



Universiteit
Leiden
The Netherlands

Avoiding redundant restarts in multimodal global optimization

Nobel, J.P. de; Vermetten, D.L.; Kononova, A.V.; Shir, O.M.; Bäck, T.H.W.; Affenzeller, M.; ... ; Machado, P.

Citation

Nobel, J. P. de, Vermetten, D. L., Kononova, A. V., Shir, O. M., & Bäck, T. H. W. (2024). Avoiding redundant restarts in multimodal global optimization. *Lecture Notes In Computer Science*, 268-283. doi:10.1007/978-3-031-70068-2_17

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/4254645>

Note: To cite this publication please use the final published version (if applicable).

Avoiding Redundant Restarts in Multimodal Global Optimization

Jacob de Nobel¹, Diederick Vermetten¹, Anna V. Kononova¹ Ofer M. Shir², and Thomas Bäck¹

¹ LIACS, Leiden University, Leiden, The Netherlands

² Tel-Hai College and Migal Institute, Upper Galilee, Israel

Abstract. Naïve restarts of global optimization solvers when operating on multimodal search landscapes may resemble the Coupon’s Collector Problem, with a potential to waste significant function evaluations budget on revisiting the same basins of attractions. In this paper, we assess the degree to which such “duplicate restarts” occur on standard multimodal benchmark functions, which defines the *redundancy potential* of each particular landscape. We then propose a repelling mechanism to avoid such wasted restarts with the CMA-ES and investigate its efficacy on test cases with high redundancy potential compared to the standard restart mechanism.

Keywords: numerical optimization, multimodal landscapes, CMA-ES

1 Introduction

Finding an effective balance between exploring the domain and exploiting promising regions is one of the primary challenges when designing any iterative optimization heuristic, being subject to an underlying hard conflict. This design choice is especially important in multimodal search landscapes, where premature convergence leads to poor overall performance.

Because of its challenging nature, a large number of algorithms have been developed for multimodal optimization [22], yet consisting of two main branches with different goals — (i) niching methods [24], whose target is adjusted from locating the globally optimal solution to finding a wider set of high-quality optima, and (ii) “upgraded global solvers”, which still target a single global optimum, but are better equipped to handle multimodality (see, e.g., [9,3]).

While niching methods offer multiple approaches to treating such multimodal landscapes, primarily in promoting population diversity, these practices are usually not transferred to global optimization for various reasons (for instance, since focusing on solution diversity can hamper performance by slowing down convergence). Instead, global solvers often rely on rebalancing the search from global to local over time, e.g., by adjusting step sizes in evolutionary algorithms. To prevent convergence to sub-optimal solutions in these methods, restart mechanisms can be utilized to reset the algorithm’s state, essentially starting a new optimization process by using the remaining function evaluation budget. This process is usually independent of the previous trajectory but can update the algorithm’s strategy parameters, such as the population size within CMA-ES.

However promising a restart scheme can be, it still carries the potential to encounter the equivalent of the **Coupon Collector’s Problem**, which is rooted in the retrials’ naïvity. Given q unique coupons, the expected number of trials needed to collect them all, with replacement, is $q \cdot H_q$ (with $H_q := \sum_{t=1}^q \frac{1}{t}$ being the q^{th} -harmonic number). Similarly, in the context of multimodal optimization, a straightforward approach of *iteration* can be used to locate sequentially multiple peaks in the landscape by means of an *iterative local search* (ILS) [14]. If the procedure is blind to any information accumulated throughout previous runs, and it sequentially restarts stochastic search processes, the ambition to hit a different peak in every run resembles the collector’s hope to obtain all the coupons in only q trials. Overall, it is likely to encounter *redundancy*, and the number of expected iterations is then increased by a factor. A **redundancy factor** can be derived if the peaks are of equal height (an *equi-fitness landscape*, i.e., the probability to converge into any of the q peaks is uniform and equal to $1/q$), simply by normalizing the expected number of trials with respect to q - that is an overall redundancy factor of H_q [24]. Importantly, when dropping the equi-fitness assumption, this factor is expected to increase.

Mind should be given to this analogy and to the actual target of the “collector”: the careful reader must note that a *global* multimodal solver is not necessarily concerned with “collecting” all

the optima (as in niching), since it is targeted at attaining only the best (i.e., picking only the “top coupon”). At the same time, when operating in a black-box fashion, the global quality of such optima is often *undecidable*. Therefore, the “collector” has no choice but to start a campaign of iterative restarts.

The idea of “avoiding duplicates” is a fundamental concept in heuristic search (e.g., the classical Tabu Search [5]), and even in statistical sampling (e.g., the conditioned Latin Hypercube Sampling that produces sampling designs while avoiding redundancy by accounting for external information [16]). To investigate this idea in the context of multimodal optimization, we propose an alternative restart strategy in the context of the renowned CMA-ES algorithm.

