



Universiteit
Leiden

The Netherlands

Algorithm design for mixed-integer black-box optimization problems with uncertainty

Thomaser, A.M.

Citation

Thomaser, A. M. (2024, October 22). *Algorithm design for mixed-integer black-box optimization problems with uncertainty*. Retrieved from <https://hdl.handle.net/1887/4104741>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4104741>

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

Chapter 6

Uncertainty Quantification

Optimization algorithms, such as CMA-ES, are designed to minimize an objective function by selecting the top- μ individuals from a population (Section 2.4). However, due to measurement errors, external disturbances or the inherent stochastic nature of the system under consideration, real-world optimization problems often exhibit noise or non-deterministic characteristics. This means that evaluating the objective with identical inputs can yield different results. Consequently, selecting the top- μ individuals based solely on a single evaluation of the objective function per individual can lead to inaccuracies. Section 2.7 provides a comprehensive overview of uncertainty quantification and allocation strategies to increase the certainty of this selection process.

This chapter introduces a novel Dynamic Allocation (DA) methodology that integrates uncertainty quantification with CMA-ES [50, 51] to mitigate the uncertainty in selecting the top- μ individuals in each generation (Section 6.2). The uncertainty quantification method utilized is based on the recently introduced concepts of Confidence Interval Sequences (CISs) (Section 2.7.3) and Rank Intervals (RIs) (Section 2.7.4). First, an artificial objective function characterized by inherent noise is developed to facilitate the design and testing of the proposed methodology (Section 6.1). The analysis evaluates the performance of this methodology and compares it to Static Allocation (SA) schemes within the CMA-ES algorithm. Finally, the methodology is applied to a real-world problem that exhibits noise, demonstrating its practical utility (Section 6.3).

6.1 Noisy Test Function

The investigation and analysis of various DA methods with minimal computational resources require test functions that are inexpensive to evaluate. To this end, a normalized ellipsoid function is defined, characterized by the dimensionality d and a distinct weight w_j for each dimension:

$$f(\mathbf{x}) = \frac{\sum_j^d w_j x_j^2}{\sum_j^d w_j}, \quad \text{with } w_j \in \{1, 2, \dots, d\}. \quad (6.1)$$

To introduce asymmetry into the landscape, each x_j is transformed according to the following equation:

$$x_j = \begin{cases} x_j, & \text{if } x_j \leq 0 \\ (1 + p_{\text{skew}})x_j, & \text{if } x_j > 0 \end{cases}, \quad (6.2)$$

where p_{skew} governs the degree of skewness in the objective landscape. The landscape reverts to that of the original ellipsoid function for $p_{\text{skew}} = 0$.

The black-box nature of many real-world optimization problems necessitates minimizing assumptions about the distribution of objective function values across repeated evaluations. Therefore, this thesis adopts a frequentist perspective. Each function evaluation corresponds to a specific scenario z_i , which is randomly sampled without replacement from the set of all possible scenarios Ω_Z . After evaluating all possible scenarios, the sample mean is identical to the true mean of the objective function (Equation 2.23).

For a given input \mathbf{x}_k , a set of $|\Omega_Z|$ values is sampled from a normal distribution with standard deviation σ_Z representing the scenarios z_i . After being sampled, this set of $|\Omega_Z|$ values remains fixed and is unique to each input \mathbf{x}_k . Furthermore, the mean across all $|\Omega_Z|$ values is subtracted from each individual value:

$$n(\mathbf{x}_k, z_i) := z_i - \frac{1}{|\Omega_Z|} \sum_{i=1}^{|\Omega_Z|} z_i, \quad \text{where } \Omega_Z := \{z_1, \dots, z_{|\Omega_Z|} \mid z_i \stackrel{s}{\sim} \sigma_Z \mathcal{N}(0, \mathbf{I})\}, \quad (6.3)$$

This ensures that the sample mean over all possible scenarios Ω_Z is equal to the true mean, which is always zero for each input \mathbf{x}_k :

$$\mathbb{E}_Z [n(\mathbf{x}_k, Z)] = \hat{\mathbb{E}}_Z [n(\mathbf{x}_k, Z)] = \frac{1}{|\Omega_Z|} \sum_{i=1}^{|\Omega_Z|} n(\mathbf{x}_k, Z) = 0. \quad (6.4)$$

Equation 6.3 defines the values used to model both additive and proportional noise within the test function. The parameter r_{noise} governs the balance between these two noise types. Consequently, for a given input \mathbf{x}_k and scenario z_i , the noisy test function $f(\mathbf{x}_k, z_i)$ can be expressed as follows:

$$f(\mathbf{x}_k, z_i) := f(\mathbf{x}_k) \cdot (1 + r_{\text{noise}} \cdot n(z_i)) + (1 - r_{\text{noise}}) \cdot n(z_i). \quad (6.5)$$

The test function is with no noise original unimodal. The multimodality results only from the noise introduced. Therefore, if an optimization algorithm fails to locate the global optimum, the failure cannot be attributed to local optima inherent to the original function. When all possible scenarios are sampled, the noisy test function converges to the original function.

The noisy test function employed in this chapter is defined with $p_{\text{skew}} = 0.5$, $r_{\text{noise}} = 0.7$ and $|\Omega_Z| = 50$. Figure 6.1 illustrates the one-dimensional noisy test function with $\sigma_Z = 0.5$ and different numbers of evaluated scenarios.

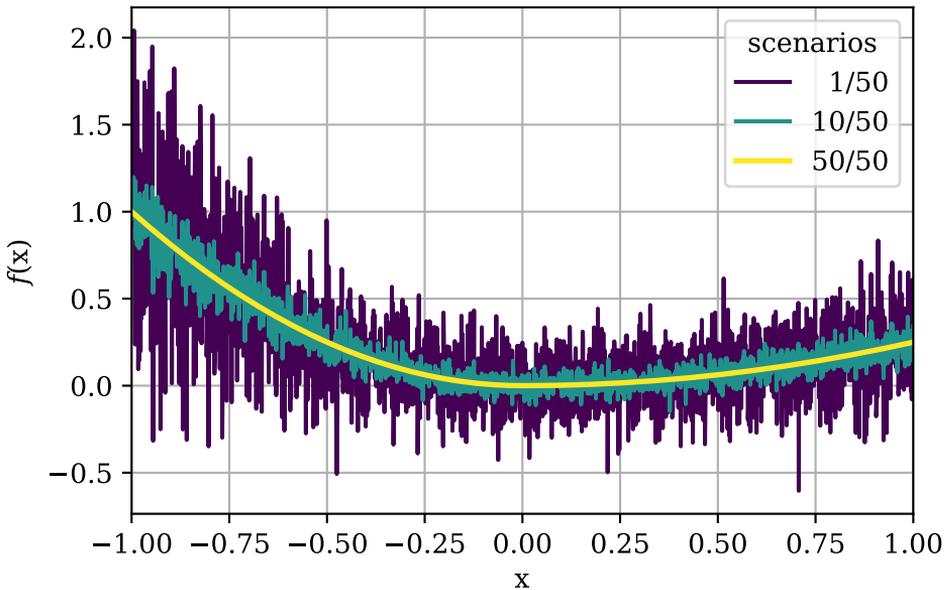


Figure 6.1: Mean values across 1, 10 and 50 evaluated scenarios on the noisy test function (Equation 6.5). Due to Equation 6.4, the noisy test function is equivalent to the original function as defined by Equation 6.1, when all 50 possible scenarios are evaluated. The degree of skewness p_{skew} is set to 0.5 (Equation 6.2). The standard deviation σ_Z of the normally distributed noise is set to 0.5 (Equation 6.3) and split to proportional and additional noise with the ratio $r_{\text{noise}} = 0.7$.

6.2 Methodology

The goal of the proposed methodology is to reliably identify the top- μ individuals within a population of λ individuals while minimizing the number of evaluations required. To achieve this, a DA methodology is proposed.

