



HAL
open science

Maximum Likelihood-based Online Adaptation of Hyper-parameters in CMA-ES

Ilya Loshchilov, Marc Schoenauer, Michèle Sebag, Nikolaus Hansen

► **To cite this version:**

Ilya Loshchilov, Marc Schoenauer, Michèle Sebag, Nikolaus Hansen. Maximum Likelihood-based Online Adaptation of Hyper-parameters in CMA-ES. 13th International Conference on Parallel Problem Solving from Nature (PPSN 2014), Sep 2014, Ljubljana, Slovenia. hal-01003504v1

HAL Id: hal-01003504

<https://inria.hal.science/hal-01003504v1>

Submitted on 10 Jun 2014 (v1), last revised 11 Jun 2014 (v2)

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.

Maximum Likelihood-based Online Adaptation of Hyper-parameters in CMA-ES

Ilya Loshchilov¹, Marc Schoenauer^{2,3}, Michèle Sebag^{3,2}, and
Nikolaus Hansen^{2,3}

¹ Laboratory of Intelligent Systems,
École Polytechnique Fédérale de Lausanne, Switzerland, Ilya.Loshchilov@epfl.ch

² TAO Project-team, INRIA Saclay - Île-de-France

³ Laboratoire de Recherche en Informatique (UMR CNRS 8623)
Université Paris-Sud, 91128 Orsay Cedex, France
FirstName.LastName@inria.fr

Abstract. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is widely accepted as a robust derivative-free continuous optimization algorithm for non-linear and non-convex optimization problems. CMA-ES is well known to be almost parameterless, meaning that only one hyper-parameter, the population size, is proposed to be tuned by the user. In this paper, we propose a principled approach called self-CMA-ES to achieve the online adaptation of CMA-ES hyper-parameters in order to improve its overall performance. Experimental results show that for larger-than-default population size, the default settings of hyper-parameters of CMA-ES are far from being optimal, and that self-CMA-ES allows for dynamically approaching optimal settings.

1 Introduction

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES [5]) is a continuous optimizer which only exploits the ranking of estimated candidate solutions to approach the optimum of an objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. CMA-ES is also invariant w.r.t. affine transformations of the decision space, explaining the known robustness of the algorithm. An important practical advantage of CMA-ES is that all hyper-parameters thereof are defined by default with respect to the problem dimension n . Practically, only the population size λ is suggested to be tuned by the user, e.g. when a parallelization of the algorithm is considered or the problem at hand is known to be multi-modal and/or noisy [1, 8]. Other hyper-parameters have been provided robust default settings (depending on n and λ), in the sense that their offline tuning allegedly hardly improves the CMA-ES performance for unimodal functions. In the meanwhile, for multi-modal functions it is suggested that the overall performance can be significantly improved by offline tuning of λ and multiple stopping criteria [16, 11]. Additionally, it is shown that CMA-ES can be improved by a factor up to 5-10 by the use of surrogate models on unimodal ill-conditioned functions [14]. This suggests that

the CMA-ES performance can be improved by better exploiting the information in the evaluated samples $(\mathbf{x}, f(\mathbf{x}))$.

This paper focuses on the automatic online adjustment of the CMA-ES hyper-parameters. The proposed approach, called self-CMA-ES, relies on a second CMA-ES instance operating on the hyper-parameter space of the first CMA-ES, and aimed at increasing the likelihood of generating the most successful samples \mathbf{x} in the current generation. The paper is organized as follows. Section 2 describes the original $(\mu/\mu_w, \lambda)$ -CMA-ES. self-CMA-ES is described in Section 3 and its experimental validation is discussed in Section 4 comparatively to related work. Section 5 concludes the paper.

2 Covariance Matrix Adaptation Evolution Strategy

The Covariance Matrix Adaptation Evolution Strategy [6, 7, 5] is acknowledgedly the most popular and the most efficient Evolution Strategy algorithm.

The original $(\mu/\mu_w, \lambda)$ -CMA-ES (Algorithm 1) proceeds as follows. At the t -th iteration, a Gaussian distribution $\mathcal{N}(\mathbf{m}^t, \sigma^{t2} \mathbf{C}^t)$ is used to generate λ candidate solution $\mathbf{x}_k \in \mathbb{R}^n$, for $k = 1 \dots \lambda$ (line 5):

$$\mathbf{x}_k^t = \mathcal{N}(\mathbf{m}^t, \sigma^{t2} \mathbf{C}^t) = \mathbf{m}^t + \sigma^t \mathcal{N}(\mathbf{0}, \mathbf{C}^t), \quad (1)$$

where the mean $\mathbf{m}^t \in \mathbb{R}^n$ of the distribution can be interpreted as the current estimate of the optimum of function f , $\mathbf{C}^t \in \mathbb{R}^{n \times n}$ is a (positive definite) covariance matrix and σ^t is a mutation step-size. These λ solutions are evaluated according to f (line 6). The new mean \mathbf{m}^{t+1} of the distribution is computed as a *weighted sum* of the best μ individuals out of the λ ones (line 7). Weights $w_1 \dots w_\mu$ are used to control the impact of selected individuals, with usually higher weights for top ranked individuals (line 1).

