



CMA-ES with Two-Point Step-Size Adaptation

Nikolaus Hansen

► To cite this version:

Nikolaus Hansen. CMA-ES with Two-Point Step-Size Adaptation. [Research Report] RR-6527, 2008. inria-00276854v4

HAL Id: inria-00276854

<https://inria.hal.science/inria-00276854v4>

Submitted on 13 May 2008 (v4), last revised 18 May 2008 (v5)

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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

CMA-ES with Two-Point Step-Size Adaptation

Nikolaus Hansen

N° 6527

May 2008

Thème COG



*rapport
de recherche*

CMA-ES with Two-Point Step-Size Adaptation

Nikolaus Hansen*

Thème COG — Systèmes cognitifs
Équipes-Projets Adaptive Combinatorial Search et TAO

Rapport de recherche n° 6527 — May 2008 — 9 pages

Abstract: We combine a refined version of *two-point step-size adaptation* with the *covariance matrix adaptation evolution strategy* (CMA-ES). Additionally, we suggest polished formulae for the learning rate of the covariance matrix and the recombination weights. In contrast to *cumulative step-size adaptation* or to the *1/5-th success rule*, the refined two-point adaptation (TPA) does not rely on any internal model of optimality. In contrast to conventional *self-adaptation*, the TPA will achieve a better target step-size in particular with large populations. The disadvantage of TPA is that it relies on two additional objective function evaluations.

Key-words: optimization, evolutionary algorithms, covariance matrix adaptation, step-size control, self-adaptation, two-point adaptation

* Adaptive Combinatorial Search Team, Microsoft Research–INRIA Joint Centre. 28, rue Jean Rostand, 91893 Orsay Cedex, France. email:forename.name@inria.fr.

Adaptation du Pas Deux-Point dans CMA-ES

Résumé : Pas de résumé

Mots-clés : Pas de motclef

1 Introduction

In the Covariance Matrix Evolution Strategy (CMA-ES) [8] two separate adaptation mechanism are performed to determine variances and covariances of the search distribution. One for (overall) step-size control, a second for adaptation of a covariance matrix. The mechanisms are mainly independent and can therefore, in principle, be replaced separately. While the standard step-size control is *cumulative step-size adaptation* (CSA), also a success-based control was successfully introduced for the $(1+\lambda)$ -CMA-ES in [9].

The CSA has a few drawbacks.

- For very large noise levels the target step-size becomes zero, while the optimal step-size is still positive [2].
- For large population sizes ($\lambda > 10n$) the original parameter setting seemed not to work properly [6]. An improved setting shortens the backward time horizon for the cumulation and performs well also with large population sizes [5]. Nonetheless, we tend to conclude that the notion of tracking a (long) path history does not mate well with a large population size (large compared to the search space dimension).
- The expected value for the displacement of the population mean under random selection is required. For this reason, the principle axes of the search distribution are in demand, but they are more expensive to acquire than a simple matrix decomposition. The latter is necessary to sample a multivariate normal distribution with given covariance matrix.

Despite these disadvantages, CSA is regarded as first choice for step-size control in the $(\mu/\mu_w, \lambda)$ -ES, but nonetheless they rise motivation to search for alternatives. Here, we suggest *two-point step-size adaptation* (TSA) as one such alternative.

Two-point self-adaptation was introduced for backpropagation in [11] and later applied in Evolutionary Gradient Search [10]. In evolutionary search, two-point adaptation resembles self-adaptation on the population level. Two different step lengths are tested for the mean displacement and the better one is chosen. In the next section, we integrate a slightly refined TSA in the CMA-ES additionally using polished formulae for the recombination weights and the learning rates of the covariance matrix.

2 The Algorithm: CMA-ES with TPA

Our description of the CMA-ES closely follows [3, 5, 7] and replaces CSA with TSA. Given an initial mean value $\mathbf{m} \in \mathbb{R}^n$, the initial covariance matrix $\mathbf{C} = \mathbf{I}$ and the initial step-size $\sigma \in \mathbb{R}_+$, the new candidate solutions \mathbf{x}_k obey

$$\mathbf{x}_k = \mathbf{m} + \sigma \mathbf{y}_k, \quad \text{for } k = 1, \dots, \lambda, \quad (1)$$

where $\mathbf{y}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$ denotes the realization of a normally distributed random vector with zero mean and covariance matrix \mathbf{C} . The solutions \mathbf{x}_k are evaluated and ranked such that $\mathbf{x}_{i:\lambda}$ becomes the i -th best solution vector and $\mathbf{y}_{i:\lambda}$ the corresponding random vector realization.

