)	214 12(3)	300 11(6)	335 11(6)	369 11(6)	15/15
f₁₀	349 2.0(0.8)	500 1.8(0.7)	574 1.8(0.6)	626 2.0(0.5)	829 1.9(0.4)	880 2.0(0.3)	15/15
f₁₁	143 4.2(3)	202 4.3(1)	763 1.3(0.3)	1177 1.1(0.2)	1467 1.0(0.1)	1673 1.1(0.1)	15/15
f₁₂	108 16(28)	268 36(66)	371 36(58)	461 51(97)	1303 21(35)	1494 31(34)	15/15
f₁₃	132 3.3(0.6)	195 3.7(0.5)	250 4.0(0.5)	1310 1.3(0.2)	1752 1.4(0.2)	2255 1.5(0.2)	15/15
f₁₄	10 2.0(2)	41 1.9(0.9)	58 3.9(1)	139 4.9(1)	251 4.6(0.5)	476 3.3(0.3)	15/15
f₁₅	511 3.6(6)	9310 4.3(4)	19369 18(20)	20073 18(19)	20769 17(19)	21359 16(18)	14/15
f₁₆	120 2.3(2)	612 5.3(8)	2662 1.7(3)	10449 1.8(2)	11644 1.6(2)	12095 1.6(2)	15/15
f₁₇	5.2 6.8(7)	215 1.0(0.7)	899 1.1(0.7)	3669 0.81(0.7)	6351 1.4(1)	7934 2.0(3)	15/15
f₁₈	103 0.80(0.5)	378 1.4(0.3)	3968 0.25(0.1)	9280 0.43(0.5)	10905 1.4(0.9)	12469 2.0(3)	15/15
f₁₉	1 17(18)	1 3280(5087)	242 542(792)	1.2e5 11(11)	1.2e5 20(21)	1.2e5 21(23)	6/15
f₂₀	16 3.2(3)	851 12(22)	38111 21(24)	54470 15(16)	54861 15(17)	55313 14(16)	14/15
f₂₁	41 12(1)	1157 23(25)	1674 36(57)	1705 35(56)	1729 35(55)	1757 35(54)	14/15
f₂₂	71 46(61)	386 44(57)	938 39(42)	1008 37(39)	1040 36(37)	1068 35(37)	14/15
f₂₃	3.0 2.2(2)	518 97(97)	14249 ∞	31654 ∞	33030 ∞	34256 ∞	15/15
f₂₄	1622 6.2(8)	2.2e5 26(27)	6.4e6 ∞	9.6e6 ∞	1.3e7 ∞	1.3e7 ∞	3/15