Algorithm 2 outlines the process of the proposed DA methodology. First, each individual \mathbf{x}_k is evaluated on the objective function $f(\mathbf{x}_k, z_i)$ for n_0 scenarios z_i , which are sampled WoR from the set of all possible scenarios Ω_Z . Scenarios are continuously allocated WoR to individuals until either the uncertainty in the selection of the top- μ individuals, quantified by the $UQiS_\mu$ value (Equation 2.41), falls below a predefined threshold u_{thr} or the maximum number of evaluations N is reached. In each iteration, the CISs CIS_k and RIs RI_k of each individual \mathbf{x}_k are computed based on the until then determined K_k objective function values $f(\mathbf{x}_k, z_i)$ of each individual \mathbf{x}_k . In order to apply the CISs introduced in Section 2.7.3, the objective function values are min-max scaled to $[0, 1]$ with the provided lower bounds lb and upper bounds ub . The $UQiS_\mu$ value is computed from the RIs. If the $UQiS_\mu$ value is not below the threshold u_{thr} , a DA policy $\pi_{\text{top-}\mu}$ allocates n_{allocate} more scenarios to the individuals. Each individual \mathbf{x}_k is then evaluated with the scenarios z_i allocated to that individual, and K_k is adjusted accordingly. Finally, the sample mean across the K_k evaluated scenarios $\hat{E}_Z [f(\mathbf{x}_k, Z)]$ is returned for each individual.

Algorithm 2 Proposed DA methodology to minimize UQiS.

Require: $\lambda, \{\mathbf{x}_k\}_{k=1}^\lambda, \mu, |\Omega_Z|, n_{\text{allocate}}, u_{\text{thr}}, \alpha, lb, ub$

$n_0 \leftarrow 1$

$UQiS_\mu \leftarrow \text{inf}$

Evaluate all individuals $\{\mathbf{x}_k\}_{k=1}^\lambda$ for n_0 scenarios z_k ▷ sampled WoR from Ω_Z

$n \leftarrow \lambda \cdot n_0$

$N \leftarrow \lambda \cdot |\Omega_Z|$

while $n \leq N - n_{\text{allocate}}$ and $UQiS_\mu > u_{\text{thr}}$ **do**

Compute CIS_k ▷ according to Equation 2.29 and 2.33

Compute RI_k and $UQiS_\mu$ ▷ according to Equation 2.39 and 2.41

if $UQiS_\mu > u_{\text{thr}}$ **then**

Allocate n_{allocate} scenarios ▷ according to DA policy $\pi_{\text{top-}\mu}$ (Sections 6.2.1)

Evaluate individuals with allocated scenarios

Adjust K_k ▷ according to the number of allocated scenarios to \mathbf{x}_k

$n \leftarrow \sum_{k=1}^\lambda K_k$ ▷ total number of allocated scenarios

end if

end while

return $\hat{E}_Z [f(\mathbf{x}_k, Z)] = \frac{1}{K_k} \sum_{i=1}^{K_k} f(\mathbf{x}_k, z_i)$

6.2.1 Dynamic Allocation Policy

A DA policy proposed by Ellmaier [32] for the top- μ selection task is utilized. This DA policy, denoted as $\pi_{\text{top-}\mu}$, allocates additional scenarios to specific individuals based on the computed RIs RI_k of each individual \mathbf{x}_k . For the top- μ selection task, only those individuals need to be considered that, according to their RI, are neither clearly among the top- μ nor the bottom- $(\lambda-\mu)$ subset.

$$\Lambda_\mu := \left\{ k \in \{1, \dots, \lambda\} \mid RI_k \in \{1, \dots, \mu\} \text{ and } RI_k \in \{\mu + 1, \dots, \lambda\} \right\}, \quad (6.6)$$

consist of the indices of these individuals. The excess of the RI RI_k of individual \mathbf{x}_k , denoted as ξ_k , is defined as follows:

$$\xi_k := \begin{cases} \max(RI_k) - \mu, & \text{if } \hat{\text{rank}}_{\text{freq}}(\mathbf{x}_k) \leq \mu \\ \mu + 1 - \min(RI_k), & \text{else} \end{cases}. \quad (6.7)$$

Figure 6.2 exemplarily shows the excesses of the RI ξ_k for the selection of the top-3 individuals from a population of six individuals. According to their RIs all individuals can be among the top-3 individuals. Thus, the $UQiS_3$ equals three, and selecting the top-3 individuals is ambiguous. If all the excesses of the RIs ξ_k are zero, the $UQiS_\mu$ value is zero, and the selection of the top- μ individuals is unambiguous.

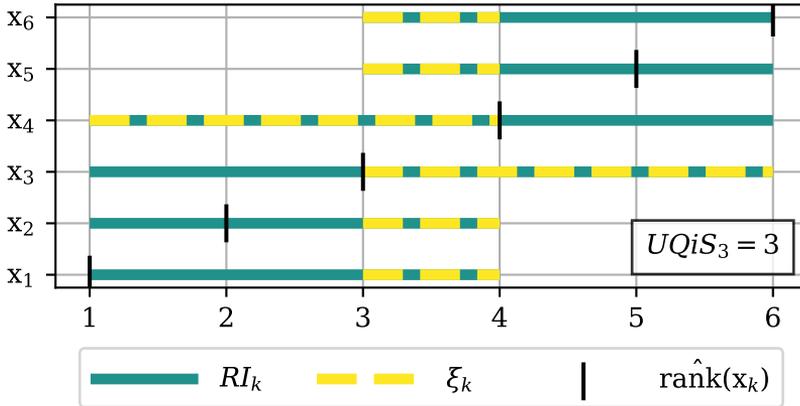


Figure 6.2: Given the RIs RI_k for each individual in $\{\mathbf{x}_1, \dots, \mathbf{x}_6\}$ the excess of the RIs ξ_k computed according to Equation 6.7 is illustrated for the selection of the top-3 individuals. According to the RIs all individuals can be among the top-3 individuals, and therefore the $UQiS_3$ equals three.

Algorithm 3 outlines the process of the proposed DA policy $\pi_{\text{top-}\mu}$. In each iteration, the policy can allocate n_{allocate} additional scenarios among the $|\Lambda_\mu|$ individuals that play a pivotal role in differentiating the top- μ individuals from the rest of the population according to Equation 6.6. The policy prioritizes individuals based on the excesses of their RIs ξ_k . This assumes that the magnitude of this excess is indicative of the uncertainty in the selection process. Then, the policy allocates additional objective values to these individuals to thereby narrow the width of their CIs and, consequently, their RIs.

To distribute more than one scenario per iteration in proportion to the magnitude of the excesses of the RIs, the policy employs the Imperiali method. The Imperiali method [39] originally designed to allocate parliamentary seats to political parties based on their proportion of votes, is repurposed here to allocate scenarios. With the Imperiali method, an integer number of scenarios $K_{\text{allocated},i}$ is allocated to each of the $|\Lambda_\mu|$ individuals not clearly assigned to either the top- μ or bottom- $(\lambda-\mu)$ subset. Thereby, the sum of the allocated scenarios $K_{\text{allocated},i}$ to the individuals is equal to the total number of allocated scenarios:

$$n_{\text{allocate}} = \sum_k^\lambda K_{\text{allocated},k}. \quad (6.8)$$

The allocation is based on the proportional excess of the RI RI_k of individual \mathbf{x}_k in percent of the total excess of the RIs of the $|\Lambda_\mu|$ individuals, denoted as $\xi_{\text{percent},k}$:

$$\xi_{\text{percent},k} := \frac{100 \cdot \xi_k}{\sum_{j \in \Lambda_\mu} \xi_j}. \quad (6.9)$$

Algorithm 3 DA policy $\pi_{\text{top-}\mu}$ to allocate n_{allocate} additional scenarios [32].

Require: $\lambda, \{RI_k\}_{k=1}^\lambda, \mu, n_{\text{allocate}}$

Determine Λ_μ ▷ according to Equation 6.6

Compute ξ_k ▷ according to Equation 6.7

Compute $\xi_{\text{percent},k}$ ▷ according to Equation 6.9

$\{K_{\text{allocated},k}\}_{k=1}^\lambda \leftarrow \text{Imperiali}(\text{totalscenario} = n_{\text{allocate}}, \text{percentages} = \xi_{\text{percent},k})$

return $\{K_{\text{allocated},k}\}_{k=1}^\lambda$

If the excess of the RIs of two individuals is equal, the width of their CIs is considered as a secondary attribute for the allocation. Apart from the proportional excess of the RIs, other metrics based on, for example, the excess of the CIs can be used as an alternative.