The adaptation of the step-size σ^t , inherited from the Cumulative Step-Size Adaptation Evolution Strategy (CSA-ES [6]), is controlled by the evolution path \mathbf{p}_σ^{t+1} . Successful mutation steps $\frac{\mathbf{m}^{t+1} - \mathbf{m}^t}{\sigma^t}$ (line 8) are tracked in the sampling space, i.e., in the isotropic coordinate system defined by the eigenvectors of the covariance matrix \mathbf{C}^t . To update the evolution path \mathbf{p}_σ^{t+1} , i) a decay/relaxation factor c_σ is used to decrease the importance of previous steps; ii) the step-size is increased if the length of the evolution path \mathbf{p}_σ^{t+1} is longer than the expected length of the evolution path under random selection $\mathbb{E} \|\mathcal{N}(\mathbf{0}, \mathbf{I})\|$; iii) otherwise it is decreased (line 13). Expectation of $\|\mathcal{N}(\mathbf{0}, \mathbf{I})\|$ is approximated by $\sqrt{n}(1 - \frac{1}{4n} + \frac{1}{21n^2})$. A damping parameter d_σ controls the change of the step-size.

The covariance matrix update consists of two parts (line 12): a *rank-one update* [7] and a *rank- μ update* [5]. The rank-one update computes the evolution path \mathbf{p}_c^{t+1} of successful moves of the mean $\frac{\mathbf{m}^{t+1} - \mathbf{m}^t}{\sigma^t}$ of the mutation distribution in the given coordinate system (line 10), along the same lines as the evolution path \mathbf{p}_σ^{t+1} of the step-size. To stall the update of \mathbf{p}_c^{t+1} when σ increases rapidly, a h_σ trigger is used (line 9).

Algorithm 1 The $(\mu/\mu_w, \lambda)$ -CMA-ES

- 1: **given** $n \in \mathbb{N}_+$, $\lambda = 4 + \lfloor 3 \ln n \rfloor$, $\mu = \lfloor \lambda/2 \rfloor$, $w_i = \frac{\ln(\mu + \frac{1}{2}) - \ln i}{\sum_{j=1}^{\mu} (\ln(\mu + \frac{1}{2}) - \ln j)}$ for $i = 1 \dots \mu$,
 $\mu_w = \frac{1}{\sum_{i=1}^{\mu} w_i^2}$, $c_\sigma = \frac{\mu_w + 2}{n + \mu_w + 3}$, $d_\sigma = 1 + c_\sigma + 2 \max(0, \sqrt{\frac{\mu_w - 1}{n + 1}} - 1)$, $c_c = \frac{4}{n + 4}$,
 $c_1 = \frac{2}{(n + 1.3)^2 + \mu_w}$, $c_\mu = \frac{2(\mu_w - 2 + 1/\mu_w)}{(n + 2)^2 + \mu_w}$
 - 2: **initialize** $\mathbf{m}^{t=0} \in \mathbb{R}^n$, $\sigma^{t=0} > 0$, $\mathbf{p}_\sigma^{t=0} = \mathbf{0}$, $\mathbf{p}_c^{t=0} = \mathbf{0}$, $\mathbf{C}^{t=0} = \mathbf{I}$, $t \leftarrow 0$
 - 3: **repeat**
 - 4: **for** $k = 1, \dots, \lambda$ **do**
 - 5: $\mathbf{x}_k = \mathbf{m}^t + \sigma^t \mathcal{N}(\mathbf{0}, \mathbf{C}^t)$
 - 6: $\mathbf{f}_k = f(\mathbf{x}_k)$
 - 7: $\mathbf{m}^{t+1} = \sum_{i=1}^{\mu} w_i \mathbf{x}_{i:\lambda}$ // the symbol $i : \lambda$ denotes i -th best individual on f
 - 8: $\mathbf{p}_\sigma^{t+1} = (1 - c_\sigma) \mathbf{p}_\sigma^t + \sqrt{c_\sigma(2 - c_\sigma)} \sqrt{\mu_w} \mathbf{C}^t - \frac{1}{2} \frac{\mathbf{m}^{t+1} - \mathbf{m}^t}{\sigma^t}$
 - 9: $h_\sigma = \mathbb{1}_{\|\mathbf{p}_\sigma^{t+1}\| < \sqrt{1 - (1 - c_\sigma)^{2(t+1)}} (1.4 + \frac{2}{n+1}) \mathbb{E}\|\mathcal{N}(\mathbf{0}, \mathbf{I})\|}$
 - 10: $\mathbf{p}_c^{t+1} = (1 - c_c) \mathbf{p}_c^t + h_\sigma \sqrt{c_c(2 - c_c)} \sqrt{\mu_w} \frac{\mathbf{m}^{t+1} - \mathbf{m}^t}{\sigma^t}$
 - 11: $\mathbf{C}_\mu = \sum_{i=1}^{\mu} w_i \frac{\mathbf{x}_{i:\lambda} - \mathbf{m}^t}{\sigma^t} \times \frac{(\mathbf{x}_{i:\lambda} - \mathbf{m}^t)^T}{\sigma^t}$
 - 12: $\mathbf{C}^{t+1} = (1 - c_1 - c_\mu) \mathbf{C}^t + c_1 \underbrace{\mathbf{p}_c^{t+1} \mathbf{p}_c^{t+1T}}_{\text{rank-one update}} + c_\mu \underbrace{\mathbf{C}_\mu}_{\text{rank-}\mu \text{ update}}$
 - 13: $\sigma^{t+1} = \sigma^t \exp(\frac{c_\sigma}{d_\sigma} (\frac{\|\mathbf{p}_\sigma^{t+1}\|}{\mathbb{E}\|\mathcal{N}(\mathbf{0}, \mathbf{I})\|} - 1))$
 - 14: $t = t + 1$
 - 15: **until** *stopping criterion is met*
-