We provide the necessary background, especially concerning basins of attraction of local optima, and an illustrative motivation, in Section 2, and identify the potential evaluation redundancy problem caused by restarts in global optimization. Then, we analyze the extent of this redundancy phenomenon on standard continuous landscapes, which motivates our proposal for an alternative restart strategy (Section 3). We propose and benchmark a restart strategy that integrates ideas from multimodal optimization, specifically *repelling subpopulations in niching*, with the existing restart practices of CMA-ES. This new algorithm, the *CMA-ES with repelling restarts* (RR-CMA-ES) is introduced in Section 4. This method is benchmarked on a wide set of problems, which shows that while wasted evaluations from “duplicate” restarts can be avoided, the precise methods to achieve this potential must be carefully calibrated to prevent deteriorating performance on some types of landscapes (Section 5). Conclusions and future work are discussed in Section 6.

2 Preliminaries and Problem Formulation

2.1 Basins of Attraction in Global Optimization

We consider the *global optimization* challenge of a single-objective, continuous minimization problem; the aim is to identify a *single* solution $\mathbf{x}^* \in \mathcal{X} \subseteq \mathbb{R}^d$ from the feasible region $\mathcal{X}$, which minimizes a given objective function $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$:

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \quad (1)$$

When the convexity property does not hold for f [4], it is often the case that other candidate solutions $\mathbf{x}_\ell \neq \mathbf{x}^*$ minimize certain neighborhoods of radii ϵ , and thus each forms a *local minimum*:

$$\exists \epsilon > 0 \quad \forall \mathbf{x} \in \mathcal{X} : \|\mathbf{x} - \mathbf{x}_\ell\| < \epsilon \Rightarrow f(\mathbf{x}_\ell) \leq f(\mathbf{x})$$

$\mathbf{x}^*$ is called the *global optimizer*. Accordingly, the process of global optimization [27] is considered successful when it concludes with locating it while escaping *local optimizers* $\mathbf{x}_\ell$ (also referred to as *traps*). Next, we would like to define the attraction basin of an optimizer (denoted as $\hat{\mathbf{x}}$, being either local or global), following [27] (alternative definitions exist - see, e.g., [2]). Given the standard Gradient Descent Algorithm [4], which is defined by its variation step from iteration i to $i+1$ (with $\sigma(i)$ being the step-size),

$$\mathbf{x}(i+1) := \mathbf{x}(i) - \sigma(i) \cdot \nabla f(\mathbf{x}(i)), \quad (2)$$

we denote its initial search-point $\mathbf{x}(0) := \mathbf{x}_0$. The following set of points is defined whenever the limit $\lim_{i \rightarrow \infty} \mathbf{x}(i)$ exists:

$$\Omega = \left\{ \mathbf{x} \in \mathbb{R}^d \mid \mathbf{x}(0) = \mathbf{x} \wedge \mathbf{x}(i)|_{i \geq 0} \text{ satisfies (2)} \wedge \lim_{i \rightarrow \infty} \mathbf{x}(i) \text{ exists} \right\} \quad (3)$$

Then, given an optimizer $\hat{\mathbf{x}}$, we define its region of attraction using Ω :

$$\mathcal{A}(\hat{\mathbf{x}}) = \left\{ \mathbf{x} \in \Omega \mid \mathbf{x}(0) = \mathbf{x} \wedge \mathbf{x}(i)|_{i \geq 0} \text{ satisfies (2)} \wedge \lim_{i \rightarrow \infty} \mathbf{x}(i) = \hat{\mathbf{x}} \right\} \quad (4)$$

Finally, and most importantly, the *basin* of $\hat{\mathbf{x}}$ is the **maximal level set** that is fully contained in $\mathcal{A}(\hat{\mathbf{x}})$.

We are particularly interested in addressing this global optimization challenge within black-box optimization, where the analytical form of the objective function f is unknown to the solver and is thus treated as a black box, which, upon receiving input, yields a (continuous) output.

2.2 CMA-ES

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [10] is a state-of-the-art method for single objective black-box optimization. The CMA-ES is a stochastic method that evolves a population of candidate solutions $\mathbf{x}$ to an optimization problem $f : \mathbb{R}^d \rightarrow \mathbb{R}$. To guide the search, it uses a parameterized multivariate normal distribution $\mathcal{N}(\mathbf{m}, \sigma\mathbf{C})$. The defining feature of the CMA-ES is that it can adapt the parameters of its mutation distribution to an arbitrary shape during optimization to guide the search. This makes it invariant to linear transformations of the search space and underpins its effectiveness on non-separable and ill-conditioned problems.