Δf	1e+1	1e+0	1e-1	1e-3	1e-5	1e-7	#succ
f₁	43 7.3(2)	43 41(9)	43 61(16)	43 88(23)	43 110(25)	43 128(32)	15/15
f₂	385 26(1)	386 31(2)	387 34(3)	390 38(4)	391 41(6)	393 43(6)	15/15
f₃	5066 1055(1500)	7626 ∞	7635 ∞	7643 ∞	7646 ∞	7651 ∞	15/15
f₄	4722 4193(4430)	7628 ∞	7666 ∞	7700 ∞	7758 ∞	1.4e5 ∞	9/15
f₅	41 8.6(1)	41 10(1)	41 11(2)	41 11(2)	41 11(2)	41 11(2)	15/15
f₆	1296 4.8(0.2)	2343 4.6(0.2)	3413 4.5(0.1)	5220 4.8(0.1)	6728 5.2(0.1)	8409 5.3(0.1)	15/15
f₇	1351 1.9(0.2)	4274 1.5(0.1)	9503 1.0(0.1)	16524 0.89(0.1)	16969 0.91(0.1)	17476 1.0(0.1)	15/15
f₈	2039 7.2(0.6)	3871 7.9(2)	4040 9.1(3)	4219 9.4(4)	4371 10(4)	4484 10(4)	15/15
f₉	1716 8.1(1)	3102 8.9(2)	3277 10(2)	3455 10(2)	3594 11(2)	3727 11(2)	15/15
f₁₀	7413 1.3(0.1)	8661 1.4(0.1)	10735 1.3(0.1)	14920 1.0(0.1)	17073 0.99(0.1)	17476 1.0(0.1)	15/15
f₁₁	1002 4.8(0.3)	2228 3.1(0.2)	6278 1.4(0.1)	9762 1.1(0.2)	12285 1.00(0.2)	14831 0.91(0.2)	15/15
f₁₂	1042 6.6(4)	1938 16(8)	2740 18(18)	4140 35(38)	12407 21(21)	13827 22(21)	15/15
f₁₃	652 7.0(3)	2021 4.6(3)	2751 19(28)	18749 17(21)	24455 40(33)	30201 81(83)	0/15
f₁₄	75 2.6(0.8)	239 7.6(1)	304 14(3)	932 12(1)	1648 10(1)	15661 1.5(0.1)	15/15
f₁₅	30378 44(52)	1.5e5 ∞	3.1e5 ∞	3.2e5 ∞	4.5e5 ∞	4.6e5 ∞	0/15
f₁₆	1384 20(10)	27265 4.2(5)	77015 9.2(7)	1.9e5 108(131)	2.0e5 133(153)	2.2e5 119(128)	6/15
f₁₇	63 2.1(1.0)	1030 3.5(0.4)	4005 2.9(0.2)	30677 0.92(0.0)	56288 1.4(0.7)	80472 12(9)	15/15
f₁₈	621 1.2(0.5)	3972 1.9(0.1)	19561 0.81(0.0)	67569 0.48(0.0)	1.3e5 4.5e5	1.5e5 5.6(6)	15/15
f₁₉	1 108(35)	1 ∞	3.4e5 ∞	6.2e6 ∞	6.7e6 ∞	6.7e6 ∞	0/15
f₂₀	82 5.4(3)	46150 1.1(0.7)	3.1e6 ∞	5.5e6 ∞	5.6e6 ∞	5.6e6 ∞	14/15
f₂₁	561 45(71)	6541 66(98)	14103 37(44)	14643 35(55)	15567 33(51)	17589 30(35)	30/30
f₂₂	467 50(86)	5580 132(251)	23491 226(268)	24948 213(241)	26847 198(206)	1.3e5 39(41)	22/30
f₂₃	3.2 1.5(2)	1614 99427(1e5)	67457 ∞	4.9e5 ∞	8.1e5 ∞	8.4e5 ∞	0/30
f₂₄	1.3e6 ∞	7.5e6 ∞	5.2e7 ∞	5.2e7 ∞	5.2e7 ∞	5.2e7 ∞	3/15

Table 2: Expected running time (ERT in number of function evaluations) divided by the best ERT measured during BBOB-2009 (given in the respective first row) for different Δf values for functions f_1-f_{24} . The median number of conducted function evaluations is additionally given in *italics*, if $\text{ERT}(10^{-7}) = \infty$. #succ is the number of trials that reached the final target $f_{\text{opt}} + 10^{-8}$.

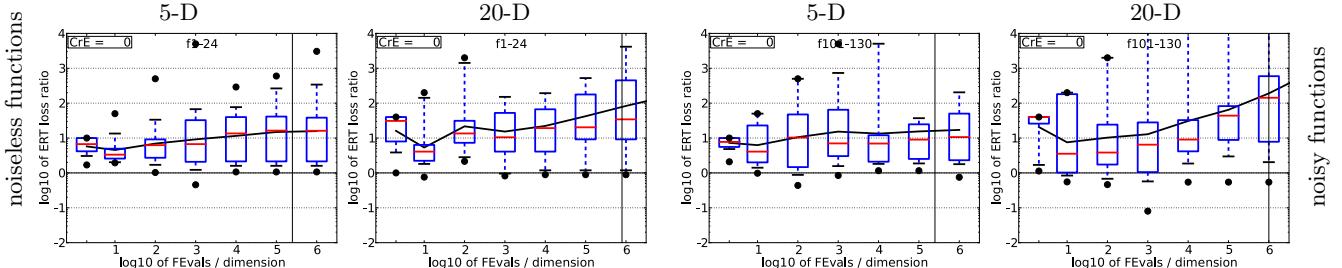


Figure 5: ERT loss ratio vs. a given budget FEvals. The target value f_t used for a given FEvals is the smallest (best) recorded function value such that $\text{ERT}(f_t) \leq \text{FEvals}$ for the presented algorithm. Shown is FEvals divided by the respective best ERT(f_t) from BBOB-2009 for all functions (noiseless f_1-f_{24} , left columns, and noisy $f_{101}-f_{130}$, right columns) in 5-D and 20-D. Line: geometric mean. Box-Whisker error bar: 25-75%-ile with median (box), 10-90%-ile (caps), and minimum and maximum ERT loss ratio (points). The vertical line gives the maximal number of function evaluations in a single trial in this function subset.