The rank- μ update computes a covariance matrix \mathbf{C}_μ as a weighted sum of the covariances of successful steps of the best μ individuals (line 11). Covariance matrix \mathbf{C} itself is replaced by a weighted sum of the rank-one (weight c_1 [7]) and rank- μ (weight c_μ [5]) updates, with c_1 and c_μ positive and $c_1 + c_\mu \leq 1$.

While the optimal parameterization of CMA-ES remains an open problem, the default parameterization is found quite robust on noiseless unimodal functions [5], which explains the popularity of CMA-ES.

3 The self-CMA-ES

The proposed self-CMA-ES approach is based on the intuition that the optimal hyper-parameters of CMA-ES at time t should favor the generation of the best individuals at time t , under the (strong) assumption that an optimal parameterization and performance of CMA-ES in each time t will lead to the overall optimal performance.

Formally, this intuition leads to the following procedure. Let θ_f^t denote the hyper-parameter vector used for the optimization of objective f at time t (CMA-ES stores its state variables and internal parameters of iteration t in θ^t and the ‘.’-notation is used to access them). At time $t + 1$, the best individuals generated according to θ_f^t are known to be the top-ranked individuals $\mathbf{x}_{1:\lambda}^t \dots \mathbf{x}_{\mu:\lambda}^t$, where $\mathbf{x}_{i:\lambda}^t$ stands for the i -th best individual w.r.t. f . Hyper-parameter vector θ_f^t would

thus have been all the better, if it had maximized the probability of generating these top individuals.

Along this line, the optimization of θ_f^t is conducted using a second CMA-ES algorithm, referred to as *auxiliary* CMA-ES as opposed to the one concerned with the optimization of f , referred to as *primary* CMA-ES. The objective of the auxiliary CMA-ES is specified as follows:

Given: hyper-parameter vector θ_f^i and points $(\mathbf{x}_{1:\lambda}^i, f(\mathbf{x}_{1:\lambda}^i))$ evaluated by primary CMA-ES at steps $i = 1, \dots, t$ (noted as $\theta_f^{i+1}.f(\mathbf{x}_{1:\lambda})$ in Algorithm 2),

Find: $\theta_f^{t,*}$ such that i) backtracking the primary CMA-ES to its state at time $t-1$; ii) replacing θ_f^t by $\theta_f^{t,*}$, would maximize the likelihood of $\mathbf{x}_{i:\lambda}^t$ for $i = 1 \dots \mu$.