2.3 Restart Mechanisms

Restarts are commonly employed to prevent premature convergence in optimization algorithms. A whole class of algorithms relies on restart mechanisms to adapt local search methods to global optimization problems, dubbed as *Iterative Local Search* (ILS) [14]. While ILS is mostly prevalent in discrete domains, techniques for continuous domains such as Multi-Level Single Linkage (MLSL) [23] have shown promising performance on commonly used benchmark problems [21]. While restart methods are common in combination with local search, they can also be applied to more global optimizers, such as evolutionary algorithms. In its simplest form, for the CMA-ES, restarting involves resetting its mutation distribution to a standard multivariate normal distribution and setting the center of mass $\mathbf{m}$ to a new location u.a.r. whenever *local* restart criteria have been met. Additionally, the parameters of the algorithm can be altered to change its behavior following a restart, of which the most well-known strategies are IPOP [3] and BIPOP [6]. The IPOP-CMA-ES increases its population size by a factor of 2 on every consecutive restart. The BIPOP-CMA-ES alternates between two regimes; the first uses an increased population size and a larger value for the initial step size σ^0 , while the second uses a smaller population size and σ^0 . While other restart strategies have been proposed [13,17], IPOP and BIPOP remain among the most commonly used methods.

2.4 Motivation

While the aforementioned restart mechanisms have been shown to be effective on multimodal problems [6], no inherent method is implemented to ensure a subsequent restart does not converge to the same basin of attraction. By restarting u.a.r. it could very well be possible that several runs will be drawn to the same attractor. When this happens, the restart effectively spends more of the evaluation budget to potentially find an already identified solution. If this could be avoided by ensuring each consecutive restart only explores a previously unvisited region of the search space, we could potentially save a portion of the evaluation budget.

Consider the example shown in Figure 1, which shows a single run of the CMA-ES with a simple restart strategy solving a modified version of the 2-dimensional Himmelblau function:

$$f_{mh}(\mathbf{x}) = (x_1^2 + x_2 - 11) + (x_1 + x_2^2 - 7)^2 + \min\{0.01, \|\mathbf{x} - \mathbf{x}^*\|_2\} \quad (5)$$

Here, $\mathbf{x}^*$ is the location of the global optimum, and $\|\cdot\|_2$ denotes the Euclidean norm. From the figure, we can observe that the CMA-ES restarts several times during the optimization process. Only during one of the restarts does the CMA-ES converge to the global optimum (green trace), while the other three local optima are found during other restarts. Notably, it can be observed that a considerable number of restarts are converging to local optima that have already been identified (red traces). In fact, this is the case for eight out of twelve restarts for this specific run. If we had been able to avoid revisiting previously found optima, considerable savings to the evaluation budget could have been realized. While one might argue that this is just a bad run or that a change to the initial σ could avoid this type of behavior, we aim to demonstrate in the next section that this is not the case.

2.5 Repelling Subpopulations

The CMSA with *Repelling Subpopulations* (RS-CMSA) [1] has been proposed as a niching strategy that uses a number of parallel subpopulations to find multiple local optima in multimodal optimization. Our approach is inspired by the RS-CMSA, but it focuses on avoiding repeatedly sampling




Fig. 1: A single run of the CMA-ES with a simple u.a.r. restart strategy ($\sigma^0 = 2$), optimizing a modified version of the Himmelblau function (f_{mh} , see eq. 5). Every restart is visualized as a line, which shows the trace followed by the CMA-ES, where the circle shows the start of the trace and the star shows the finally obtained solution. The restart that converges to the global optimum is shown in green, and restarts converging to a new local optimum are shown in yellow. The traces in red show restarts that converge to local optima, which have already been found during previous restarts (yellow).

points in the basin of attraction of a local optimum that the algorithm has already identified. The overall goal is to make the CMA-ES more sample-efficient for the goal of finding a single best solution (i.e., the best local optimum approximation, given a budget B of function evaluations), rather than multiple solutions representing different local optima (i.e., multiple niches). The underlying methods, such as the usage of the Hill-Valley heuristic, restarts, and tabu regions, are inspired by the RS-CMSA.

2.6 Hill-Valley Function as a Boolean Heuristic

The Hill-Valley function $HV(\mathbf{x}^i, \mathbf{x}^j, f) \in \{0, 1\}$, listed as Algorithm 1, introduced as a part of the multinational EA [28], can be used as a heuristic to determine whether two points $\mathbf{x}^i$ and $\mathbf{x}^j$ belong to the same basin of attraction [15]. It calculates the objective function value for a maximum of $N_t = 10$ points on a line drawn between $\mathbf{x}^i$ and $\mathbf{x}^j$. If any of these points are of a higher objective function value, assuming minimization, a hill is assumed to be present between $\mathbf{x}^i$ and $\mathbf{x}^j$, which then do not share a basin of attraction.