6.2.2 Top- μ Selection and Ranking

The mechanism of the DA methodology in selecting the top- μ individuals from a population, as outlined in Algorithm 2, is exemplarily demonstrated by utilizing the noisy test function defined in Section 6.1 with a dimensionality of two and a standard deviation σ_Z of 0.5. A population is generated by randomly sampling six individuals $\{\mathbf{x}_1, \dots, \mathbf{x}_6\}$ from a uniform distribution within the input space $[-1, 1]^2$. In each iteration, six additional scenarios can be allocated to the individuals, with a total limit of 50 scenarios per individual. The allocation is stopped when the UQiS of the top- μ individuals $UQiS_\mu$ is zero. The significance level α for the computation of the CISs is set to 0.3. Furthermore, for the uncertainty quantification, the lower bound lb and upper bound ub are set to the true minimum and maximum possible objective function values, respectively, by evaluating all possible scenarios for each individual in advance. The provision of broader lower and upper bounds does not affect the core principles of the qualitative findings that emerge from the results presented with the true lower and upper bounds. Table 6.1 summarizes the chosen hyperparameters of Algorithm 2 for the top- μ selection task and the utilized noisy test function.

Table 6.1: Chosen hyperparameters of the utilized noisy test function (Section 6.1) and of the DA (Algorithm 2) for the top- μ selection task.

Component	Parameter Description	Parameter Value
	Dimensionality	2
Noisy	Standard deviation σ_Z	0.5
Test	Degree of skewness p_{skew}	0.5
Function	Ratio proportional noise r_{noise}	0.7
	Possible scenarios Ω_Z	$\{z_1, \dots, z_{50}\}$
	Population X	$\{\mathbf{x}_1, \dots, \mathbf{x}_6\}$
	Significance level α	0.3
	UQiS of top- μ individuals	μ
DA	UQiS threshold u_{thr}	0
	Allocations per iteration n_{allocate}	6
	Lower bound lb	$\min_{\mathbf{x}_k \in X, z_i \in \Omega_Z} (f(\mathbf{x}_k, z_i))$
	Upper bound ub	$\max_{\mathbf{x}_k \in X, z_i \in \Omega_Z} (f(\mathbf{x}_k, z_i))$

First, the task of selecting the top-1 individual out of the population of the six individuals $\{\mathbf{x}_1, \dots, \mathbf{x}_6\}$ is examined. Figure 6.3 illustrates the state after the final iteration of the DA by Algorithm 2, showcasing the determined objective function values $f(\mathbf{x}_k, z_i)$, along with the computed CIs CI_k , RIs RI_k and the resultant UQiS of the top-1 individual $UQiS_1$.

A total of 114 scenarios are allocated across all individuals. The top-1 individual \mathbf{x}_1 is evaluated on all 50 scenarios, which narrows the CI for this individual to a width of zero. The number of scenarios allocated to an individual K_k decreases along with the rank of the individual. This allocation pattern arises because the employed DA policy $\pi_{\text{top-}\mu}$ (Algorithm 3) stops allocating additional scenarios to an individual once this individual can be definitively assigned to the top-1 or bottom-5 subset based on the RIs RI_k . This trend is especially noticeable for the individuals in the lower ranks, which can be assigned to the bottom-5 subset with only a few scenarios, resulting in comparatively large CIs. However, in the final iteration, the CIs of the individuals in the bottom-5 subset are tight enough to avoid overlapping with the CI of the top-1 individual. Thus, the UQiS of the top-1 individual $UQiS_1$ is zero, and the selection of the top-1 individual is clear.

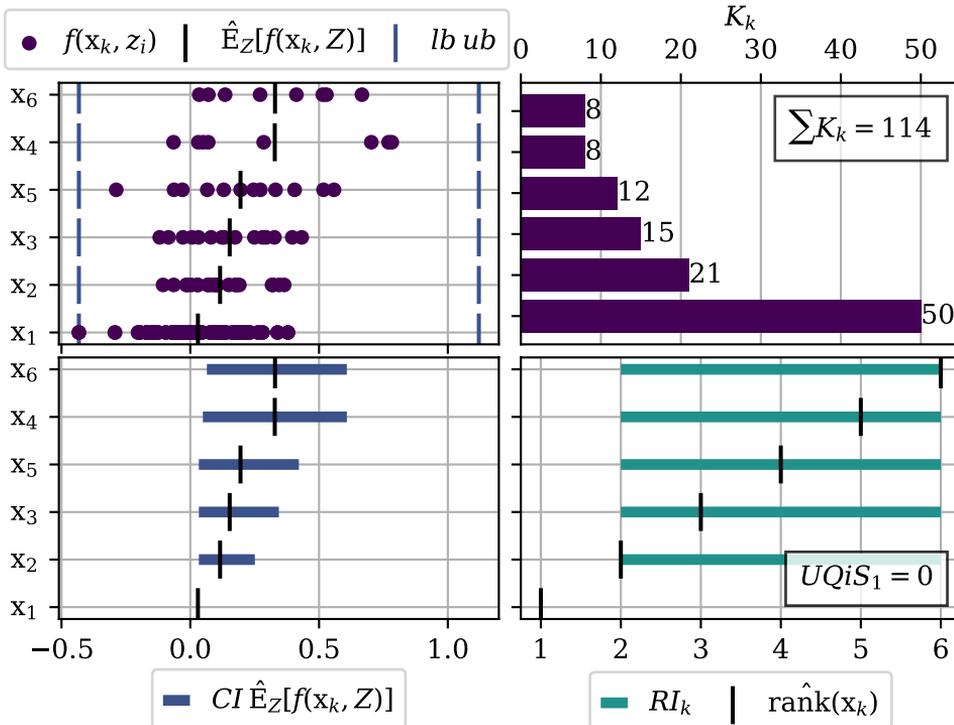


Figure 6.3: State after the final iteration of the DA by Algorithm 2 with hyperparameters of Table 6.1 for the selection of the top-1 individual from the population $X = \{\mathbf{x}_1, \dots, \mathbf{x}_6\}$. *Top Left:* Determined objective function values $f(\mathbf{x}_k, z_i)$, sample means and given bounds. *Top Right:* Number of allocated scenarios K_k sampled WoR from the 50 possible scenarios. *Bottom Left:* Computed CIs from the determined objective function values and sample means. *Bottom Right:* Resulting RIs, sample ranks and UQiS of the top-1 individual $UQiS_1$.

Subsequently, the task of selecting the top-3 individuals from the population of the six individuals $\{\mathbf{x}_1, \dots, \mathbf{x}_6\}$ is considered. Figure 6.4 illustrates the final state. A total of 180 scenarios are allocated across the population, requiring 66 additional evaluations compared to selecting only the top-1 individual. Moreover, a significant shift in the allocation pattern is observed. The number of scenarios allocated to individuals in the middle ranks is increased. This is necessary to narrow the CIs for an unambiguous assignment of these individuals to either the top-3 or bottom-3 subsets. In comparison, the individuals at the extreme ranks need fewer allocations for an unambiguous assignment. In the final iteration, the RIs of individuals in the top-3 subset do not overlap with those in the bottom-3 subset. Therefore, the selection of the top-3 individuals is unambiguous, and the UQiS of the top-3 individuals $UQiS_3$ is zero.

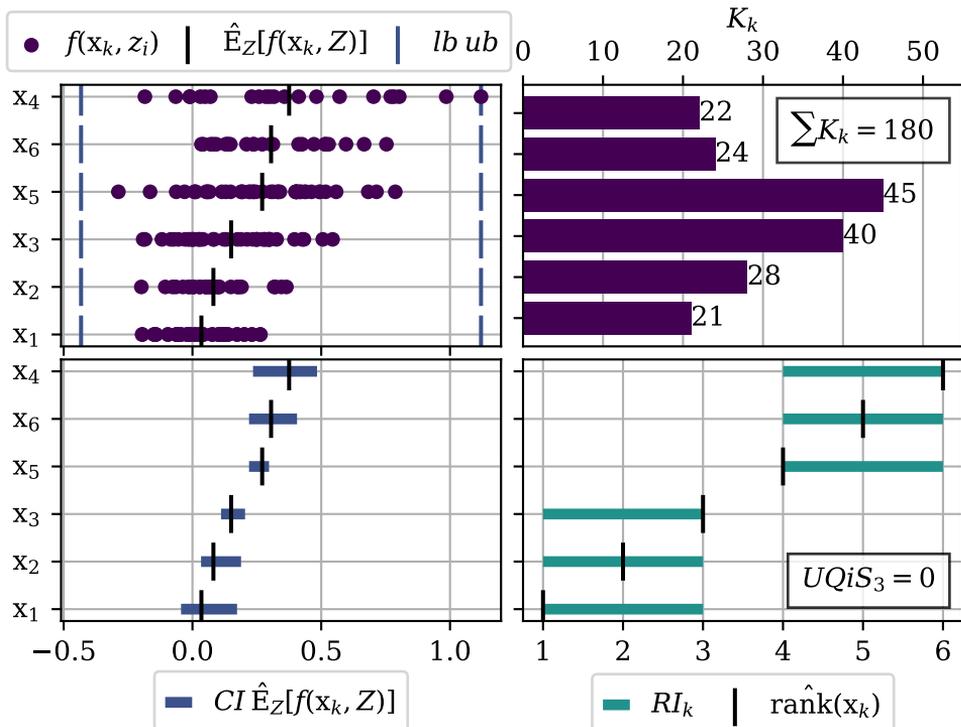


Figure 6.4: State after the final iteration of the DA by Algorithm 2 with hyperparameters of Table 6.1 for the selection of the top-3 individual from the population $X = \{\mathbf{x}_1, \dots, \mathbf{x}_6\}$. *Top Left:* Determined objective function values $f(\mathbf{x}_k, z_i)$, sample means and given bounds. *Top Right:* Number of allocated scenarios K_k sampled WoR from the 50 possible scenarios. *Bottom Left:* Computed CIs from the determined objective function values and sample means. *Bottom Right:* Resulting RIs, sample ranks and UQiS of the top-3 individual $UQiS_3$.