The auxiliary CMA-ES might thus tackle the maximization of $g_t(\theta)$ computed as the weighed log-likelihood of the top-ranked μ_{sel} individuals at time t :

$$g_t(\theta) = \sum_{i=1}^{\mu_{sel}} w_{sel,i} \log (P(\mathbf{x}_{i:\lambda}^t | \theta_f^t = \theta)), \quad (2)$$

where $w_{sel,i} \geq 0$, $i = 1 \dots \mu_{sel}$, $\sum_{i=1}^{\mu_{sel}} w_{sel,i} = 1$, and by construction

$$P(\mathbf{x}_i | \mathbf{m}^t, \mathbf{C}^t) = \frac{1}{\sqrt{(2\pi)^n |\mathbf{C}^t|}} \exp(-0.5(\mathbf{m}^t - \mathbf{x}_i^t) \mathbf{C}^{t-1} (\mathbf{m}^t - \mathbf{x}_i^t)), \quad (3)$$

where \mathbf{C}^t is the covariance matrix multiplied by σ^{t^2} and $|\mathbf{C}^t|$ is its determinant.

While the objective function for the auxiliary CMA-ES defined by Eq. (2) is mathematically sound, it yields a difficult optimization problem; firstly the probabilities are scale-sensitive; secondly and overall, in a worst case scenario, a single good but unlikely solution may lead the optimization of θ_f^t astray.

Therefore, another optimization objective $h_t(\theta)$ is defined for the auxiliary CMA-ES, where $h_t(\theta)$ measures the agreement on $\mathbf{x}_{i:\lambda}^t$ for $i = 1 \dots \mu$ between i) the order defined from f ; ii) the order defined from their likelihood conditioned by $\theta_f^t = \theta$ (Algorithm 3). Procedure *ReproduceGenerationCMA* in Algorithm 3 updates the strategy parameters described from line 7 to line 14 in Algorithm 1 using already evaluated solutions stored in $\theta_f^t.\mathbf{x}_{i:\lambda}$. Line 4 computes the Mahalanobis distance, division by step-size is not needed since only ranking will be considered in line 5 (decreasing order of Mahalanobis distances corresponds to increasing order of log-likelihoods). Line 6 computes a weighted sum of ranks of likelihoods of best individuals.

Finally, the overall scheme of self-CMA-ES (Algorithm 2) involves two interdependent CMA-ES optimization algorithms, where the primary CMA-ES is concerned with optimizing objective f , and the auxiliary CMA-ES is concerned with optimizing objective h_t , that is, optimizing the hyper-parameters of the primary CMA-ES⁴. Note that self-CMA-ES is not *per se* a “more parameterless“

⁴ This scheme is actually inspired from the one proposed for surrogate-assisted optimization [13], where the auxiliary CMA-ES was in charge of optimizing the surrogate learning hyper-parameters.

Algorithm 2 The self-CMA-ES

```
1:  $t \leftarrow 1$ 
2:  $\theta_f^t \leftarrow \text{InitializationCMA}()$  { primary CMA-ES aimed at optimizing  $f$  }
3:  $\theta_h^t \leftarrow \text{InitializationCMA}()$  { auxiliary CMA-ES aimed at optimizing  $h_t$  }
4: fill  $\theta_f^t$  with corresponding parameters stored in mean of distribution  $\theta_h^t.m$ 
5:  $\theta_f^{t+1} \leftarrow \text{GenerationCMA}(f, \theta_f^t)$ 
6:  $t \leftarrow t + 1$ 
7: repeat
8:    $\theta_f^{t+1} \leftarrow \text{GenerationCMA}(f, \theta_f^t)$ 
9:    $\theta_h^{t+1} \leftarrow \text{GenerationCMA}(h_t, \theta_h^t)$ 
10:  fill  $\theta_f^{t+1}$  with corresponding parameters stored in mean of distribution  $\theta_h^{t+1}.m$ 
11:   $t \leftarrow t + 1$ 
12: until stopping criterion is met
```

algorithm than CMA-ES, in the sense that the user is still invited to modify the population size λ . The main purpose of self-CMA-ES is to achieve the online adaptation of the other CMA-ES hyper-parameters.

Specifically, while the primary CMA-ES optimizes $f(\mathbf{x})$ (line 8), the auxiliary CMA-ES maximizes $h_t(\theta)$ (line 9) by sampling and evaluating λ_h variants of θ_f^t . The updated mean of the auxiliary CMA-ES in the hyper-parameter space is used as a local estimate of the optimal hyper-parameter vector for the primary CMA-ES. Note that the auxiliary CMA-ES achieves a single iteration in the hyper-parameter space of the primary CMA-ES, with two motivations: limiting the computational cost of self-CMA-ES (which scales as λ_h times the time complexity of the CMA-ES), and preventing θ_f^t from overfitting the current sample $\mathbf{x}_{i:\lambda}^t, i = 1 \dots \mu$.