- [5] T. Glasmachers, T. Schaul, and J. Schmidhuber. A Natural Evolution Strategy for Multi-Objective Optimization. In *Parallel Problem Solving from Nature (PPSN)*, 2010.
- [6] T. Glasmachers, T. Schaul, Y. Sun, D. Wierstra, and J. Schmidhuber. Exponential Natural Evolution Strategies. In *Genetic and Evolutionary Computation Conference (GECCO)*, Portland, OR, 2010.
- [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] N. Hansen, S. Finck, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking

5-D											20-D										
Δf	1e+1	1e+0	1e-1	1e-3	1e-5	1e-7	#succ	Δf	1e+1	1e+0	1e-1	1e-3	1e-5	1e-7	#succ						
f_{101}	11 2.8(2)	37 2.4(1)	44 5.2(2)	62 7.9(2)	69 10(2)	75 13(3)	15/15 15/15	f_{101}	59 4.8(2)	425 5.1(0.8)	571 5.7(2)	700 6.7(3)	739 7.7(2)	783 8.4(2)	15/15 30/30						
f_{102}	11 3.3(2)	35 2.7(2)	50 4.7(1)	72 7.3(1)	86 9.2(1)	99 10(2)	15/15 15/15	f_{102}	231 1.5(0.4)	399 5.8(2)	579 6.2(2)	921 5.6(2)	1157 4.1(2)	1407 5.5(3)	15/15 30/30						
f_{103}	11 2.7(2)	28 2.9(2)	30 7.2(2)	31 16(3)	35 24(5)	115 11(2)	15/15 15/15	f_{103}	65 5.0(2)	417 5.0(1)	629 4.9(2)	1313 4.1(2)	1893 5.5(3)	2464 6.8(2)	14/15 30/30						
f_{104}	173 1.4(0.4)	773 2.4(0.6)	1287 2.7(6)	1768 2.3(4)	2040 2.1(4)	2284 2.0(3)	15/15 15/15	f_{104}	23690 24(29)	85656 ∞	1.7e5 ∞	1.8e5 ∞	1.9e5 ∞	2.0e5 ∞	15/15 0/30						
f_{105}	167 1.4(0.5)	1436 4.0(4)	5174 14(13)	10388 11(7)	10824 11(7)	11202 10(7)	15/15 15/15	f_{105}	1.9e5 24(23)	6.1e5 ∞	6.3e5 ∞	6.5e5 ∞	6.6e5 ∞	6.7e5 ∞	15/15 0/21						
f_{106}	92 2.7(0.8)	529 3.3(0.9)	1050 3.7(5)	2666 2.2(3)	2887 2.7(3)	3087 15/15	f_{106}	11480 1.2(0.2)	21668 1.4(0.3)	23746 1.5(0.3)	25470 1.7(0.3)	26492 1.8(0.8)	27360 2.0(0.7)	15/15 30/30							
f_{107}	40 3.2(5)	228 7.1(1)	453 4.2(0.5)	940 6.9(11)	1376 7.2(8)	1850 7.0(10)	15/15 15/15	f_{107}	8571 1.6(0.9)	13582 44(41)	16226 786(980)	27357 ∞	52486 ∞	65052 ∞	15/15 0/15						
f_{108}	87 69(120)	5144 12(17)	14469 56(60)	30935 ∞	58628 ∞	80667 ∞	15/15 0/15	f_{108}	58063 ∞	97228 ∞	2.0e5 ∞	4.5e5 ∞	6.3e5 ∞	9.0e5 ∞	15/15 0/15						
f_{109}	11 3.6(2)	57 1.7(0.7)	216 1.3(0.5)	572 1.8(0.5)	873 2.3(0.6)	946 4.1(0.8)	15/15 15/15	f_{109}	333 1.0(0.2)	632 4.7(0.6)	1138 6.9(0.8)	2287 9.4(0.7)	3583 10(0.5)	4952 13(8)	15/15 15/15						
f_{110}	949 0.44(0.2)	33625 1.9(2)	1.2e5 3.6(5)	5.9e5 1.2(1)	6.0e5 1.6(2)	6.1e5 2.3(3)	15/15 8/15	f_{110}	∞ ∞	∞ ∞	∞ ∞	∞ ∞	∞ ∞	∞ ∞	0 0/15						
f_{111}	6856 5.4(8)	6.1e5 1.8(2)	8.8e6 0.96(1)	2.3e7 ∞	3.1e7 ∞	3.1e7 ∞	3/15 0/15	f_{111}	∞ ∞	∞ ∞	∞ ∞	∞ ∞	∞ ∞	∞ ∞	0 0/15						
f_{112}	107 2.3(1)	1684 2.6(5)	3421 7.9(10)	4502 17(18)	5132 28(23)	5596 37(38)	15/15 15/15	f_{112}	25552 0.99(0.2)	64124 59(67)	69621 952(972)	73557 3878(3942)	76137 3747(3913)	78238 3646(3966)	15/15 1/15						
f_{113}	133 1.2(1)	1883 2.4(4)	8081 2.5(4)	24128 1.4(2)	24128 1.4(2)	24402 1.5(2)	15/15 15/15	f_{113}	50123 1.4(1)	3.6e5 179(181)	5.6e5 ∞	5.9e5 ∞	5.9e5 ∞	5.9e5 ∞	15/15 0/15						
f_{114}	767 8.0(13)	14720 9.1(13)	56311 157(169)	83272 ∞	83272 ∞	84949 ∞	15/15 0/15	f_{114}	2.1e5 ∞	1.1e6 ∞	1.4e6 ∞	1.6e6 ∞	1.6e6 ∞	1.6e6 ∞	15/15 0/15						