In each generation of CMA-ES, the top- μ individuals from a population are selected. Additionally, in the subsequent recombination step, the individuals are weighted based on their rank (Equation 2.5). Therefore, besides selecting the top- μ individuals, the ranking within these top- μ individuals is required. To achieve this, the DA by Algorithm 2 is conducted several times, sequentially increasing the selection of the top- μ individuals from 1 to 3. If $UQiS_1$ is zero, then $UQiS_2$ is used to allocate additional scenarios until it is zero. Once $UQiS_3$ is also zero, a total of 204 scenarios are allocated (Figure 6.5). As the top three individuals have similar objective function values, all three are evaluated on all 50 possible scenarios. The remaining individuals are assigned to the bottom three sets with fewer allocated scenarios.

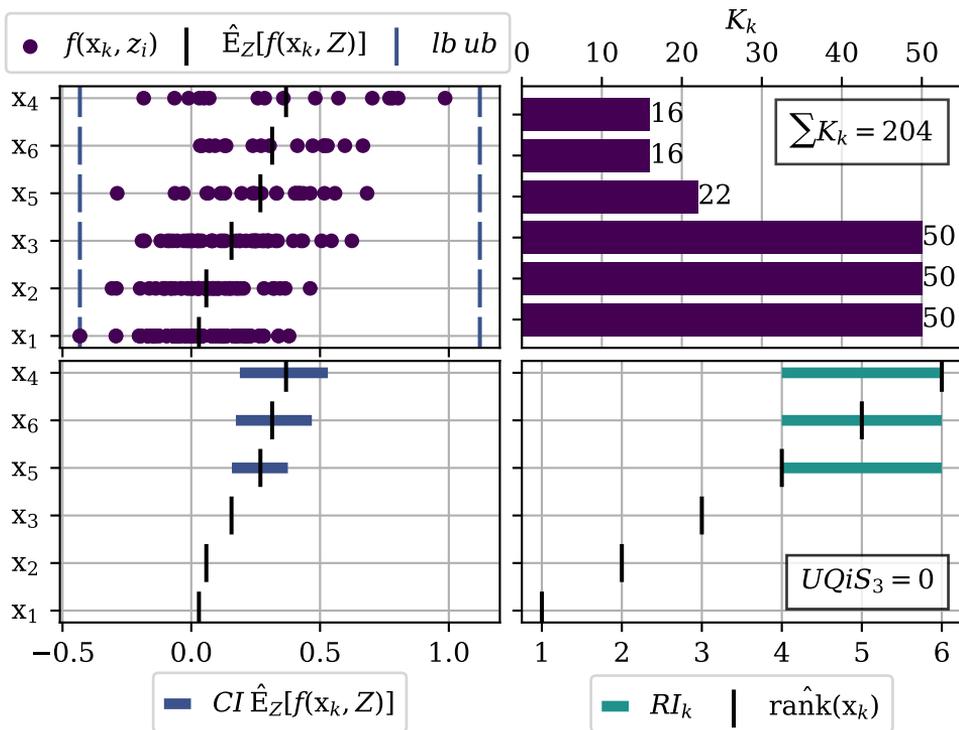


Figure 6.5: State after the final iteration of the DA by Algorithm 2 with hyperparameters of Table 6.1 for the ranking of the top-3 individual from the population $X = \{x_1, \dots, x_6\}$. *Top Left:* Determined objective function values $f(x_k, z_i)$, sample means and given bounds. *Top Right:* Number of allocated scenarios K_k sampled WoR from the 50 possible scenarios. *Bottom Left:* Computed CIs from the determined objective function values and sample means. *Bottom Right:* Resulting RIs, sample ranks and UQiS of the top-3 individual $UQiS_3$.

6.2.3 CMA-ES and Bound Adaption for UQiS

The DA methodology (Algorithm 2) can be combined with CMA-ES to identify the top- μ individuals within each generation g of CMA-ES. However, since CMA-ES minimizes the objective function, the values of the objective function are likely to decrease as the generations progress. Moreover, the amount of noise may change during the optimization process. Therefore, using static lower and upper bounds for the UQiS throughout the entire optimization procedure is suboptimal.

To address this issue, a bounds adaption mechanism is introduced. This mechanism adjusts the bounds for the UQiS dynamically after each generation based on the minimum and maximum objective function values from the last n_{gen} generations, denoted as $y_{\min, n_{\text{gen}}}$ and $y_{\max, n_{\text{gen}}}$. To account for potential future increases or decreases in the bounds, the range between the determined minimum and maximum values, scaled by a factor b_{factor} , is added to the current minimum and maximum values, respectively:

$$\begin{aligned} ub^{(g+1)} &= y_{\max, n_{\text{gen}}} + b_{\text{factor}} \cdot (y_{\max, n_{\text{gen}}} - y_{\min, n_{\text{gen}}}), \\ lb^{(g+1)} &= y_{\min, n_{\text{gen}}} + b_{\text{factor}} \cdot (y_{\max, n_{\text{gen}}} - y_{\min, n_{\text{gen}}}), \end{aligned} \quad (6.10)$$

Figure 6.6 illustrates the bound adaption with $b_{\text{factor}} = 0.25$ and $n_{\text{gen}} = 3$ during a CMA-ES run on the noisy test function with $\sigma_Z = 0.5$ and $r_{\text{noise}} = 0.9$.

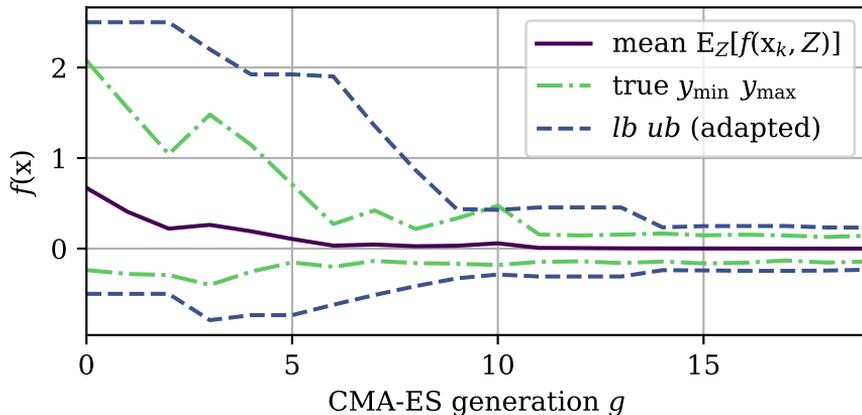


Figure 6.6: Adaption of the lower bound lb and the upper bound ub according to Equation 6.10 with $b_{\text{factor}} = 0.25$ and $n_{\text{gen}} = 3$ in each generation g of a CMA-ES run on the noisy test function (Section 6.1) with $\sigma_Z = 0.5$ and $r_{\text{noise}} = 0.9$. In addition, for the DA methodology, unknown true minimum y_{\min} and maximum y_{\max} objective function values are provided along with the true mean objective function value across the CMA-ES population $X = \{\mathbf{x}_1, \dots, \mathbf{x}_\mu\}$.

The significance level α is set to 0.99, and the UQiS threshold u_{thr} is set to 1. This configuration results in fewer scenarios being allocated for selecting the top- μ individuals, which resembles a more complex case for bound adaptation due to the less information available from past generations. Starting with an initial lower bound of -0.5 and an upper bound of 2.5 , the first adaptation is conducted after three generations. During the 20 generations of CMA-ES, the mean objective function value and thus also the proportional noise component decreases, which results in a continuous adjustment of the bounds.