Algorithm 3 Objective function $h_t(\theta)$

```
1: Input:  $\theta, \theta_f^{t+1}, \theta_f^{t-1}, \theta_f^t, \mu, w_{sel,i}$  for  $i = 1, \dots, \mu$ 
2:  $\theta_f^{t-1} \leftarrow \theta$ 
3:  $\theta_f^t \leftarrow \text{ReproduceGenerationCMA}(f, \theta_f^{t-1})$  using already evaluated  $\theta_f^t.\mathbf{x}_{i:\lambda}$ 
4:  $d_i \leftarrow \left\| \theta_f^t \cdot \sqrt{C^{-1}} \cdot (\theta_f^{t+1}.\mathbf{x}_i^t - \theta_f^t.m) \right\|$ ; for  $i = 1, \dots, \theta_f^{t+1}.\lambda$ 
5:  $p_i \leftarrow \text{rank of } d_i, i = 1 \dots \lambda$  sorted in decreasing order
6:  $h(\theta) \leftarrow \sum_{i=1}^{\mu} w_{sel,i} p_{i:\lambda}$  {  $i : \lambda$  denotes the rank of  $\theta_f^{t+1}.\mathbf{x}_i$  }
7: Output:  $h(\theta)$ 
```

4 Experimental Validation

The experimental validation of self-CMA-ES investigates the performance of the algorithm comparatively to CMA-ES on the BBOB noiseless problems [4]. Both algorithms are launched in IPOP scenario of restarts when the CMA-ES