For $\mu < \lambda$ let

$$\langle \mathbf{y} \rangle = \sum_{i=1}^{\mu} w_i \mathbf{y}_{i:\lambda}, \quad w_1 \geq \dots \geq w_{\mu} > 0, \quad \sum_{i=1}^{\mu} w_i = 1 \quad (2)$$

be the weighted mean of the μ best ranked \mathbf{y}_k vectors. The recombination weights sum to one. The *variance effective selection mass* is defined as

$$\mu_w = \frac{\sum_{i=1}^{\mu} w_i}{\sum_{i=1}^{\mu} w_i^2} = \frac{1}{\sum_{i=1}^{\mu} w_i^2} \geq 1. \quad (3)$$

From the definition follows that $1 \leq \mu_w \leq \mu$ and $\mu_w = \mu$ for equal recombination weights. The role of μ_w is analogous to the role of the parent number μ when the recombination weights are all equal. Usually $\mu_w \approx \lambda/4$ is appropriate. Weighted recombination is discussed in more detail in [1].

We compute two additional objective function values,

$$f_+ = f(\mathbf{m} + \exp(\alpha) \sigma \langle \mathbf{y} \rangle) \quad (4)$$

$$f_- = f(\mathbf{m} + (2 - \exp(\alpha)) \sigma \langle \mathbf{y} \rangle), \quad (5)$$

where f is the objective function to be minimized, and $\alpha \approx 0.5 \gg \ln(2)$ is the step-size change parameter. The factor $2 - \exp(\alpha)$ in Equation 5 is chosen due to geometrical symmetry considerations. Step-size changes on the one hand should be multiplicative. A factor $\exp(\pm\alpha)$ should be used to realize changes of σ , which is symmetric in the log scale. In contrary, using those factors for generating *test* steps extends one direction, $\exp(+\alpha)$, further than the other, $\exp(-\alpha)$. Assuming the most simple spherical objective function model and optimal step-size, *i.e.* $f(\mathbf{m} + \sigma \langle \mathbf{y} \rangle)$ is minimal for the given σ , the larger test step is disfavored and the step-size will be systematically decreased. Target step-size and optimal step-size disagree. On simple functions, like the sphere model, this effect might even lead to a performance improvement, because the optimum can be approached quickly and therefore the optimal step-size decreases fast. The sub-optimal target step-size “anticipates” this change. Nevertheless, in general, we tend to favor an agreement of target and optimal step-size.

In the remainder, the generation step is completed with the updates of \mathbf{m} , σ , and \mathbf{C} , where two additional state variables will be introduced, $\alpha_s \in \mathbb{R}$ and $\mathbf{p}_c \in \mathbb{R}^n$.

The mean The distribution mean is updated according to

$$\mathbf{m} \leftarrow \mathbf{m} + \sigma \langle \mathbf{y} \rangle. \quad (6)$$

Given σ from Equation (1), Equation (6) can also be written as

$$\mathbf{m} \leftarrow \sum_{i=1}^{\mu} w_i \mathbf{x}_{i:\lambda} . \quad (7)$$

Step-size control: Two-Point Adaptation (TPA) A two-point self-adaptive scheme is implemented based on [10]. Two different step lengths were tested in Equations (4) and (5). Here, the step-size is changed toward the better. Using the values f_+ and f_- we have

$$\alpha_{\text{act}} = \begin{cases} -\alpha + \beta & < 0, \text{ if } f_- \text{ is better (smaller) than } f_+ \\ \alpha & > 0, \text{ otherwise} \end{cases} \quad (8)$$

Initializing $\alpha_s = 0$, the new step-size is calculated according to

$$\alpha_s \leftarrow (1 - c_\alpha) \alpha_s + c_\alpha \alpha_{\text{act}} \quad (9)$$

$$\sigma \leftarrow \sigma \times \exp\left(\frac{\alpha_s}{d_\alpha}\right) \quad (10)$$

where $1/c_\alpha \geq 1$ determines the backward time horizon for smoothing the step-size changes in the generation sequence, and damping $d_\alpha \geq 1$ controls the magnitude of realized changes of σ .

Following [10], the update of \mathbf{m} in Equation (6) could be postponed until after the step-size is updated in Equation (10). Whether or not this results in a better \mathbf{m} cannot be decided without additional costs, because neither the original step-size nor the updated step-size (given $c_\alpha < 1$ or $d_\alpha > 1$) are tested. Equation (7) would not hold anymore. Empirically, using the new step-size leads to slightly higher convergence rates in norm optimization (sphere function) in small dimensions.