3 The Potential Gain of Avoiding Redundant Restarts

3.1 Defining the Redundancy Measure

Motivated by the *redundancy factor* mentioned earlier in the context of the Coupon’s Collector Problem, we would like to define a measure that quantifies the gain potential for having effective restarts. The trials of the Coupon’s Collector must be adapted to the notion of restarts and to account for the global optimum, which constitutes the “top priority coupon” when following the analogy. To this end, we define the *restarts’ redundancy factor* as the proportion of function

Algorithm 1 Hill-Valley test $\text{HV}(\mathbf{x}^i, \mathbf{x}^j, f)$

Require: $\mathbf{x}^i \in \mathbb{R}^d, \mathbf{x}^j \in \mathbb{R}^d, N_t \in \mathbb{N}, f : \mathbb{R}^d \rightarrow \mathbb{R}$
for $k \in 1, \dots, N_t$ **do**
 $\mathbf{x}^{test} = \mathbf{x}^i + \frac{k}{N_t+1}(\mathbf{x}^j - \mathbf{x}^i)$
 if $\max\{f(\mathbf{x}^i), f(\mathbf{x}^j)\} \leq f(\mathbf{x}^{test})$ **then**
 return 0
 end if
end for
return 1

evaluations spent by *duplicate restarts* (different restarts that converged to a basin of attraction that was previously visited). This calculation excludes function evaluations spent within the global basin (regardless of being duplicates), denoted as $\mathbf{x}^*$ – and it accumulates such function evaluations and normalizes them by the total budget. With r denoting the current number of restarts, having the current restart converging to a point $\mathbf{x}^{(r)}$, we use $\text{red}(\mathbf{x}^{(r)}) \in \{0, 1\}$ to determine whether the restart is redundant, relying on an oracle that answers the boolean query `sameBasin()` indicative of two points sharing a basin of the function f :

$$\text{red}(\mathbf{x}^{(r)}) := \left[\neg \text{sameBasin}(\mathbf{x}^*, \mathbf{x}^{(r)}, f) \right] \wedge \left[\bigvee_{1 \leq k < r} \text{sameBasin}(\mathbf{x}^{(k)}, \mathbf{x}^{(r)}, f) \right] \quad (6)$$

Importantly, the previously defined Hill-Valley heuristic may approximate the boolean query `sameBasin()`, and will play this role in our implementation.

We then define the *restarts' redundancy factor* as the normalized accumulated redundancy of a given set of restarts:

$$\text{RRF} \left(\left\{ \mathbf{x}^{(r)} \right\}_{r=1}^R \right) := \frac{\sum_{r=1}^R \text{red}(\mathbf{x}^{(r)}) b^{(r)}}{B}, \quad (7)$$

where $b^{(r)}$ is the number of function evaluations spent by a restart, R is the total number of completed restarts, and B is the total budget of function evaluations spent. Importantly, this factor accounts for a given set of restarts rather than measuring a search algorithm/mechanism (i.e., an implicit algorithmic measure).

3.2 Numerical Assessment of Expected Redundancy Factors

We analyze the CMA-ES with three different restart strategies for several benchmark functions to identify the potential performance gains from avoiding duplicated restarts in global optimization. We compare a naive restart strategy (labeled ‘RESTART’) alongside the popular IPOP and BIPOP restart strategies. Each strategy places the center of mass $\mathbf{m}$ uniformly at random in the domain on each restart and uses saturation for bound correction. The remaining settings of CMA-ES are left as default in the modCMA package [19]. For each benchmark function, we perform 100 independent runs. In addition to logging the performance trajectories with IOHExperimenter [20], we log the center of mass and the number of evaluations used each time a restart is triggered. Based on this information, we can use Equation 7 to calculate the RRF for a given experiment.

BBOB We benchmark on the noiseless, single-objective BBOB suite [8]. Originally proposed as part of the COCO benchmarking platform [7], this set of 24 continuous optimization problems has been one of the last decade’s most commonly used benchmarks for iterative optimization heuristics. For completeness, we perform our redundancy analysis on all 24 BBOB functions, even though a relatively large fraction consists of unimodal problems or problems with strong global structures where we don’t expect to see any redundant restart. We run each problem for dimensionality $d \in \{2, 3, 4, 5, 6, 7, 8, 9, 10, 20\}$ for 10 instances.

Figure 2 shows the RRF for each function, aggregated over all dimensions, runs, and instances for each of the three restart strategies. From the figure, as expected, most redundant restarts occur




Fig. 2: Boxplot showing the fraction of the total budget spent by restarts converging to a previously visited basin of attraction, the RRF (Equation 7). All 24 objective functions from the BBOB benchmark are aggregated over all dimensions and instances for each tested restart strategy.

in functions f_{21} – f_{24} , multimodal functions with weak global structure. We expect this is related to the weak structure since the chance we go to any given optima over one of the others should be related to the ‘direction’ of the function’s global structure. Since there is no global structure and several local optima of equal height, we would be equally likely to end up in any of them randomly.

The other functions showing some redundancy in Figure 2 are f_3 , f_4 , f_{15} and f_{19} , which also exhibit some multimodality. Notably, f_7 is an interesting outlier, which, due to its plateauing landscape, causes several restarts for the ‘RESTART’ strategy to be classified as redundant. Figure 3 provides a closer look at the functions with any redundant restarts, specifically showing the relation between dimensions for the ‘RESTART’ strategy. Here we can observe that for some functions, such as f_{19} and f_{22} , the RRF seems to increase with dimension; the opposite is true for functions f_3 and f_4 .