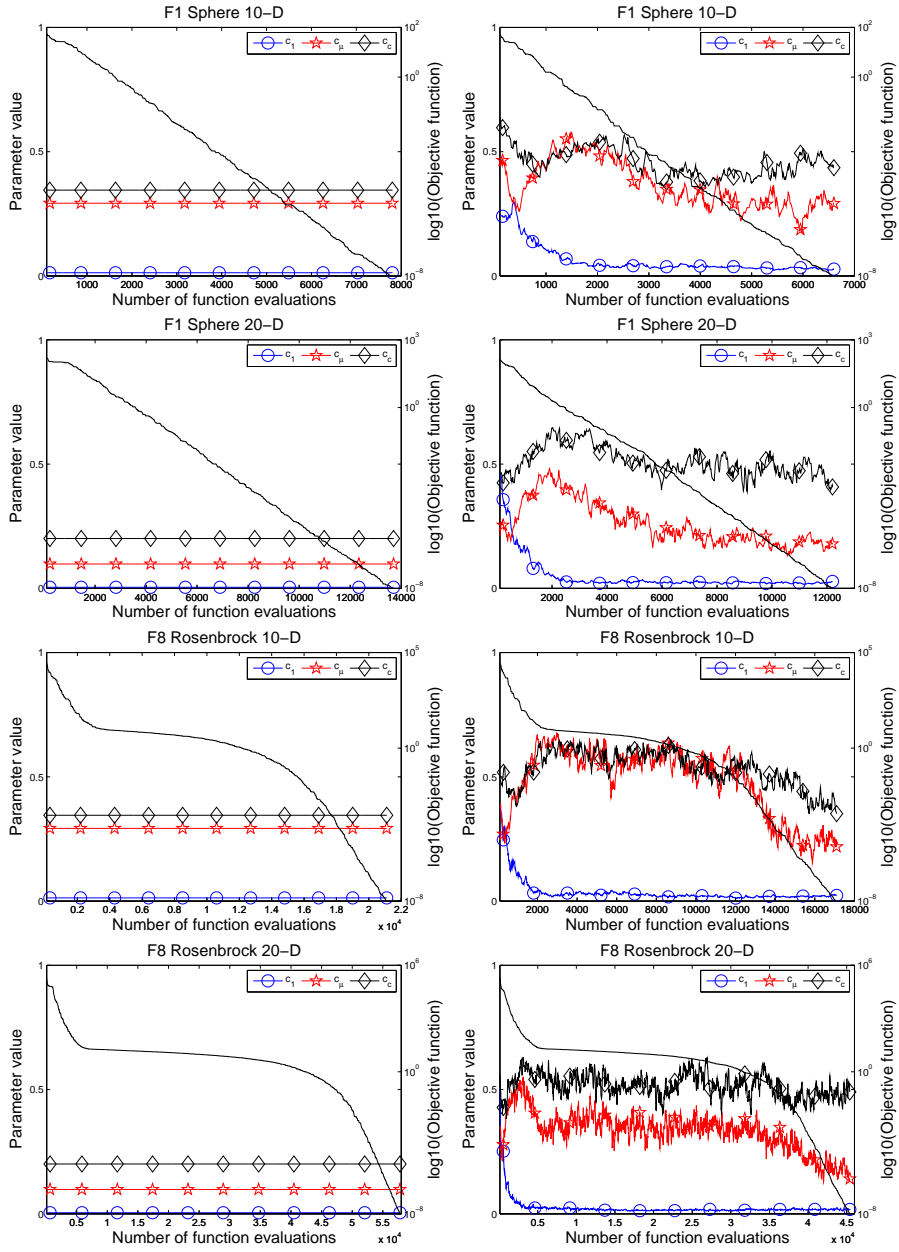


Fig. 1. Evolution of c_1 , c_μ , c_c (left y-axis) and \log_{10} (objective function) (right y-axis) of CMA-ES (left column) and self-CMA-ES (right column) on 10- and 20-dimensional Sphere and Rosenbrock functions from [4]. The medians of 15 runs are shown.

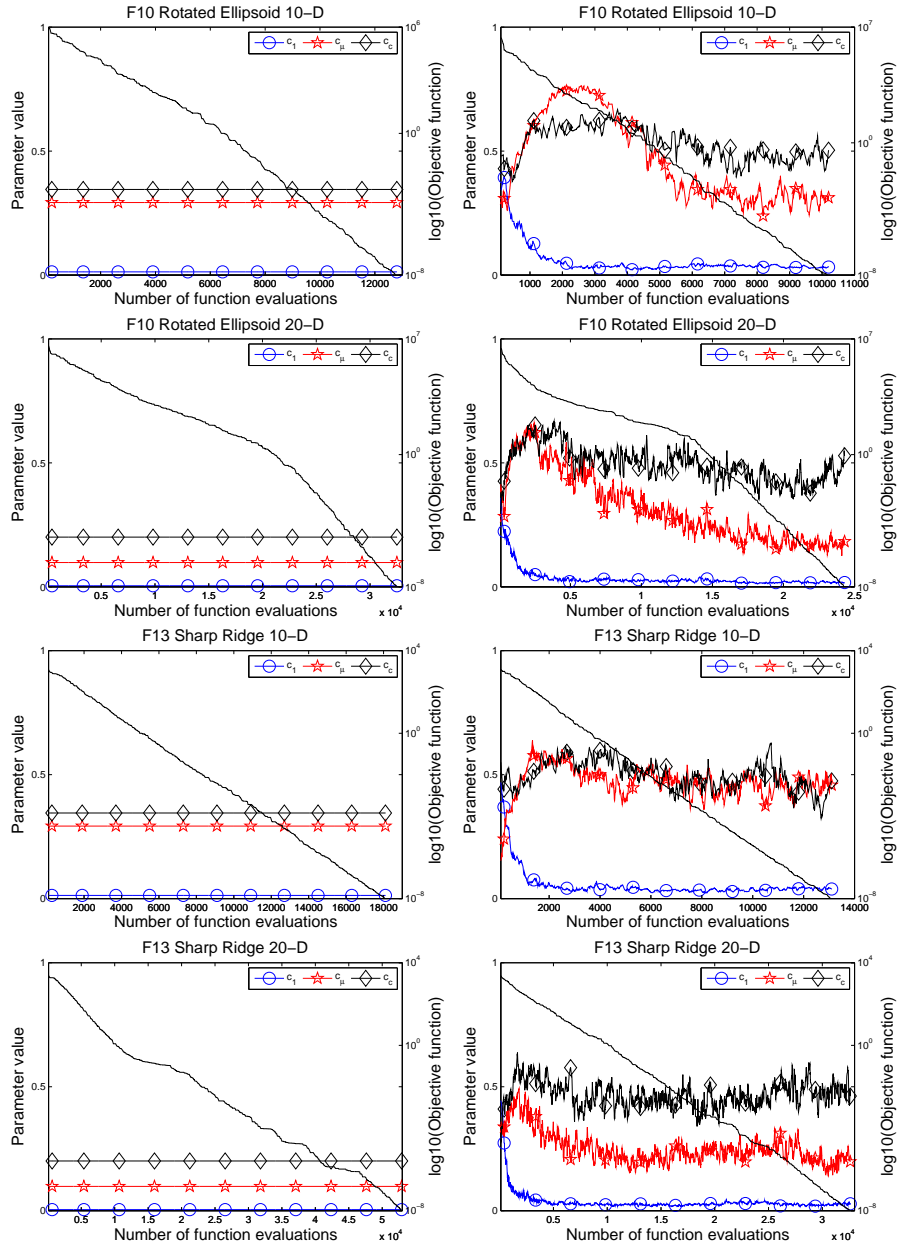


Fig. 2. Evolution of c_1 , c_μ , c_c (left y-axis) and $\log_{10}(\text{objective function})$ (right y-axis) of CMA-ES (left column) and self-CMA-ES (right column) on 10- and 20-dimensional Rotated Ellipsoid and Sharp Ridge functions [4]. The medians of 15 runs are shown.

is restarted with doubled population size once stopping criteria [3] are met⁵. The population size λ is chosen to be 100 for both CMA-ES and self-CMA-ES. We choose this value (about 10 times larger than the default one, see the default parameters of CMA-ES in Algorithm 1) to investigate how sub-optimal the other CMA-ES hyper-parameters, derived from λ , are in such a case, and whether self-CMA-ES can recover from this sub-optimality.