For each generation g , the bounds are considered correctly adapted if the lower bound $lb^{(g)}$ is less than the true minimum objective function value $y_{\min}^{(g)}$ and the upper bound $ub^{(g)}$ is greater than the true maximum objective function value $y_{\max}^{(g)}$. Over the course of 20 generations, the bounds are correct in 95% of the cases (Figure 6.6). Only in generation 10 is the upper bound slightly too low compared to the true maximum objective function value. This violates the assumptions for the UQiS, which requires that all objective function values fall within the given lower and upper bounds. Increasing the factor b_{factor} leads to looser bounds and a higher percentage of correct bounds. However, looser bounds also result in larger CIs, necessitating more function evaluations to unambiguously select the top- μ individuals. To investigate this trade-off, an experiment is conducted with varying the values of the factor b_{factor} . Table 6.2 summarizes the chosen hyperparameters, resulting in 30 configurations. For each configuration, 100 CMA-ES runs are conducted for 20 generations on the noisy test function, defined in Section 6.1.

Table 6.2: Chosen hyperparameters of the utilized noisy test function (Section 6.1) and of the DA (Algorithm 2) for the analysis of the bound adaption mechanism.

Component	Parameter Description	Parameter Values
	Dimensionality	2
Noisy	Standard deviation σ_Z	{0.1, 0.25, 0.5}
Test	Degree of skewness p_{skew}	0.5
Function	Ratio proportional noise r_{noise}	0.9
	Possible scenarios Ω_Z	{ z_1, \dots, z_{50} }
	Population X	{ $\mathbf{x}_1, \dots, \mathbf{x}_6$ }
	Significance level α	0.99
	UQiS of top- μ individuals	{1, 3}
DA	UQiS threshold u_{thr}	1
	Allocations per iteration n_{allocate}	6
	Initial lower bound lb	-0.5
	Initial upper bound ub	2.5
	Bound adaption factor b_{factor}	{0.0, 0.1, 0.25, 0.5, 1.0}

Figure 6.7 illustrates the relationship between the ratio of correct bounds and the ratio of evaluated scenarios, averaged across 200 CMA-ES runs, for various combinations of standard deviation σ_Z and factor b_{factor} values. As expected, a higher factor, which corresponds to looser bounds, results in a higher percentage of correct bounds. With a factor of zero, slightly more than half of the bounds are correct. Increasing the factor to 0.1 leads to approximately 90% of the bounds being correct. However, as the factor increases, the increase in accuracy decreases while the number of scenarios evaluated continues to increase. Therefore, for the subsequent experiments, the factor b_{factor} is set to 0.1 as a compromise between accuracy and computational efficiency.

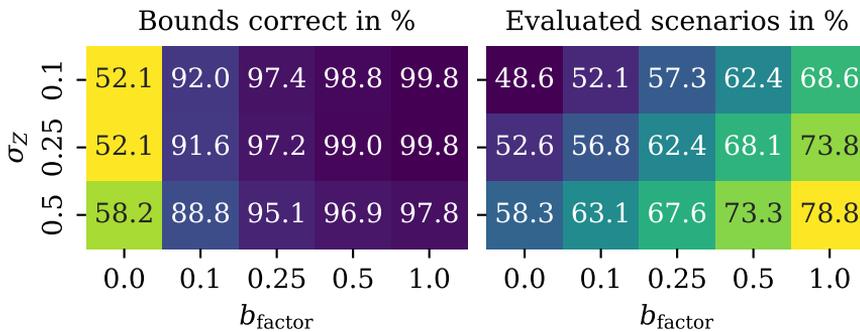


Figure 6.7: The ratio of correct bounds in percent (*left*) and the ratio of evaluated scenarios in percent (*right*) on average across 200 CMA-ES runs for different combinations of the standard deviations σ_Z and factors b_{factor} for the bound adaption according to Equation 6.10.

6.2.4 CMA-ES Convergence

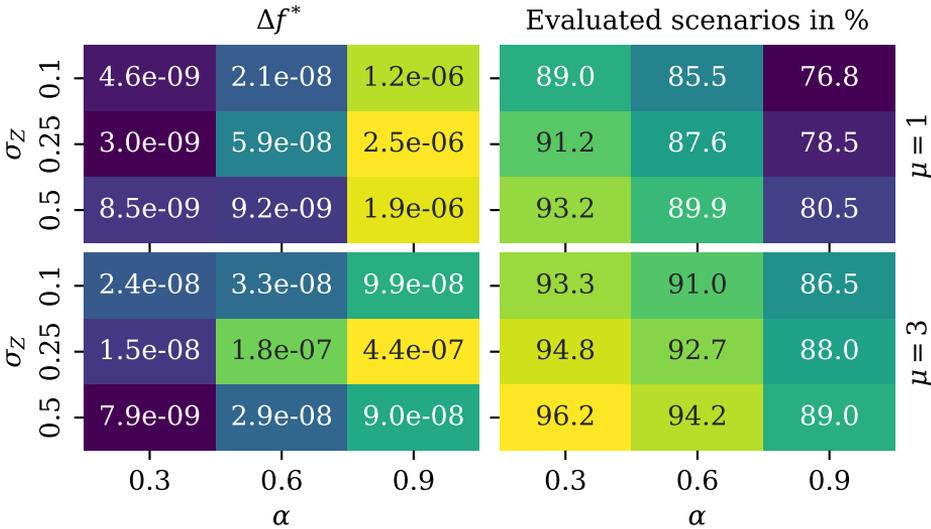
The DA methodology (Algorithm 3) can be employed to determine the ranking of the top- μ individuals within a CMA-ES population (Section 6.2.2). The magnitude of uncertainty in the ranking can be primarily controlled by the user through two parameters: the significance level α and the threshold for the UQiS u_{thr} . With higher levels of uncertainty, the probability of an incorrect ranking increases.

In this section, the effect of increased uncertainty in the ranking on the convergence of CMA-ES is investigated for $\mu = 1$ and $\mu = 3$. The significance level α is varied first, followed by the threshold u_{thr} . Finally, the DA methodology is compared to SA. Table 6.3 summarizes the chosen hyperparameters. For each configuration, 100 CMA-ES runs are conducted over 50 generations on the noisy test function, defined in Section 6.1.

Table 6.3: Chosen hyperparameters of the utilized noisy test function (Section 6.1) and of the DA (Algorithm 2) for the analysis of the convergence of CMA-ES.

Component	Parameter Description	Parameter Values
Noisy Test Function	Dimensionality	2
	Standard deviation σ_Z	{0.1, 0.25, 0.5}
	Degree of skewness p_{skew}	0.5
	Ratio proportional noise r_{noise}	0.7
	Possible scenarios Ω_Z	{ z_1, \dots, z_{50} }
DA	Population X	{ $\mathbf{x}_1, \dots, \mathbf{x}_6$ }
	Allocations per iteration n_{allocate}	6
	Initial lower and upper bound	-0.5, 2.5

Three values for the significance level $\alpha \in \{0.3, 0.6, 0.9\}$ are considered (Figure 6.8). Higher values of α correspond to increased uncertainty, which leads to larger distances to the optimal objective function value Δf^* and, therefore, slower convergence of CMA-ES. This trend is observed across all three standard deviations. The DA methodology adapts to the different noise levels for a given α by adjusting the number of evaluations. For higher values of α , fewer evaluations are required to achieve an unambiguous ranking due to the resulting tighter CIs. Determining an unambiguous ranking for $\mu = 3$ requires more evaluation than for $\mu = 1$.

**Figure 6.8:** Distance to the optimal objective function value Δf^* after 50 generations of CMA-ES and the ratio of evaluated scenarios on average across 100 runs for $\mu = 1$ (first row) and $\mu = 3$ (second row). Different configurations of the significance level $\alpha \in \{0.3, 0.6, 0.9\}$ (Equation 2.26) and the standard deviation σ_Z of the noisy test function are considered.

