Refined CMA-ES for CEC 2026 Bound-Constrained Single-Objective Optimization

Adam Stelmaszczyk*, Rafał Biedrzycki*, Jarosław Arabas*
*Institute of Computer Science, Warsaw University of Technology, Warsaw, Poland
{adam.stelmaszczyk, rafal.biedrzycki, jaroslaw.arabas}@pw.edu.pl

***Abstract*—We present a simplified and parameter-adjusted variant of CMA-ES designed for the CEC 2026 Bound-Constrained Single-Objective Optimization competition. Unlike typical algorithmic developments that increase complexity by adding mechanisms, Refined CMA-ES systematically removes components that do not improve performance, such as the flat-land escape, while retaining standard features that provide measurable benefits. Adjusting key parameters—population size and initial step-size—yields substantial gains, outperforming the baseline CMA-ES and several state-of-the-art competitors. The paper highlights the practical relevance of individual CMA-ES components under the CEC 2026 ranking, which equally weights solution accuracy and convergence speed, demonstrating that simplification combined with improved parameter settings can lead to superior performance.**

***Index Terms*—CMA-ES, evolutionary algorithms, single-objective optimization, black-box optimization**

I. Introduction

In evolutionary computation, algorithmic development typically follows an additive paradigm: a baseline method is taken and additional mechanisms are progressively introduced in order to improve performance. Examples include restart strategies, adaptive population sizes or hybridization. Over time, this process often leads to increasingly complex algorithms, with the effects of individual components not always clearly assessed.

In this work, we adopt the opposite approach. Instead of adding new components to CMA-ES, we systematically analyze the impact of its existing mechanisms using the CEC 2026 benchmark. The goal is not to increase algorithmic complexity, but to identify which components are actually beneficial in this specific evaluation setting, and which can be removed without harming performance—or even improving it.

From this perspective, the CEC 2026 benchmark is treated not only as a competition platform, but also as a diagnostic tool for performing an ablation study of CMA-ES. By measuring the effect of removing selected components and adjusting key parameters, we obtain insights into the practical relevance of several design choices in the original algorithm.

An important aspect of this study is the specific nature of the CEC 2026 ranking itself. Unlike most traditional benchmarks that focus solely on final solution quality, the CEC 2026 ranking assigns equal importance to two criteria: solution quality and convergence speed. Half of the total score is determined by the final error values, while the other half measures how quickly an algorithm reaches the solution it ultimately reports as the result of the optimization process, according to the official evaluation protocol [1].

This hybrid evaluation scheme significantly alters the optimization landscape. Mechanisms that are beneficial for long-term fine-tuning of solutions may become less relevant, or even detrimental, when early convergence is heavily rewarded. As a result, some classical components of CMA-ES, originally designed for robust performance, may no longer be useful under this ranking protocol. This makes the CEC 2026 benchmark particularly well suited for analyzing which parts of CMA-ES are truly effective in practice under time-sensitive evaluation.

The main contributions of this paper are:

- an assessment of the covariance matrix decomposition delay in CMA-ES, highlighting its limited application,
- an empirical evaluation of the flat-land escape mechanism, which was found to degrade performance on CEC 2026 benchmark,
- an empirical evaluation indicating that the h_σ correction, which stalls evolution path adaptation when the step-size becomes too small, provides a small performance improvement on the CEC 2026 benchmark,
- a demonstration that parameter settings alone yield substantial performance gains,
- the proposal of Refined CMA-ES (R-CMA-ES), a simplified and parameter-adjusted variant that outperforms the baseline CMA-ES and several state-of-the-art competitors.

The remainder of this paper is organized as follows. Section II describes our baseline, the CMA-ES algorithm. Section III details the methodology used to measure the impact of each modification. Section IV introduces R-CMA-ES, provides a detailed description of the modifications, and defines the algorithm variations. Section V presents pairwise and total-point comparisons of these variants. Section VI compares R-CMA-ES with selected existing methods. Finally, Section VII discusses implementation details and reports on the algorithmic complexity of R-CMA-ES.

II. CMA-ES

As a base algorithm, we use the state-of-the-art Covariance Matrix Adaptation Evolution Strategy (CMA-ES) as described in [2]. CMA-ES is a population-based, derivative-free optimization algorithm designed for continuous black-box problems. At each generation, a population of candidate solutions is sampled from a multivariate normal distribution

$$x_i \sim \mathcal{N}(m, \sigma^2 C),$$

where m is the current mean, σ is the global step-size, and C is the covariance matrix.