The auxiliary CMA-ES is concerned with optimizing hyper-parameters c_1 , c_μ and c_c (Algorithm 1), responsible for the adaptation of the covariance matrix of the primary CMA-ES. These parameters range in $[0, .9]$ subject to $0 \leq c_1 + c_\mu \leq 0.9$; the constraint is meant to enforce a feasible \mathbf{C} update for the primary CMA-ES (the decay factor of \mathbf{C} should be in $[0, 1]$). Infeasible hyper-parameter solutions get a very large penalty, multiplied by the sum of distances of infeasible hyper-parameters to the range of feasibility.

We set $w_{sel,i} = 1/\mu$ for $i = 1, \dots, \mu$ and $\mu = \lfloor \lambda/2 \rfloor$ to 50. The internal computational complexity of self-CMA-ES thus is $\lambda_h = 20$ times larger than the one of CMA-ES without lazy update (being reminded that the internal time complexity is usually negligible compared to the cost per objective function evaluation).

4.1 Results

Figures 1 and 2 display the comparative performances of CMA-ES (left) and self-CMA-ES (right) on 10 and 20-dimensional Sphere, Rosenbrock, Rotated Ellipsoid and Sharp ridge functions from the noiseless BBOB testbed [4] (medians out of 15 runs). Each plot shows the value of the hyper-parameters (left y-axis) together with the objective function (in logarithmic scale, right y-axis). Hyper-parameters c_1 , c_μ and c_c are constant and set to their default values for CMA-ES while they are adapted along evolution for self-CMA-ES.

In self-CMA-ES, the hyper-parameters are uniformly initialized in $[0, 0.9]$ (therefore the medians are close to 0.45) and they gradually converge to values which are estimated to provide the best update of the covariance matrix w.r.t. the ability to generate the current best individuals. It is seen that these values are problem and dimension-dependent. The values of c_1 are always much smaller than c_μ but are comparable to the default c_1 . The values of c_μ and c_c and c_1 are almost always larger than the default ones; this is not a surprise for c_1 and c_μ , as their original default values are chosen in a rather conservative way to prevent degeneration of the covariance matrix.

Several interesting observations can be made about the dynamics of the parameter values. The value of c_μ is high most of the times on the Rosenbrock functions, but it decreases toward values similar to those of the Sphere functions, when close to the optimum. This effect is observed on most problems; indeed, on most problem fast adaptation of the covariance matrix will improve

⁵ For the sake of reproducibility, the source code is available at <https://sites.google.com/site/selfcmapps/>

the performance in the beginning, while the distribution shape should remain stable when the covariance matrix is learned close to the optimum.

The overall performance of self-CMA-ES on the considered problems is comparable to that of CMA-ES, with a speed-up of a factor up to 1.5 on Sharp Ridge functions. The main result is the ability of self-CMA-ES to achieve the online adaptation of the hyper-parameters depending on the problem at hand, side-stepping the use of long calibrated default settings⁶.

4.2 Discussion

self-CMA-ES offers a proof of concept for the online adaptation of three CMA-ES hyper-parameters in terms of feasibility and usefulness. Previous studies on parameter settings for CMA-ES mostly considered offline tuning (see, e.g., [16, 11]) and theoretical analysis dated back to the first papers on Evolution Strategies. The main limitation of these studies is that the suggested hyper-parameter values are usually specific to the (class of) analyzed problems. Furthermore, the suggested values are fixed, assuming that optimal parameter values remain constant along evolution. However, when optimizing a function whose landscape gradually changes when approaching the optimum, one may expect optimal hyper-parameter values to reflect this change as well.

Studies on the online adaptation of hyper-parameters (apart from σ , \mathbf{m} and \mathcal{C}) usually consider population size in noisy [2], multi-modal [1, 12] or expensive [9] optimization. A more closely related approach was proposed in [15] where the learning rate for step-size adaptation is adapted in a stochastic way similarly to Rprop-updates [10].