To examine the convergence of CMA-ES across generations, the configuration with $\mu = 1$ and $\sigma_Z = 0.5$ is analyzed in detail. Figure 6.9 presents the distance to the optimal objective function value Δf^* , the ratio of evaluated scenarios and the ratio of correct ranking (RCR) of the top- μ individuals, denoted as RCR_μ , across 100 CMA-ES runs for each generation g . For $\alpha = 0.9$, the decrease in the distance to the optimal objective function value slows down after 10 generations compared to $\alpha = 0.3$ and $\alpha = 0.6$. Simultaneously, the RCR falls below 0.9. With $\alpha = 0.6$, the RCR remains above 0.9, and the convergence of CMA-ES is not noticeably impacted, even though fewer scenarios are evaluated compared to the case with $\alpha = 0.3$. CMA-ES can internally compensate for the incorrect rankings to some extent.

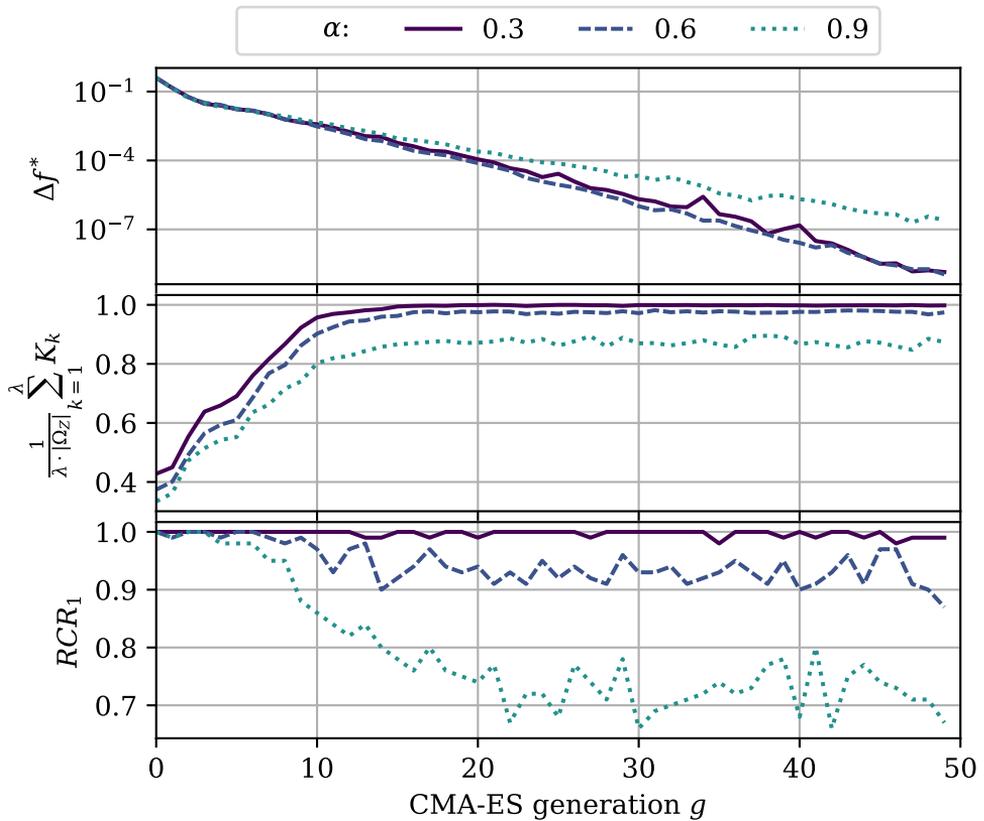


Figure 6.9: Mean distance to the optimal objective function value Δf^* , the ratio of evaluated scenarios and RCR of the top-1 individual RCR_1 for each CMA-ES generation g across 100 runs. Different values of the significance level $\alpha \in \{0.3, 0.6, 0.9\}$ (Equation 2.26) for the DA (Algorithm 2) are considered. The standard deviation of the noisy test function is set to 0.5.

However, there is no significant difference in convergence among the different considered significance levels α until the first five to ten generations. The RCR is also 1.0 for all three significance levels. Therefore, a significance level of 0.3 is too conservative in the beginning and leads to unnecessary evaluations. Conversely, a significance level of 0.9 is too high after the beginning, as not enough scenarios are allocated, and the RCR falls to 0.7.

Subsequently, the UQiS threshold u_{thr} is varied between zero and three, with the significance level α set to 0.3. The DA methodology allocates scenarios until the UQiS does not exceed the set threshold. Therefore, for thresholds greater than zero, the DA methodology allocates fewer scenarios, and the resulting ranking includes $\mu + u_{\text{thr}}$ individuals within the top- μ set. Thus, higher thresholds result in a higher uncertainty in the ranking of the top- μ individuals, leading to poorer convergence across all three standard deviations (Figure 6.10). For $\mu = 1$ and $\mu = 3$, the distance to the optimal objective function value after 50 generations of CMA-ES is similar for the different considered thresholds. However, for $\mu = 3$, more scenarios are required.

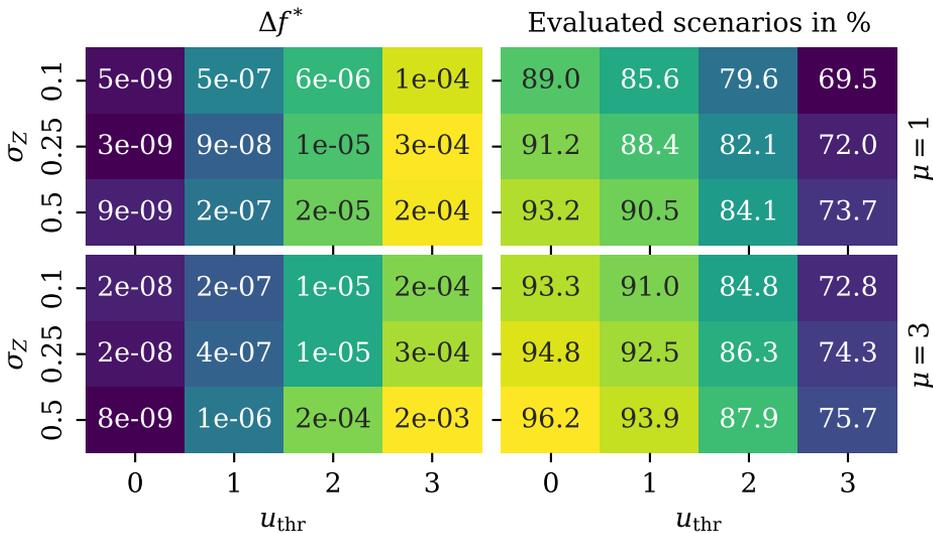


Figure 6.10: Distance to the optimal objective function value Δf^* after 50 generations of CMA-ES and the ratio of evaluated scenarios on average across 100 runs for $\mu = 1$ (first row) and $\mu = 3$ (second row). Different configurations of the UQiS threshold $u_{\text{thr}} \in \{0, 1, 2, 3\}$ and the standard deviation $\sigma_Z \in \{0.1, 0.25, 0.5\}$ of the noisy test function are considered.

For the configuration with $\mu = 1$ and $\sigma_Z = 0.5$, Figure 6.11 presents the distance to the optimal objective function value Δf^* , the ratio of evaluated scenarios and the ratio of correct ranking (RCR) of the top- μ individuals across 50 generations of CMA-ES for the different considered thresholds $u_{\text{thr}} \in \{0, 1, 2, 3\}$. With higher thresholds, the convergence slows down with further generations. After a continuous increase in the number of allocated scenarios in the first ten generations, the allocation stagnates. With a threshold of one, almost all possible scenarios are allocated. However, the RCR is only around 0.8 due to the two possible individuals within the top-1 individual according to the UQiS. Thus, the threshold needs to be zero to receive a correct ranking according to the UQiS.