The mean vector m is updated using a weighted recombination of the best individuals in the population, and the covariance matrix C is adapted over time, enabling an implicit form of second-order optimization.

In addition to mean and covariance matrix adaptation, CMA-ES employs step-size control based on the evolution path mechanism, which adjusts σ to regulate the overall scale of the search steps. Beyond these core principles, several auxiliary mechanisms have been incorporated into the standard algorithm over time, including covariance matrix decomposition delay, flat-land escape, and the h_σ correction.

III. Methodology

The CEC 2026 ranking [1] evaluates algorithms based on both solution accuracy and convergence speed, assigning equal weight to each criterion. For pairwise comparisons, each algorithm is run for 25 independent trials on each of the 29 benchmark functions. This yields $50 \cdot 49/2 = 1225$ pairwise trial comparisons per function, resulting in a maximum of 35525 points per objective. Combining accuracy and speed objectives, a total of 71050 points can be awarded for each pair of algorithms.

Table I reports pairwise differences, computed as the difference between the CEC 2026 ranking points achieved by algorithm 1 and algorithm 2. Positive values indicate that the first algorithm outperformed the second. In contrast, Tables II and IV present the CEC 2026 ranking for several algorithms, where total points across all functions and objectives are aggregated to rank all compared algorithms simultaneously.

This methodology allows us to quantify both absolute performance and relative differences between algorithms, providing a clear view of which modifications and parameter settings yield measurable improvements under the hybrid accuracy-speed evaluation.

IV. Refined CMA-ES

R-CMA-ES is a parameter-adjusted CMA-ES in which the covariance matrix decomposition delay and flat-land escape mechanisms have been removed. The following subsections provide a detailed description of these modifications and define the A, B, C algorithm variations.

A. No covariance matrix decomposition delay

In CMA-ES, the covariance matrix decomposition is delayed by $\max(1, \lfloor 0.1 \cdot n \cdot c_{cov} \rfloor)$ generations, where the learning coefficient c_{cov} depends on the problem dimension n. With the default CMA-ES parameters, $\lambda = 4 + \lfloor 3 \log(n) \rfloor$, this delay is greater than one generation only for dimensions $n \geq 190$. When $\lambda \geq n$, decomposition occurs every generation, resulting in a per-generation computational complexity of $O(n^3)$. Therefore, for $n < 190$, the delay mechanism is effectively inactive, and we simplify the implementation by removing it without affecting the observed behavior.

B. No flat-land escape

Step-size σ is increased when multiple individuals share the same fitness, aiming to escape flat regions of the fitness landscape more quickly. However, as noted in [2], this mechanism has certain drawbacks: "This method interferes with the termination criterion `TolFun`. In practice, observation of a flat fitness should rather be a termination criterion and consequently lead to a reconsideration of the objective function formulation." By removing the flat-land escape, these issues are avoided, and we observed an improvement in overall performance on the CEC 2026 benchmark.

C. Non-default parameter settings

By default, CMA-ES uses a population size $\lambda = 4 + \lfloor 3 \log(n) \rfloor$ and an initial step-size $\sigma = 0.5$. Inspired by RB-IPOP-CMA-ES [3], we adjusted these two parameters. First, the population size was increased to $\lambda = 4n$, corresponding to an increase from 14 to 120 in the 30-dimensional CEC 2026 benchmark. Second, the initial step-size σ was increased from 0.5 to 7. Both changes result in substantial performance gains, considerably larger than removing the flat-land escape, as shown later.

D. Other modifications

In addition to the modifications reported above, we also examined the impact of removing the correction of the evolution path adaptation h_σ, which is a standard component of the original CMA-ES. In contrast to decomposition delay and flat-land escape, removing h_σ caused a slight performance decrease. Consequently, h_σ was retained in the final R-CMA-ES.

Furthermore, we evaluated several mechanisms commonly reported in the CMA-ES literature, such as IPOP restarts [4], bound constraint handling by resampling [5], hybridization with a local optimizer (L-BFGS-B), the use of midpoint evaluation, and starting point presampling [3]. None of these extensions led to performance improvements under the CEC 2026 evaluation. In particular, restart strategies resulted in a minor improvement in accuracy, but this was offset by a substantial drop in convergence speed, which is heavily weighted in the CEC 2026 ranking.

Table I presents pairwise comparisons of the CMA-ES variants using the CEC 2026 ranking. Algorithm A is omitted, as its observed behavior is identical to the baseline CMA-ES in 30 dimensions.