The two-point step-size adaptation described here differs from [10] in that smoothing and damping are introduced and the original step-size is used for updating \mathbf{m} . Setting $\alpha = \ln(1.8) \approx 0.588$, $\beta = 0$, $c_\alpha = 1$ and $d_\alpha = 1$, replacing $2 - \exp(\alpha)$ with $\exp(-\alpha)$ in Equation (5) and using the new step-size for updating \mathbf{m} , recovers the step-size adaptation from [10]. We do not expect an essentially different behavior due to our refinements in most cases.

Covariance Matrix Adaptation (CMA) The covariance matrix admits a rank-one and a rank- μ update. For the rank-one update an evolution path \mathbf{p}_c is constructed.

$$\mathbf{p}_c \leftarrow (1 - c_c) \mathbf{p}_c + h_\sigma \sqrt{c_c(2 - c_c)\mu_w} \langle \mathbf{y} \rangle \quad (11)$$

$$\mathbf{C} \leftarrow (1 - c_1 - c_\mu) \mathbf{C} + c_1 \underbrace{\mathbf{p}_c \mathbf{p}_c^T}_{\text{rank-one update}} + c_\mu \underbrace{\sum_{i=1}^{\mu} w_i \mathbf{y}_{i:\lambda} \mathbf{y}_{i:\lambda}^T}_{\text{rank-}\mu \text{ update}} , \quad (12)$$

where $h_\sigma = 0$ if $\alpha_s > (1 - (1 - c_\alpha)^g)(1 - (1 - c_\alpha)^g)\alpha$, and 1 otherwise, where g is the generation counter. The update of \mathbf{p}_c is stalled when α_s is large. The stall is decisive after a change in the environment which demands a significant increase of the step-size. Fast changes of the distribution shape are postponed until after the step-size has increased to a reasonable value.

For the covariance matrix update, the cumulation in (11) serves to capture dependencies between consecutive steps. Dependency information would be lost for $c_c = 1$, because a change in sign of \mathbf{p}_c or $\mathbf{y}_{i:\lambda}$ does not matter in (12).

Parameters The default values for all parameters are shown in Table 1. The changes of parameters compared to [3, 5, 7] are minor polishings.

Recombination weights Compared to [3, 5, 7], where $\mu' = \lceil(\lambda - 1)/2\rceil$ we have chosen $\mu' = (\lambda - 1)/2$. The small difference occurs only for even λ . In the former version, given odd population size λ , the recombination weights did not change when λ was reduced by one. In the present version the recombination weights always adjust to changes of λ .

c_1 and c_μ are the learning rates for the rank-one and rank- μ update of the covariance matrix respectively. In [3, 5, 6, 7], a learning rate $c_{\text{cov}} \approx c_1 + c_\mu$ is used such that $c_1 \approx c_{\text{cov}}/\mu_{\text{cov}}$ and $c_\mu \approx c_{\text{cov}}(\mu_{\text{cov}} - 1)/\mu_{\text{cov}}$. In the former formulation, c_1 was almost two times smaller for values of $\mu_{\text{cov}} \approx 2$ than for $\mu_{\text{cov}} = 1$ and did not monotonously decrease with larger μ_{cov} .

c_α determines the smoothing of α_s . Smoothing and damping suppress stochastic fluctuations of σ . In contrast to damping, the smoothing does not affect the maximal possible change rate for σ . For $c_\alpha \geq 0.5$ we have $\alpha_{\text{act}}\alpha_s > 0$. Signs of the recent measurement and the actual change always agree and the smoothing cannot lead to oscillations. For $c_\alpha \geq 0.3$ only after a second agreeing measure for α_{act} we have always $\alpha_{\text{act}}\alpha_s > 0$. Even smaller values for c_α might be useful, but for much smaller values, presumably the damping must be chosen more carefully.

d_α is the damping parameter and not in use by default. Arguably, for smaller c_α the damping must be increased.

β is the bias parameter for the step-size change. On potentially noisy or highly rugged functions β should be set to 0.1 which results in an effective noise handling.