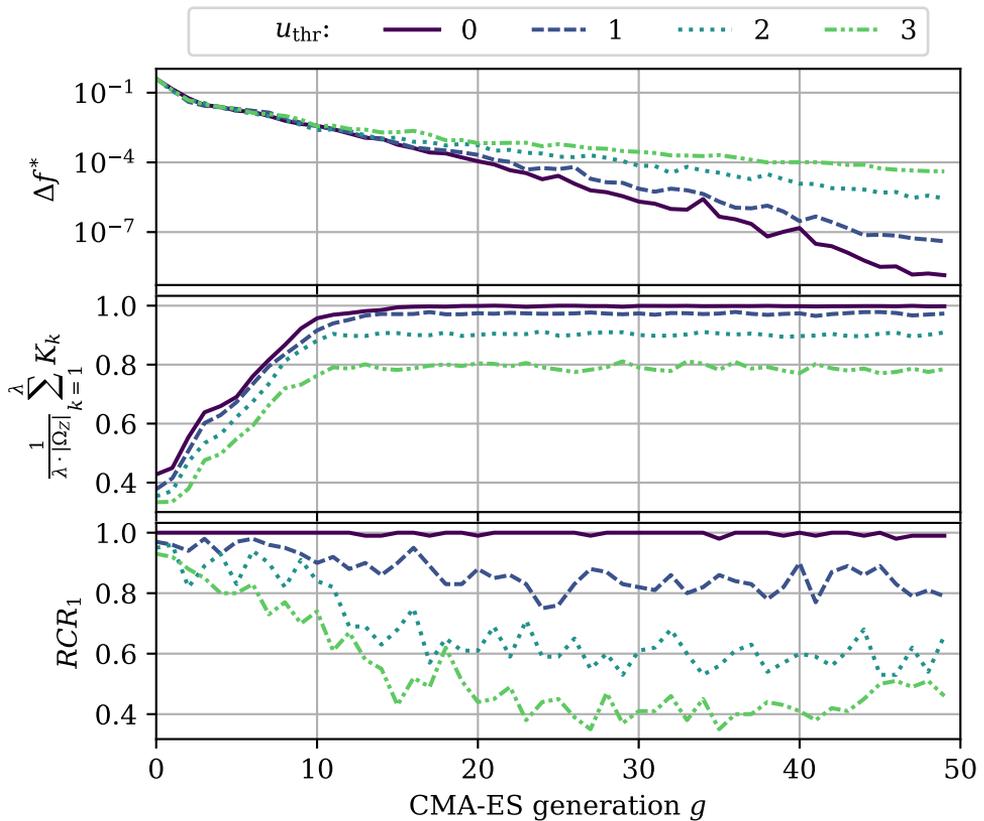


Figure 6.11: Mean distance to the optimal objective function value Δf^* , the ratio of evaluated scenarios and RCR of the top-1 individual RCR_1 for each CMA-ES generation g across 100 runs. Different values of the UQiS threshold $u_{\text{thr}} \in \{0, 1, 2, 3\}$ for the DA (Algorithm 2) are considered. The standard deviation of the noisy test function is set to 0.5.

However, to reduce the number of evaluations, the threshold can be set greater than zero in the beginning. This will not significantly impact the RCR and convergence. The reason for this is that in the beginning, the objective function values of the individuals overlap less, resulting in fewer rank changes as more scenarios are allocated. Therefore, the initial ranking is often already correct. Furthermore, selecting the second-best individual can also lead to significant improvement, and the CMA-ES population evolves even with an incorrect ranking provided.

Finally, the DA is compared to SA with different numbers of allocated scenarios $K_k \in \{45, 40, 30, 20, 5\}$ per individual considered (Figure 6.12).

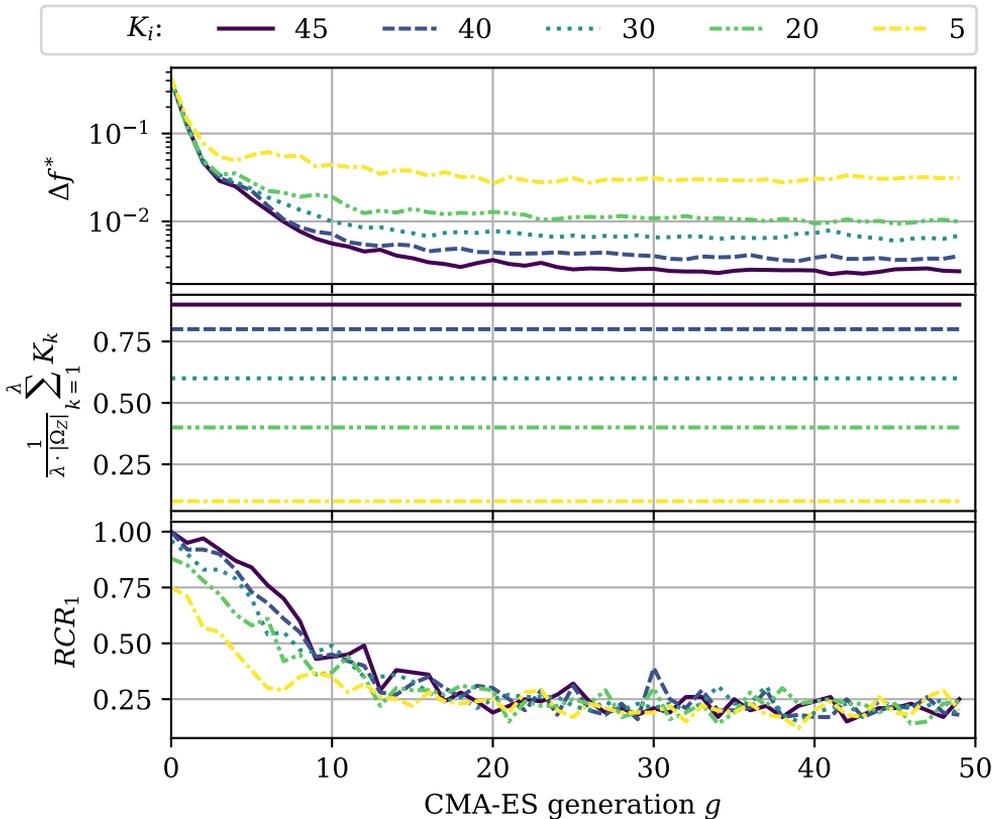


Figure 6.12: Mean distance to the optimal objective function value Δf^* , ratio of evaluated scenarios and RCR of the top-1 individual RCR_1 for each CMA-ES generation g across 100 runs. Different numbers of allocated scenarios $K_k \in \{45, 40, 30, 20, 5\}$ per individual for the SA are considered. The standard deviation of the noisy test function is set to 0.5.

Due to the SA, the number of allocated scenarios is constant across all generations. Thus, no adaptation to the current uncertainty in the ranking is conducted. Concerning convergence of CMA-ES, in the beginning, only five allocated scenarios to each individual are enough for progress, even with an RCR of below 0.75. However, the RCR of the other SA with more scenarios also falls below 0.5 within the first ten generations. Here, two clear drawbacks of the SA are present. To reach efficiently a constant high RCR, in the beginning, fewer evaluations are needed, and with further progress, more and more evaluations are needed for a correct ranking. Thus, a DA is required. A further drawback of SA is the allocation of the same number of scenarios to each individual, which wastes evaluations on the last ranked individuals with no information gain for ranking the top-1 individual.

6.2.5 Conclusion

The DA methodology (Algorithm 2) introduced in this work, which is based on uncertainty quantification, effectively addresses the drawbacks associated with SA. By employing a policy (Algorithm 3), scenarios are dynamically allocated to individuals within a population, thereby reducing the uncertainty in the ranking of these individuals. This methodology enables the ranking or selection of the top- μ individuals with a high degree of confidence. The DA methodology distributes function evaluations more effectively to the individuals than SA. However, the uncertainty quantification method requires the user to provide upper and lower bounds for the objective function values, which are a priori unknown and must be estimated.

The combination of the DA methodology and CMA-ES is not without its imperfections. Initially, the DA tends to allocate too many scenarios to achieve a correct ranking, which is unnecessary for the early stages of CMA-ES. The algorithm is capable of evolving effectively even with a ranking that is not perfectly accurate. Therefore, allowing for higher uncertainty, for example, by setting a higher threshold in the beginning, can reduce the number of allocated scenarios. Furthermore, no information about the uncertainty from past generations is used, the uncertainty quantification starts anew for each population. Only the bounds are adjusted based on information from previous generations. Using more information from past generations could be a valuable direction for future research.

6.3 Application to Real-World Problems

The optimization of parameters in vehicle dynamics control systems, such as ABS and ARP, is a challenging task due to the inherent noise. The objective function value for a given parameter configuration is determined by the sample mean across several repeated maneuvers, with each maneuver considered a distinct scenario.

In this section, the uncertainty quantification and DA methodology is applied to the two-dimensional real-world problem y_1 , which involves a partially loaded vehicle with high-performance tires (Section 3.4). For each of the 10101 possible parameter configurations, 30 braking maneuvers are simulated, and the resulting braking distances are obtained.

The objective function for the real-world problem y_1 returns the distance of the braking distance in meters to the global optimal average braking distance across all 30 considered scenarios. Figure 6.13 illustrates these distances, which are averaged over various numbers of randomly selected scenarios. The first input parameter x_1 is varied, while the second input parameter x_2 is held at its optimal value.