Variant AB (no delay + no escape) won against the original CMA-ES by 1615 ranking points. AC (no delay + parameter settings) outperformed the original CMA-ES by a much larger margin of 27816 points, showing that parameter settings alone provide a significant performance improvement. This is further confirmed by comparing AC and AB: AC won by 27529

TABLE I
PAIRWISE COMPARISONS OF CMA-ES VARIANTS USING CEC 2026 RANKING POINTS

algorithm 1 \ algorithm 2	CMA-ES	AB	AC
AB	1 615	—	—
AC	27 816	27 529	—
R-CMA-ES (ABC)	28 308	28 069	1 275

points. Finally, R-CMA-ES (no delay + no escape + parameter settings) beats AC by 1275 points, a much smaller margin.

Table II presents the CEC 2026 ranking for the considered CMA-ES variants based on total CEC 2026 points. R-CMA-ES takes first place with 100601 points. The AC variant follows in second place with 98810 points, trailing by 1791 points, confirming that the improved parameter choices account for most of the performance gain. The AB variant ranks third with 44783 points, showing that the removal of the flat-land escape provides only a small gain of 1878 points compared to the baseline CMA-ES.

TABLE II
CEC 2026 RANKING OF CONSIDERED CMA-ES VARIANTS

place	algorithm	points
1	**R-CMA-ES (ABC)**	100 601
2	AC	98 810
3	AB	44 783
4	CMA-ES	42 905

Table III presents the mean of the minimum error values (Min_EV) obtained across 25 independent trials for 30 dimensional problems. The "Test" column reports the result of the Wilcoxon signed-rank test comparing the 25 Min_EV values. "=" indicates no statistically significant difference (p = 0.05). "+" indicates that R-CMA-ES performed significantly better, whereas "–" indicates it performed worse. Overall, in 30 dimensions, R-CMA-ES performed significantly better on 23 out of 29 functions and worse on one function (F3).

TABLE III
MINIMUM ERROR VALUES AVERAGED OVER 25 TRIALS IN 30 DIMENSIONS

Function	R-CMA-ES	CMA-ES	Test
1	**0.0e+00**	1.1e-14	+
2	1.0e+02	1.0e+02	=
3	1.6e+02	**1.4e+02**	–
4	**1.1e+02**	7.8e+02	+
5	**1.0e+02**	1.9e+02	+
6	**1.4e+02**	2.8e+03	+
7	**1.1e+02**	7.0e+02	+
8	**1.0e+02**	1.5e+04	+
9	**2.1e+03**	5.3e+03	+
10	**2.0e+02**	3.1e+02	+
11	1.5e+03	1.6e+03	=
12	**1.9e+02**	2.4e+03	+
13	**2.1e+02**	3.0e+02	+
14	**3.6e+02**	4.7e+02	+
15	**3.1e+02**	9.3e+02	+
16	**1.9e+02**	6.3e+02	+
17	3.3e+02	3.3e+02	=
18	**2.0e+02**	3.2e+02	+
19	**4.2e+02**	1.4e+03	+
20	**3.1e+02**	6.2e+02	+
21	**1.1e+03**	6.1e+03	+
22	**4.6e+02**	2.1e+03	+
23	**5.2e+02**	6.2e+02	+
24	**4.9e+02**	4.9e+02	+
25	9.8e+02	1.3e+03	=
26	**6.1e+02**	9.0e+02	+
27	4.5e+02	4.4e+02	=
28	**5.6e+02**	1.1e+03	+
29	**2.3e+03**	2.7e+03	+
Wins	23	1	

TABLE IV
COMPARISON WITH SELECTED EXISTING METHODS ON THE CEC 2026 BENCHMARK

place	algorithm	points
1	**R-CMA-ES**	109 816
2	L-SRTDE	98 387
3	RB-IPOP-CMA-ES	85 123
4	CMA-ES	79 213
5	EA4EigSimpTowardsIDE	76 960

V. COMPARISON WITH SELECTED EXISTING METHODS

Table IV presents a comparison of Refined CMA-ES with the baseline CMA-ES and three other notable algorithms: L-SRTDE [6], EA4EigSimpTowardsIDE [7], and RB-IPOP-CMA-ES [3]. L-SRTDE was the winner of the CEC 2024 competition, EA4EigSimpTowardsIDE is an improved simplification of EA4Eig, which won CEC 2022, and RB-IPOP-CMA-ES reached 7th place in CEC 2017.