CEC 2013 While the overall potential to be gained from avoiding duplicated restarts is limited on the BBOB suite, we still see relatively large values in some multimodal problems. As such, we extend our setup to the CEC’13 suite [11], which is designed specifically for multimodal optimization. We multiply each function with -1 to enable minimization. To transform the functions to global optimization problems, we select one of the existing global optima to turn into the only global one by adding a norm of distance as follows: $f'(\mathbf{x}) = -f(\mathbf{x}) + \min(0.01, \|\mathbf{x} - \mathbf{x}^*\|^2)$. Given this “globalization” procedure, we create 10 instances of each of the 16 CEC problems (where each problem has a predefined dimensionality). The remainder of the setup is equivalent to the one used for the BBOB problems. The resulting redundancy factors are plotted in Figure 4, from which we observe that most problems defined for multimodal optimization show significant potential for saving redundant evaluations. We also observe that the distribution for IPOP is much wider and has a higher mean than the other methods. This happens because the population sizes are increased, leading to individual redundant runs containing a larger fraction of the total budget. This also occurs to a lesser extent in BIPOP since the larger population sizes are interleaved with low population size restarts.

4 Combating the Redundancy: Repelling CMA-ES

Based on the RS-CMSA[1], which uses the concept of repelling subpopulations in a niching context, we introduce a CMA-ES with repelling restarts (RR-CMA-ES). Building on the standard restart strategies, the algorithm’s idea is to define regions in the space where the CMA-ES cannot sample, as they were already visited during previous restarts. It maintains a set of tabu points $\mathbf{T}$, see Section 4.1, which describe a region. To determine the shape and size of this rejection region, we



Fig. 3: The average Relative Redundancy Factor over all instances for the CMA-ES using the ‘RESTART’ strategy for the BBOB functions with any redundant restarts. The grid shows the RRF (Eq. 7) per dimension and function individually.




Fig. 4: Fraction of budget which could be saved by avoiding convergence to redundant regions of the search space in subsequent restarts (as defined in Equation 6), for 16 functions from the CEC 2013 benchmark, aggregated over all instances for runs with the CMA-ES using different restart strategies.

use the ideas from the RS-CMSA [1] and use the current covariance matrix $\mathbf{C}$ and step size σ . The size of the region is unique for each tabu point and controlled by $\delta(\mathbf{T})$. Algorithm 2 gives a general overview of the method.

4.1 Tabu Points

The CMA-ES samples λ points at every generation g . In the RR-CMA-ES, each newly sampled point $\mathbf{x}$ is tested against the archive $\mathcal{T}$ of tabu points before being accepted for evaluation (line 9). A tabu point defines a location in the search space where the optimizer is not allowed to go. While initially imagined for combinatorial optimization, for the continuous spaces the CMA-ES deals with, it defines a hyper-ellipsoid centered around a point where the optimizer cannot sample.

Algorithm 2 RR-CMA-ES

Require: $\sigma_0 \in \mathbb{R}, \gamma \in \mathbb{R}, f : \mathbb{R}^d \rightarrow \mathbb{R}$

```

1:  $\mathcal{T} \leftarrow \emptyset$  ▷ Tabu point archive
2:  $r \leftarrow 0$ 
3: while not happy do
4:    $\sigma \leftarrow \sigma_0, \mathbf{C} \leftarrow \mathbf{I}, \mathbf{m} \sim \mathcal{U}(lb, ub)$ 
5:   while not any restart conditions do
6:      $\mathcal{X} \leftarrow \emptyset, n_{\text{rej}} \leftarrow 0$ 
7:     while  $|\mathcal{X}| < \lambda$  do
8:        $\mathbf{x} \leftarrow \mathbf{m} + \mathcal{N}(0, \sigma \mathbf{C})$ 
9:       if  $(\forall \mathbf{T} \in \mathcal{T} : \Delta_{\text{rej}}(\mathbf{x}, \mathbf{T}, \gamma, n_{\text{rej}}) = 0)$  then ▷ Check if  $\mathbf{x}$  can be accepted
10:         $\mathcal{X} \leftarrow \mathcal{X} \cup \{\mathbf{x}\}$ 
11:      else
12:         $n_{\text{rej}} \leftarrow n_{\text{rej}} + 1$  ▷ Increase rejection count for  $\mathbf{m}$ 
13:      end if
14:    end while
15:     $\mathcal{F} \leftarrow \text{evaluate}(\mathcal{X})$  ▷ Evaluate all offspring
16:     $\text{cmaUpdate}(\mathcal{X}, \mathcal{F}, \mathbf{m}, \sigma, \mathbf{C})$  ▷ Continue the regular CMA-ES procedure
17:  end while
18:   $r \leftarrow r + 1$ 
19:   $\mathcal{T} \leftarrow \mathcal{T} \cup \{\mathbf{m}\}$  ▷ Add  $\mathbf{m}$  to the tabu point archive
20: end while