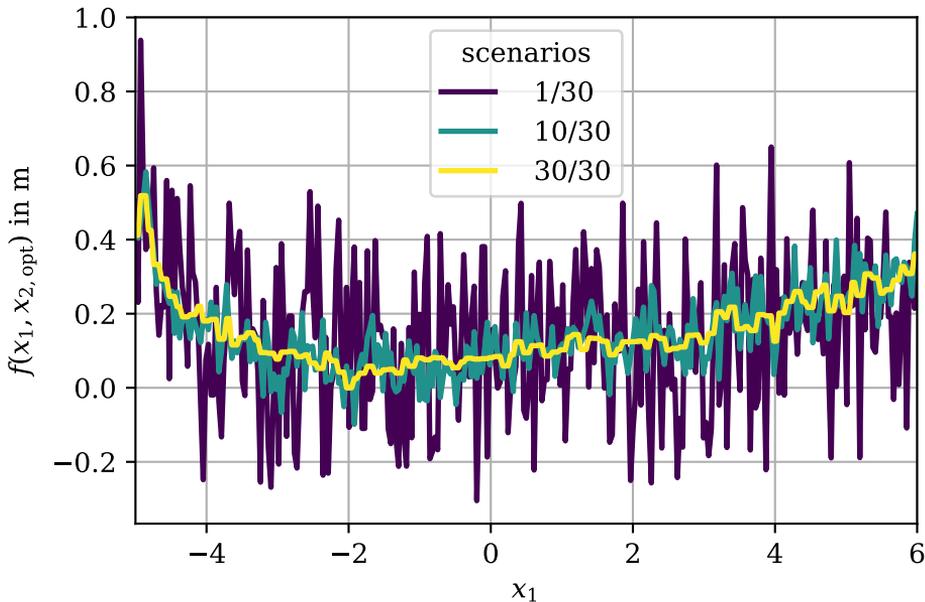


Figure 6.13: Distance of the braking distance in meters to the global optimal average braking distance across different numbers of randomly selected scenarios for the two-dimensional real-world problem y_1 (Table 3.1). The first input parameter x_1 is varied, while the second input parameter x_2 is held at its optimal value.

The dataset for the two-dimensional real-world problem y_1 consists of 303030 objective function values, which maps the objective function landscape completely. Evaluating the objective function by using the dataset is inexpensive and eliminates the need for further costly simulations. This allows to analyze the performance and convergence of CMA-ES in combination with the DA methodology on the real-world problem, similar to Section 6.2.4. For the analysis, CMA-ES is configured with a population size of six individuals. The DA methodology allocates scenarios during each generation of CMA-ES until the UQiS of the top-1 individual reaches zero. For each significance level $\alpha \in \{0.3, 0.6, 0.9\}$, 100 CMA-ES runs of 50 generations each are performed. The initial lower and upper bounds of the objective function values are set between -0.3 and 1.2 m. Figure 6.14 presents the results.

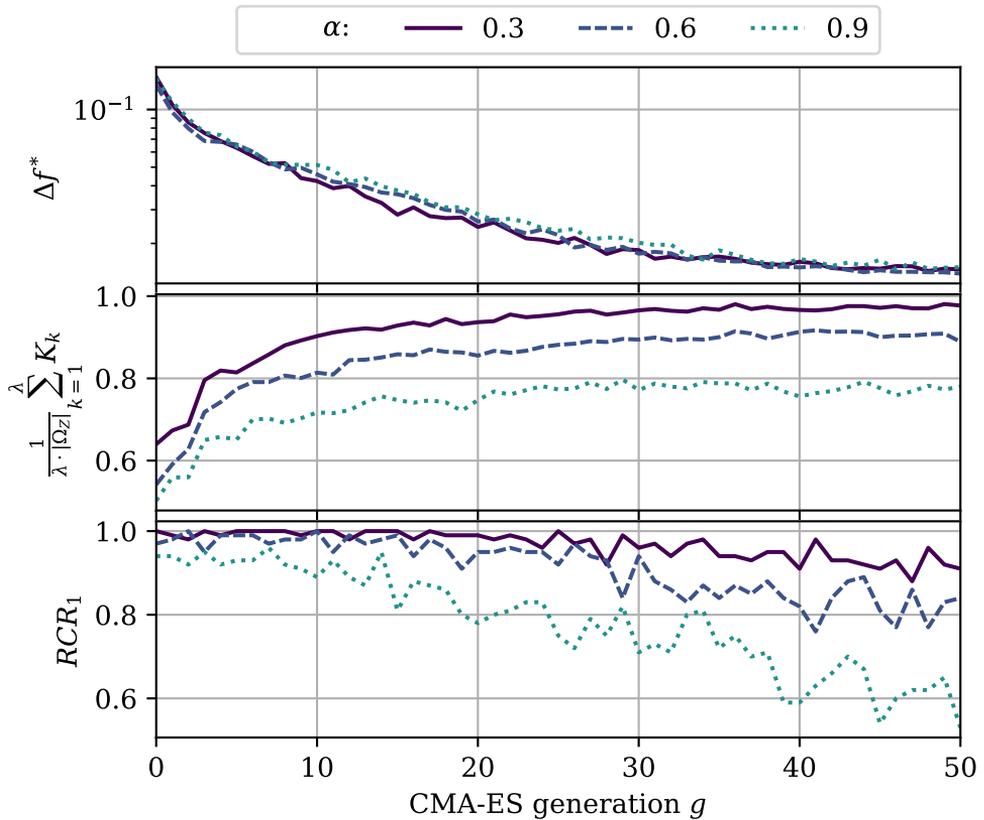


Figure 6.14: Mean distance to the optimal objective function value Δf^* , the ratio of evaluated scenarios and RCR of the top-1 individual RCR_1 for each CMA-ES generation g across 100 runs on the two-dimensional real-world problem y_1 (Table 3.1). Different values of the significance level $\alpha \in \{0.3, 0.6, 0.9\}$ (Equation 2.26) for the DA (Algorithm 2) are considered.

Across the 50 generations of CMA-ES, the distance to the optimal objective function value decreases consistently for all three significance levels $\alpha \in \{0.3, 0.6, 0.9\}$ considered, with no discernible differences in convergence. However, the ratio of evaluated scenarios is lower for higher significance levels across all 50 generations. Thus, on the real-world problem, a higher significance level for the uncertainty quantification is more efficient when applied in combination with CMA-ES. The lower RCR results from higher significance levels and decreases further across generations. This does not negatively impact the convergence of CMA-ES and is similar to the results on the noisy test function in Section 6.2.4. The DA methodology demonstrates that evaluations can be significantly reduced compared to evaluating all individuals across all scenarios. Specifically, with significance levels of 0.3 and 0.9, reductions of 8% and 26%, respectively, are achieved.

Regardless of the selected significance level, not every CMA-ES run reaches the optimal value of the objective function. On average, the deviation from the optimal value is approximately 2 cm. The reason for this is the presence of noise, even if all possible scenarios are evaluated. This noise contributes to a multimodal objective function landscape and can cause CMA-ES to converge to a local optimum instead of the global optimum. Two potential strategies can be employed to address this challenge: increasing the population size or the number of possible scenarios. First, a larger population size enhances the exploratory capabilities of the optimization algorithm. Thus, the optimization algorithm is more likely to escape a local optima. Second, a larger number of scenarios reduces the impact of noise and smoothes the objective function landscape. This simplifies the identification of the true global optimum for the optimization algorithm.

The RCR is with a significance level of 0.3 above 0.9 across all 50 generations. Thus, a significance level of 0.3 is effective for selecting the best parameter configuration with high probability, regardless of the generation. In the last generations of CMA-ES, almost all scenarios must be allocated to achieve correct ranking. Thus, even low levels of noise can complicate the differentiation between individuals. This could serve as a termination criterion or a sign to increase the number of scenarios.

Combining the DA methodology with CMA-ES can reduce the required evaluations by up to 26% in a real-world setting. However, the precision of uncertainty quantification for CMA-ES may be excessive, as the algorithm can progress even with some incorrect rankings, suggesting potential for further efficiency gains. Nonetheless, the DA methodology is adept at ranking individuals correctly and efficiently within a population.

The DA methodology and associated uncertainty quantification are not restricted to simulated objective function values but can also be applied to real-world vehicle measurements on a test track. Uncertainty quantification can provide information to reduce the number of maneuvers required. Furthermore, this information can guide engineers when choosing between parameter configurations of comparable quality.