5 Conclusion and Perspectives

This paper proposes a principled approach for the self-adaptation of CMA-ES hyper-parameters, tackled as an auxiliary optimization problem: maximizing the likelihood of generating the best sampled solutions. The experimental validation of self-CMA-ES shows that the learning rates involved in the covariance matrix adaptation can be efficiently adapted on-line, with comparable or better results than CMA-ES. It is worth emphasizing that matching the performance of CMA-ES, the default setting of which represent a historical consensus between theoretical analysis and offline tuning, is nothing easy.

The main novelty of the paper is to offer an intrinsic assessment of the algorithm internal state, based on retrospective reasoning (given the best current solutions, how could the generation of these solutions have been made easier) and on one assumption (the optimal hyper-parameter values at time t are "sufficiently good" at time $t + 1$). Further work will investigate how this intrinsic assessment can support the self-adaptation of other continuous and discrete hyper-parameters used to deal with noisy, multi-modal and constrained optimization problems.

⁶ $c_c = \frac{4}{n+4}$, $c_1 = \frac{2}{(n+1.3)^2 + \mu_w}$, $c_\mu = \frac{2(\mu_w - 2 + 1/\mu_w)}{(n+2)^2 + \mu_w}$.

Acknowledgments We acknowledge anonymous reviewers for their constructive comments. This work was supported by the grant ANR-2010-COSI-002 (SIMINOLE) of the French National Research Agency.

References

1. A. Auger and N. Hansen. A Restart CMA Evolution Strategy With Increasing Population Size. In *IEEE Congress on Evolutionary Computation*, pages 1769–1776. IEEE Press, 2005.
2. H.-G. Beyer and M. Hellwig. Controlling population size and mutation strength by meta-es under fitness noise. In *Proceedings of the Twelfth Workshop on Foundations of Genetic Algorithms XII, FOGA XII '13*, pages 11–24. ACM, 2013.
3. N. Hansen. Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed. In *GECCO Companion*, pages 2389–2396, 2009.
4. N. Hansen, A. Auger, S. Finck, and R. Ros. Real-Parameter Black-Box Optimization Benchmarking 2010: Experimental Setup. Technical report, INRIA, 2010.
5. N. Hansen, S. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 2003.
6. N. Hansen and A. Ostermeier. Adapting Arbitrary Normal Mutation Distributions in Evolution Strategies: The Covariance Matrix Adaptation. In *International Conference on Evolutionary Computation*, pages 312–317, 1996.
7. N. Hansen and A. Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2):159–195, June 2001.
8. N. Hansen and R. Ros. Benchmarking a weighted negative covariance matrix update on the BBOB-2010 noisy testbed. In *GECCO '10: Proceedings of the 12th annual conference comp on Genetic and evolutionary computation*, pages 1681–1688, New York, NY, USA, 2010. ACM.
9. F. Hoffmann and S. Holemann. Controlled Model Assisted Evolution Strategy with Adaptive Preselection. In *International Symposium on Evolving Fuzzy Systems*, pages 182–187. IEEE, 2006.
10. C. Igel and M. Hüsken. Empirical evaluation of the improved rprop learning algorithms. *Neurocomputing*, 50:105–123, 2003.
11. T. Liao and T. Stützle. Benchmark results for a simple hybrid algorithm on the CEC 2013 benchmark set for real-parameter optimization. In *IEEE Congress on Evolutionary Computation (CEC)*, pages 1938–1944. IEEE press, 2013.
12. I. Loshchilov, M. Schoenauer, and M. Sebag. Alternative Restart Strategies for CMA-ES. In V. C. et al., editor, *Parallel Problem Solving from Nature (PPSN XII)*, LNCS, pages 296–305. Springer, September 2012.
13. I. Loshchilov, M. Schoenauer, and M. Sebag. Self-Adaptive Surrogate-Assisted Covariance Matrix Adaptation Evolution Strategy. In *Genetic and Evolutionary Computation Conference (GECCO)*, pages 321–328. ACM Press, July 2012.
14. I. Loshchilov, M. Schoenauer, and M. Sebag. Intensive Surrogate Model Exploitation in Self-adaptive Surrogate-assisted CMA-ES (saACM-ES). In *Genetic and evolutionary computation conference*, pages 439–446. ACM, 2013.
15. T. Schaul. Comparing natural evolution strategies to bipop-cma-es on noiseless and noisy black-box optimization testbeds. In *Genetic and evolutionary computation conference companion*, pages 237–244. ACM, 2012.
16. S. Smit and A. Eiben. Beating the ‘world champion’ Evolutionary Algorithm via REVAC Tuning. In *IEEE Congress on Evolutionary Computation*, pages 1–8, 2010.