3 Empirical Validation

In empirical investigations of the TPA-CMA-ES, we find the expected, feasible behavior. The comparison with CSA shows no clear winner. Depending on the objective function either TPA or CSA is faster, but the factor is seldom larger than two. Surprisingly, in our exploratory simulations, there is no clear winner

Table 1: Default Strategy Parameters of the CMA-ES, where n is the problem dimension and $\mu_w \stackrel{\text{def}}{=} \frac{1}{\sum_{i=1}^{\mu} w_i^2}$. The previously defined learning rate [4] amounts to $c_{\text{cov}} \approx c_1 + c_{\mu}$

Selection and recombination:

$$\lambda = 4 + \lfloor 3 \ln n \rfloor, \quad \mu' = \frac{\lambda - 1}{2}, \quad \mu = \lceil \mu' \rceil,$$

$$w_i = \frac{\ln(\mu' + 1) - \ln i}{\sum_{j=1}^{\mu} (\ln(\mu' + 1) - \ln j)} \quad \text{for } i = 1, \dots, \mu,$$

Step-size control:

$$\alpha = 0.5, \quad \beta = 0, \quad c_{\alpha} = 0.3, \quad d_{\alpha} = 1$$

Covariance matrix adaptation:

$$c_{\mathbf{c}} = \frac{4}{n + 4}, \quad \mu_{\text{cov}} = \mu_w$$

$$c_1 = \frac{2}{(n + 1.3)^2 + \mu_{\text{cov}}}, \quad c_{\mu} = \min \left(2 \frac{\mu_{\text{cov}} - 2 + \frac{1}{\mu_{\text{cov}}}}{(n + 2)^2 + \mu_{\text{cov}}}, 1 - c_1 \right)$$

depending on dimension or population size or noise. On noisy functions, setting $\beta = 0.1$ for TPA is quite effective, while we observe only a minor effect from this change otherwise. We did not extensively try to exploit potential weaknesses (as has been done for CSA), but we suspect that the TPA is a feasible and robust alternative to CSA.

4 Conclusion and Outlook

We see some principle **advantages** of using two-point step-size adaptation (TPA) in the CMA-ES.

- The TPA does not rely on a predefined optimality condition, like a success rate of 1/5 or conjugate-perpendicularity of consecutive steps.
- The TPA does not rely on specific properties of the sample distribution or the selection of solutions. Therefore, it is supposably less sensitive to any modifications of the underlying algorithm, in particular compared to CSA.

- The step-size change rate can be adjusted mainly independently from TPA-internal considerations. Time averaging or damping are not essentially necessary.

Even so, we see two principle **disadvantages** of TPA.

- Two additional function evaluations are needed per iteration step. This is not a grave disadvantage, in particular when the population size is not very small. As a remedy, these two points could be incorporated in the population and used to compute the (final) mean in Equation (7), while they are of limited use in the rank- μ update of the covariance matrix.
- Step-size control is based on two objective function evaluations only. Selection information from the remaining population (and history information) is somewhat disregarded. This is a conceptional defect, that might be irrelevant in practice.

In conclusion, two-point step-size adaptation is an alternative to cumulative step-size adaptation well worth of further exploration. Whether and when it should finally replace CSA in practice must be answered in future empirical studies.

References

- [1] D.V. Arnold. Weighted multirecombination evolution strategies. *Theoretical Computer Science*, 361(1):18–37, 2006.
- [2] H.-G. Beyer and D. Arnold. Qualms regarding the optimality of cumulative path length control in CSA/CMA-evolution strategies. *Evolutionary Computation*, 11(1):19–28, 2003.
- [3] N. Hansen. The CMA evolution strategy: a comparing review. In J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea, editors, *Towards a new evolutionary computation. Advances on estimation of distribution algorithms*, pages 75–102. Springer, 2006.
- [4] N. Hansen, F. Gernerle, A. Auger, and P. Koumoutsakos. When do heavy-tail distributions help? In T. P. Runarsson et al., editors, *Parallel Problem Solving from Nature PPSN IX*, volume 4193 of *Lecture Notes in Computer Science*, pages 62–71. Springer, 2006.
- [5] N. Hansen and S. Kern. Evaluating the CMA evolution strategy on multimodal test functions. In X. Yao et al., editors, *Parallel Problem Solving from Nature PPSN VIII*, volume 3242 of *LNCS*, pages 282–291. Springer, 2004.
- [6] N. Hansen, S. D. 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.

-
- [7] N. Hansen, S. P. N. Niederberger, L. Guzzella, and P. Koumoutsakos. A method for handling uncertainty in evolutionary optimization with an application to feedback control of combustion. *IEEE Transactions on Evolutionary Computation*, 200?. Accepted for publication.
 - [8] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.
 - [9] C. Igel, N. Hansen, and S. Roth. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation*, 15(1):1–28, 2007.
 - [10] R. Salomon. Evolutionary algorithms and gradient search: similarities and differences. *Evolutionary Computation, IEEE Transactions on*, 2(2):45–55, 1998.
 - [11] R. Salomon and J. L. van Hemmen. Accelerating backpropagation through dynamic self-adaptation. *Neural Networks*, 9(4):589–601, 1996.



Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399