f_{115}	64 1.3(0.9)	485 0.88(0.4)	1829 1.8(4)	2550 1.8(4)	2550 1.8(4)	2970 1.6(3)	15/15 15/15	f_{115}	2405 1.0(0.1)	30268 0.24(0.0) \downarrow^4	91749 0.14(0.0) \downarrow	1.3e5 0.37(0.3)	1.3e5 0.37(0.3)	1.3e5 0.46(0.5)	15/15 15/15						
f_{116}	5730 1.4(2)	14472 1.2(2)	22311 1.1(2)	26868 1.3(1)	30329 1.6(2)	31661 1.8(2)	15/15 15/15	f_{116}	5.0e5 39(48)	6.9e5 ∞	8.9e5 ∞	1.0e6 ∞	1.1e6 ∞	1.1e6 ∞	15/15 0/15						
f_{117}	26686 8.4(6)	76052 39(49)	1.1e5 ∞	1.4e5 ∞	1.7e5 ∞	1.9e5 ∞	15/15 0/15	f_{117}	1.8e6 ∞	2.5e6 ∞	2.6e6 ∞	2.9e6 ∞	3.2e6 ∞	3.6e6 ∞	15/15 0/30						
f_{118}	429 1.00(0.4)	1217 0.48(0.1) \downarrow^4	1555 0.48(0.1) \downarrow^4	1998 0.69(0.2) \downarrow^2	2430 0.96(0.2)	2913 1.1(0.2)	15/15 15/15	f_{118}	6908 0.94(0.1)	11786 0.77(0.1)	175514 0.80(0.1)	26342 1.1(0.1)	30062 1.1(0.1)	32659 1.7(0.7)	15/15 15/15						
f_{119}	12 3.8(4)	657 2.1(1)	1136 6.3(13)	10372 1.6(2)	35296 2.3(3)	49747 5.9(6)	15/15 15/15	f_{119}	2771 0.53(0.4) \downarrow	29365 10(10)	35930 851(775)	4.1e5 ∞	4.1e5 ∞	4.1e5 ∞	1.9e6 ∞	15/15 0/13					
f_{120}	16 51(18)	2900 13(21)	18698 25(38)	72438 ∞	3.3e5 ∞	5.5e5 ∞	15/15 0/15	f_{120}	36040 6.9(7)	1.8e5 ∞	2.8e5 ∞	1.6e6 ∞	6.7e6 ∞	1.4e7 ∞	13/15 0/15						
f_{121}	8.6 2.6(2)	111 0.85(0.6)	273 0.96(0.5)	1583 1.6(0.8)	3870 1.5(2)	6195 2.3(3)	15/15 15/15	f_{121}	249 0.81(0.3)	769 2.8(0.6)	1426 6.3(0.5)	9304 2.9(0.2)	34434 2.9(0.2)	57404 1.7(0.6)	15/15 15/15						
f_{122}	10 5.0(4)	1727 5.5(11)	9190 8.9(8)	30087 583(727)	53743 ∞	1.1e5 ∞	15/15 0/15	f_{122}	692 0.81(0.9)	52008 ∞	1.4e5 ∞	7.9e5 ∞	2.0e6 ∞	5.8e6 ∞	15/15 0/13						
f_{123}	11 6.6(9)	16066 7.2(7)	81505 ∞	3.4e5 ∞	6.7e5 ∞	2.2e6 ∞	15/15 0/15	f_{123}	1063 11(14)	5.3e5 ∞	1.5e6 ∞	5.3e6 ∞	2.7e7 ∞	1.6e8 ∞	0/15						
f_{124}	10 2.9(3)	202 1.2(0.5)	1040 2.0(0.7)	20478 1.1(1)	45337 36(41)	95200 60(66)	15/15 15/15	f_{124}	192 0.84(0.3)	1959 2.7(0.5)	40840 0.59(0.1)	1.3e5 4.0(3)	3.9e5 ∞	8.0e5 ∞	15/15 0/15						
f_{125}	1 1.3(0.5)	1 23(22)	1 7786(7916)	1 ∞	2.4e5 ∞	2.4e5 ∞	15/15 0/15	f_{125}	1 1.7(1)	1 470(238)	1 ∞	2.5e7 ∞	8.0e7 ∞	8.1e7 ∞	4/15 0/15						
f_{126}	1 1.1(0.5)	1 44(88)	1 41598(46392)	1 ∞	3.4e5 ∞	3.9e5 ∞	15/15 0/15	f_{126}	1 1.4(1)	1 26641(51744)	1 ∞	4.4e6 ∞	7.3e6 ∞	7.4e6 ∞	0 0/15						
f_{127}	1 1.1(0.5)	1 20(16)	1 3858(5345)	1 ∞	3.4e5 ∞	3.9e5 ∞	15/15 0/15	f_{127}	1 1.5(0.5)	1 131(87)	1 ∞	4.4e6 ∞	7.3e6 ∞	7.4e6 ∞	15/15 0/15						
f_{128}	111 22(67)	4248 6.2(9)	7808 5.7(7)	12447 3.6(4)	17217 2.6(3)	21162 2.1(2)	15/15 15/15	f_{128}	1.4e5 19(23)	1.3e7 0.40(0.4)	1.7e7 0.78(0.7)	1.7e7 1.0(1)	1.7e7 1.0(1)	1.7e7 1.3(1)	9/15 6/15						
f_{129}	64 17(16)	10710 11(11)	59443 8.6(11)	2.8e5 14(16)	5.1e5 ∞	5.8e5 ∞	15/15 0/15	f_{129}	7.8e6 ∞	4.1e7 ∞	4.2e7 ∞	4.2e7 ∞	4.2e7 ∞	5/15 0/15							
f_{130}	55 17(0.9)	812 67(79)	3034 28(48)	32823 2.6(4)	33889 2.5(4)	34528 15/15	f_{130}	4904 15(23)	93149 5.3(4)	2.5e5 5.4(9)	2.5e5 5.4(9)	2.6e5 5.4(9)	2.6e5 5.4(9)	7/15 15/15							