Refined CMA-ES achieved the highest score, outperforming the runner-up L-SRTDE by 11428 points. The relatively lower performance of RB-IPOP-CMA-ES can be attributed to its use of IPOP restarts: restart strategies improve the accuracy but cause a substantial drop in convergence speed, which is heavily weighted in the CEC 2026 ranking. Overall, these results demonstrate that selective simplification combined with improved initial parameter settings yields superior performance compared to both the baseline CMA-ES and other state-of-the-art methods.

VI. IMPLEMENTATION NOTES AND ALGORITHM COMPLEXITY

The full source code is available at https://github.com/AdamStelmaszczyk/refined-cma-es. The experiments were conducted on a Linux Debian system equipped with an Intel Xeon CPU E7-4830 v3, with a base clock of 2.1 GHz and turbo frequency up to 2.7 GHz. R-CMA-ES is a single-threaded R implementation based on the `cmaes` package [8]. The CEC 2026 benchmark relies on CEC 2017 functions, for which we used the `cec2017` package [9]. The algorithmic complexity was assessed following the methodology described in [1]. Table V reports the timing results. T_1 denotes the time required for 10000 function evaluations, averaged over the 29 CEC functions. T_2 denotes the time required to run R-CMA-ES with a budget of 10000 evaluations, also averaged over the 29 functions. In 30 dimensions, approximately 90% of the total runtime is spent on function evaluations.

TABLE V
ALGORITHM COMPLEXITY

T_1	T_2	$(T_2 - T_1)/T_1$
3.32s	3.66s	0.1

VII. CONCLUSIONS

This paper introduced Refined CMA-ES, a simplified and parameter-adjusted variant of CMA-ES, which achieves superior performance on the CEC 2026 benchmark compared to the state-of-the-art methods. Refined CMA-ES outperforms L-SRTDE, the winner of CEC 2024, demonstrating that simplification combined with improved parameter settings can yield substantial gains. The study also shows that certain CMA-ES mechanisms, such as covariance matrix decomposition delay and flat-land escape, are not beneficial under this evaluation, while adjusting key parameters—population size and initial step-size—has a strong positive impact on both accuracy and convergence speed.

ACKNOWLEDGMENT

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: CI TASK) for providing computer facilities and support within computational grant no. PLG/2025/018266.

REFERENCES

[1] P. Chen, H. Lin, K. Qiao, X. Ban, K. V. Price, C. Yue, J. Liang, V. Stanovov, and P. N. Suganthan, "Evaluation criteria for CEC 2026 competition and special session on bound constrained single and multi-objective optimization considering both accuracy and speed," GitHub, Tech. Rep., 2026. [Online]. Available: https://github.com/P-N-Suganthan/2026-CEC/blob/main/TR_BoundConst-MOP-SOP.pdf

[2] N. Hansen, "The CMA evolution strategy: A tutorial," *CoRR*, vol. abs/1604.00772, 2016. [Online]. Available: http://arxiv.org/abs/1604.00772

[3] R. Biedrzycki, "A version of IPOP-CMA-ES algorithm with midpoint for CEC 2017 single objective bound constrained problems," in *2017 IEEE Congress on Evolutionary Computation (CEC)*, 2017, pp. 1489–1494.

[4] A. Auger and N. Hansen, "A restart cma evolution strategy with increasing population size," in *2005 IEEE Congress on Evolutionary Computation*, vol. 2, 2005, pp. 1769–1776 Vol. 2.

[5] R. Biedrzycki, "Handling bound constraints in CMA-ES: An experimental study," *Swarm and Evolutionary Computation*, vol. 52, no. 100627, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210650219301622

[6] V. Stanovov and E. Semenkin, "Success rate-based adaptive differential evolution L-SRTDE for CEC 2024 competition," in *2024 IEEE Congress on Evolutionary Computation (CEC)*, 2024, pp. 1–8.

[7] R. Biedrzycki, "Analysis and simplification of the winner of the CEC 2022 optimization competition on single objective bound constrained search," *Evolutionary Computation*, pp. 1–18, 07 2025.

[8] H. Trautmann, O. Mersmann, and D. Arnu, *cmaes: Covariance Matrix Adapting Evolutionary Strategy*, The Comprehensive R Archive Network (CRAN), 2022, r package version 1.0-12. [Online]. Available: https://CRAN.R-project.org/package=cmaes

[9] D. Jagodziński. Common interface to the 30 CEC-2017 benchmark functions - special session and competition on single objective real-parameter numerical optimization - bound constrained case. [Online]. Available: https://staff.elka.pw.edu.pl/~djagodzi/programy.html