```

Specifically, a tabu point $\mathbf{T}$ consists of a triplet $(\mathbf{x}^T, f(\mathbf{x}^T), n^T)$, where $\mathbf{x}^T$ is the location in the search space with $f(\mathbf{x}^T)$ its corresponding fitness and n^T being the number of times the CMA-ES has converged to $\mathbf{T}$ during previous restarts. During sampling, a tabu point $\mathbf{T}$ rejects a newly sampled point $\mathbf{x}$ according to the following boolean query ($\Delta_{\text{rej}} \in \{0, 1\}$):

$$\Delta_{\text{rej}}(\mathbf{x}, \mathbf{T}, \gamma, n_{\text{rej}}) = \left(\frac{d_m(\mathbf{x}, \mathbf{x}^T, \mathbf{C}^{-1})}{\sigma} < \gamma^{n_{\text{rej}}} \delta(\mathbf{T}) \right) \quad (8)$$

Here, d_m denotes the Mahanobis distance metric, scaled by the current covariance matrix $\mathbf{C}$ and step size σ , and $\delta(\mathbf{T})$ denotes the rejection radius around the tabu point $\mathbf{T}$. To avoid stagnation, a shrinkage factor $0 < \gamma < 1$ is applied to $\delta(\mathbf{T})$ with n_{rej} denoting the number of times a point has been rejected in the current generation. Then for every newly sampled point $\mathbf{x}$ it is accepted if and only if:

$$\forall \mathbf{T} \in \mathcal{T} : \Delta_{\text{rej}}(\mathbf{x}, \mathbf{T}, \gamma, n_{\text{rej}}) = 0 \quad (9)$$

4.2 Restarting

Upon every restart r , the archive $\mathcal{T}$ is updated with the converged center of mass $\mathbf{m}$. Using the Hill Valley HV routine, we check whether $\mathbf{m}$ has converged to a new basin of attraction. If the restart converged to a new basin of attraction, that is $\text{HV}(\mathbf{m}, \mathbf{x}^T, f) = 0, \forall \mathbf{T} \in \mathcal{T}$, the new tabu point $(\mathbf{m}, f(\mathbf{m}), 1)$ gets added to the archive. If the restart converged to a point that is already present in $\mathcal{T}$, n^T gets increased by 1, and if $f(\mathbf{m}) < f(\mathbf{x}^T)$, $(\mathbf{m}, f(\mathbf{m}), n^T)$ replaces $(\mathbf{x}^T, f(\mathbf{x}^T), n^T)$.

4.3 Search space coverage

We define a *coverage factor* c , which controls the ratio of the total volume of the search space $S = \prod_{i=1}^d (ub_i - lb_i)$ that is covered by the repelling regions of all tabu points. Intuitively, $\frac{1}{c}$ denotes the maximal proportion of S , which is unavailable for the CMA-ES to sample. This is divided amongst all tabu points, where points with a higher n^T cover a larger part of this volume. Then for each tabu point, the volume of the repelling region, normalized for σ_0 is:

$$V(\mathbf{T}) = n^T \frac{S}{c \sigma_0 R} \quad (10)$$




Fig. 5: Distribution of redundant function evaluation over the BBOB functions where redundant restarts were found in Section 3.2, i.e. $f_3, f_3, f_{15}, f_{19}, f_{21}, f_{22}, f_{23}$ and f_{24} , separated by problem dimensionality. The CMA-ES with the ‘RESTART’ strategy is compared to the RR-CMA-ES, with different coverage factors c .

where R denotes the total number of restarts. This is used to calculate the rejection radius:

$$\delta(\mathbf{T}) = V(\mathbf{T})^{\frac{1}{d}} \frac{\Gamma(\frac{d}{2} + 1)^{\frac{1}{d}}}{\sqrt{\pi}} \quad (11)$$

where d denotes the dimensionality and $\Gamma(\cdot)$ the gamma function.

5 Proof-of-Concept: Repelling CMA-ES in Action

To illustrate the workings of the RR-CMA-ES, we benchmark several versions with redundancy factors ranging from 2 to 1000 on both the BBOB and CEC’13 functions. We perform 50 independent runs on each of the 10 instances used in Section 3.2 and compare the results to the original CMA-ES. In this section, we use the versions with the default restart mechanism, but the IPOP and BIPOP results are available in our reproducibility repository [18]. To gauge whether the repelling strategy prevents duplicate restarts from occurring, we perform the same expected redundancy factor calculation from Equation 7 on the runs from the RR-CMA-ES, and visualize the results in Figure 5. From this figure, we can see that on average, the RR-CMA-ES with the lowest coverage factor (the largest repelling regions) quite effectively prevents different restarts from converging to the same basins.