Table 3: ERT ratios, as in table 2, for functions f_{101} – f_{130} .

- 2009: Noisy functions definitions. Technical Report RR-6869, INRIA, 2009. Updated February 2010.
- [10] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *IEEE Transactions on Evolutionary Computation*, 9:159–195, 2001.
- [11] T. Schaul. *Studies in Continuous Black-box Optimization*. Ph.D. thesis, Technische Universität München, 2011.
- [12] T. Schaul. Natural Evolution Strategies Converge on Sphere Functions. In *Genetic and Evolutionary Computation Conference (GECCO)*, Philadelphia, PA, 2012.
- [13] T. Schaul, J. Bayer, D. Wierstra, Y. Sun, M. Felder, F. Sehnke, T. Rückstieß, and J. Schmidhuber. PyBrain. *Journal of Machine Learning Research*, 11:743–746, 2010.
- [14] T. Schaul and J. Schmidhuber. Metalearning. *Scholarpedia*, 5(6):4650, 2010.
- [15] Y. Sun, D. Wierstra, T. Schaul, and J. Schmidhuber. Stochastic search using the natural gradient. In

International Conference on Machine Learning (ICML), 2009.

- [16] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, and J. Schmidhuber. Natural Evolution Strategies. Technical report, 2011.
- [17] D. Wierstra, T. Schaul, J. Peters, and J. Schmidhuber. Natural Evolution Strategies. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC)*, Hong Kong, China, 2008.