We can perform the same redundancy-based comparisons for the CEC functions, where in Figure 6 we see how the potential changes when the repelling strategy is used. From this figure, we can see that the impact of the repelling regions is even larger than that of the BBOB problems, but with large variations between the different functions. However, the relation to the coverage factor seems consistent, with a lower factor leading to fewer redundant restarts.

Given the promising reduction in redundant evaluations, we look at the performance of the considered algorithm variants. In Figure 7, we show the empirical cumulative distribution plot³ on several BBOB problems where potential gains were identified. We use bounds 10^2 and 10^{-8} with log-scaling between them to be consistent with the common COCO setup [7]. This figure shows no difference between these methods in the early stage of the search (since no restarts have been triggered yet), but the repelling can increase or decrease overall performance depending on the function. An important aspect to note is that on functions with a clear global structure, such as F4, the repelling regions might adversely affect convergence, as shown in Figure 7a. Thus, a balance between preventing redundant restarts and avoiding steering the search away from regions near the global optimum has to be found.

³ This ECDF is based on the empirical attainment function, equivalent to infinite targets for the standard ECDF [12].




Fig. 6: Distribution of redundant function evaluation over the CEC'13 functions, only function with any redundant restarts for CMA-ES-RESTART are shown. The CMA-ES with the 'RESTART' strategy is compared to the RR-CMA-ES, with different coverage factors c .




Fig. 7: ECDF of the different versions of RR-CMA-ES compared to the original CMA-ES, for selected BBOB functions (in dimensionality 10) where potential improvement was observed from Figure 2.

6 Conclusions and Future Work

In this paper, we reflected that global optimization of multimodal problems is an area where current optimization algorithms can still be improved in light of an evident redundancy of restart mechanisms. When independent restarts converge to the same local optima, these essentially waste function evaluations, which could be prevented by more effectively guiding the search in subsequent restarts.

Our proposed repelling-based restart method for CMA-ES shows a slight performance improvement over the default restart mechanisms for specific cases and reduces redundant convergence for the BBOB functions. For the multimodal CEC functions specifically, the method shows a large reduction of redundant restarts compared to a standard restart strategy.

Further research into the precise ways of avoiding these redundancies is still required. In particular, the rejection criteria of the tabu regions could be modified to incorporate information about the fitness values found in this region to better handle global structure.

The tabu regions' shapes could also be formed based on the converged covariance matrix at the point or slightly before the restart is triggered to potentially better capture the precise basin shape. Furthermore, extended practices such as *regularization* could also be exercised to remedy numerical issues and to obtain a more accurate structure (see, e.g., [25]), relying on the theoretical relation of the ESS' covariance matrix and the landscape local Hessian [26].

In addition to more effectively preventing a restart from converging to a known basin of attraction, one could also utilize the information from prior restarts to more intelligently select the new location and initial parameterization of new restarts. In comparison to, e.g., MLSL, starting points could be selected based on which regions of the domain have already been sampled, which would

also cause the repelling mechanism to trigger less often during the initial phase of the restarted run.

In the context of global optimization, the explicit criteria used to trigger a restart mechanism also play a large role in the algorithm’s performance. Depending on the optimization goal, it might be worthwhile to restart more frequently and to exploit the best-attained basins of attraction using a local search method with a small fraction of the total budget.

Acknowledgements

References

1. Ahrari, A., Deb, K., Preuss, M.: Multimodal optimization by covariance matrix self-adaptation evolution strategy with repelling subpopulations. *Evolutionary Computation* **25**(3), 439–471 (2017). https://doi.org/10.1162/evco_a_00182
2. Antonov, K., Botari, T., Tukker, T., Bäck, T., van Stein, N., Kononova, A.V.: New solutions to Cooke triplet problem via analysis of attraction basins. In: Kress, B.C., Czarske, J.W. (eds.) *Digital Optical Technologies 2023*. vol. 12624, p. 126240T. International Society for Optics and Photonics, SPIE (2023). <https://doi.org/10.1117/12.2675836>, <https://doi.org/10.1117/12.2675836>
3. Auger, A., Hansen, N.: A restart cma evolution strategy with increasing population size. In: 2005 IEEE Congress on Evolutionary Computation. vol. 2, pp. 1769–1776 Vol. 2 (2005). <https://doi.org/10.1109/CEC.2005.1554902>
4. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, New York (2004)
5. Glover, F., Laguna, M.: *Tabu Search*. Springer US, Boston, MA (1998). https://doi.org/10.1007/978-1-4613-0303-9_33, https://doi.org/10.1007/978-1-4613-0303-9_33
6. Hansen, N.: Benchmarking a bi-population cma-es on the bbob-2009 function testbed. In: *Proceedings of the 11th annual conference companion on genetic and evolutionary computation conference: late breaking papers*. pp. 2389–2396 (2009)
7. Hansen, N., Auger, A., Ros, R., Mersmann, O., Tušar, T., Brockhoff, D.: COCO: a platform for comparing continuous optimizers in a black-box setting. *Optim. Methods Softw.* **36**(1), 114–144 (2021)
8. Hansen, N., Finck, S., Ros, R., Auger, A.: *Real-Parameter Black-Box Optimization Benchmarking 2009: Noiseless Functions Definitions*. Tech. Rep. RR-6829, INRIA (2009), <https://hal.inria.fr/inria-00362633/document>
9. Hansen, N., Kern, S.: Evaluating the CMA Evolution Strategy on Multimodal Test Functions. In: *Parallel Problem Solving from Nature - PPSN VIII*. Lecture Notes in Computer Science, vol. 3242, pp. 282–291. Springer-Verlag, Berlin Heidelberg (2004)
10. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation* **9**(2), 159–195 (2001)
11. Li, X., Engelbrecht, A., Epitropakis, M.G.: *Benchmark functions for cec’2013 special session and competition on niching methods for multimodal function optimization*. RMIT University, Evolutionary Computation and Machine Learning Group, Australia, Tech. Rep (2013)
12. López-Ibáñez, M., Vermetten, D., Dreó, J., Doerr, C.: Using the empirical attainment function for analyzing single-objective black-box optimization algorithms. *arXiv preprint arXiv:2404.02031* (2024)
13. Loshchilov, I., Schoenauer, M., Sebag, M.: Alternative restart strategies for cma-es. In: *International Conference on Parallel Problem Solving from Nature*. pp. 296–305. Springer (2012)
14. Lourenço, H.R., Martin, O.C., Stützle, T.: *Iterated local search: Framework and applications*. Handbook of metaheuristics pp. 129–168 (2019)
15. Maree, S.C., Alderliesten, T., Thierens, D., Bosman, P.A.N.: Real-valued evolutionary multi-modal optimization driven by hill-valley clustering. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. p. 857–864. GECCO ’18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3205455.3205477>, <https://doi.org/10.1145/3205455.3205477>
16. Minasny, B., McBratney, A.B.: A conditioned latin hypercube method for sampling in the presence of ancillary information. *Computers and Geosciences* **32**(9), 1378 – 1388 (2006). <https://doi.org/10.1016/j.cageo.2005.12.009>
17. Nishida, K., Akimoto, Y.: Benchmarking the psa-cma-es on the bbob noiseless testbed. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. pp. 1529–1536 (2018)
18. de Nobel, J., Shir, O., Vermetten, D., Kononova, A.V., Bäck, T.: *Reproducibility files and additional figures* (2024). <https://doi.org/10.5281/zenodo.10997200>
19. de Nobel, J., Vermetten, D., Wang, H., Doerr, C., Bäck, T.: Tuning as a means of assessing the benefits of new ideas in interplay with existing algorithmic modules. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. pp. 1375–1384 (2021)

20. de Nobel, J., Ye, F., Vermetten, D., Wang, H., Doerr, C., Bäck, T.: Iohexperimenter: Benchmarking platform for iterative optimization heuristics. *Evolutionary Computation* pp. 1–6 (2024)
21. Pál, L.: Benchmarking a hybrid multi level single linkagealgorithm on the bbob noiseless testbed. In: *Proceedings of the 15th annual conference companion on Genetic and evolutionary computation*. pp. 1145–1152 (2013)
22. Preuss, M.: *Multimodal Optimization by Means of Evolutionary Algorithms*. Natural Computing Series, Springer International Publishing (2015)
23. Rinnooy Kan, A., Timmer, G.T.: Stochastic global optimization methods part i: Clustering methods. *Mathematical programming* **39**, 27–56 (1987)
24. Shir, O.M.: Niching In Evolutionary Algorithms. In: Rozenberg, G., Baeck, T., Kok, J.N. (eds.) *Handbook of Natural Computing*. Springer Verlag (2012), <http://www.springer.com/computer/theoretical+computer+science/book/978-3-540-92909-3>
25. Shir, O.M., Roslund, J., Whitley, D., Rabitz, H.: Efficient retrieval of landscape hessian: Forced optimal covariance adaptive learning. *Physical Review E* **89**, 063306 (Jun 2014). <https://doi.org/10.1103/PhysRevE.89.063306>
26. Shir, O.M., Yehudayoff, A.: On the covariance-hessian relation in evolution strategies. *Theoretical Computer Science* **801**, 157–174 (2020). <https://doi.org/10.1016/j.tcs.2019.09.002>
27. Törn, A., Zilinskas, A.: *Global Optimization*, Lecture Notes in Computer Science, vol. 350. Springer, Berlin, Germany (1987)
28. Ursem, R.: Multinational evolutionary algorithms. In: *Proceedings of the 1999 Congress on Evolutionary Computation-CEC99* (Cat. No. 99TH8406). vol. 3, pp. 1633–1640 Vol. 3 (1999). <https://doi.org/10.1109/CEC.1999.785470>