Black-Box Min–Max Continuous Optimization Using CMA-ES with Worst-case Ranking Approximation

Atsuhiro Miyagi
Taisei Corporation & University of Tsukuba
atsuhiro@bbo.cs.tsukuba.ac.jp

Jun Sakuma
University of Tsukuba & RIKEN AIP
jun@cs.tsukuba.ac.jp

Kazuto Fukuchi
University of Tsukuba & RIKEN AIP
fukuchi@cs.tsukuba.ac.jp

Youhei Akimoto
University of Tsukuba & RIKEN AIP
akimoto@cs.tsukuba.ac.jp

ABSTRACT

In this study, we investigate the problem of min–max continuous optimization in a black-box setting $\min_x \max_y f(x, y)$. A popular approach updates x and y simultaneously or alternatingly. However, two major limitations have been reported in existing approaches. (I) As the influence of the interaction term between x and y (e.g., $x^\mathsf{T} By$) on the Lipschitz smooth and strongly convex–concave function f increases, the approaches converge to an optimal solution at a slower rate. (II) The approaches fail to converge if f is not Lipschitz smooth and strongly convex–concave around the optimal solution. To address these difficulties, we propose minimizing the worst-case objective function $F(x) = \max_y f(x, y)$ directly using the covariance matrix adaptation evolution strategy, in which the rankings of solution candidates are approximated by our proposed worst-case ranking approximation (WRA) mechanism. Compared with existing approaches, numerical experiments show two important findings regarding our proposed method. (1) The proposed approach is efficient in terms of f-calls on a Lipschitz smooth and strongly convex–concave function with a large interaction term. (2) The proposed approach can converge on functions that are not Lipschitz smooth and strongly convex–concave around the optimal solution, whereas existing approaches fail.

CCS CONCEPTS

• **Mathematics of computing → Continuous optimization**;

KEYWORDS

Black-Box Min–Max Continuous Optimization, Worst-case Ranking Approximation, Covariance Matrix Adaptation Evolution Strategy

1 INTRODUCTION

Background. Simulation-based optimization has been utilized in various fields. In such optimizations, numerical simulations are used to evaluate the objective function on a solution candidate, with the conditions decided by preliminary investigation. For example, the coefficients in the governing equation or constitution rule should be decided beforehand when using the finite element method. However, these predetermined conditions contain uncertainties in many cases. Hence, the numerical simulation results contain uncertainties [8, 21]. Therefore, finding a robust solution against these uncertainties is desirable for simulation-based optimization.

To obtain a robust solution, previous studies for the automatic berthing control problem [3] and the electromagnetic scattering design problem [6] formulated a min–max optimization determined by following:

$$\min_{x \in \mathbb{X}} \max_{y \in \mathbb{Y}} f(x, y) \ , \tag{1}$$

where $f(x, y)$ is the objective function, $x \in \mathbb{X} \subseteq \mathbb{R}^m$ is a *design variable*, and $y \in \mathbb{Y} \subseteq \mathbb{R}^n$ is an uncertain parameter that denotes the numerical simulation conditions, which is refered as the *scenario variable* in this study. A naive approach is to select the recommended scenario variable $y_{\text{est}} \in \mathbb{Y}$ based on expert judgment and obtain the following optimum solution: $x_{y_{\text{est}}} = \operatorname{argmin}_{x \in \mathbb{X}} f(x, y_{\text{est}})$. However, because of the discrepancy between y_{est} and the scenario variable in a real environment y_{real}, the performance of the solution $x_{y_{\text{est}}}$ in the simulation does not guarantee satisfactory performance in the real-world environment. The main concern is that the solution may result in $f(x_{y_{\text{est}}}, y_{\text{real}}) \gg f(x_{y_{\text{est}}}, y_{\text{est}})$. Nevertheless, the optimal solution to (1) guarantees the lower-bound of the performance in a real environment, i.e., $f(x, y_{\text{real}}) \leqslant F(x) = \max_{y \in \mathbb{Y}} f(x, y)$ as long as $y_{\text{real}} \in \mathbb{Y}$. Therefore, the solution obtained by (1) performs well in a real-world environment. When the appropriate y_{est} cannot be selected, it is important to consider the min–max problem shown in (1).

In this study, we consider min–max optimization (1) with the following properties: the gradient information of the objective function f is unavailable (derivative-free optimization), and the objective function f is not mathematically and explicitly expressed. Moreover, its characteristic constants, such as its Lipschitz constant, are

Atsuhiro Miyagi, Kazuto Fukuchi, Jun Sakuma, and Youhei Akimoto

unavailable (black-box optimization). We refer to such a problem as the *black-box min–max optimization* problem.

Related works. Liu et al. [17] proposed the zero-order projected gradient descent ascent (ZOPGDA), which searches the optimal solution using the approximated gradient of the objective function for x and y. This approach updates x and y at iteration t repeatedly as follows:

$$(x^{t+1}, y^{t+1}) = (x^t, y^t) + (-\eta_x \cdot B_x, \eta_y \cdot B_y) \ , \tag{2}$$

where t is the number of iterations, η_x and η_y are the learning rates, and B_x and B_y are the approximated gradients $(\bar{\nabla}_x f, \bar{\nabla}_y f)$ of the objective function. Numerical experiments showed that ZOPGDA is superior to STABLEOPT [7], in terms of the scalability of the optimization time against the problem dimension; STABLEOPT is a derivative-free approach based on Bayesian optimization.

In another previous study [3], the optimization approach Adversarial-CMA-ES was proposed for the black-box min–max optimization. Adversarial-CMA-ES updates x and y using (2) with $B_x = \bar{x}^t - x^t$ and $B_y = \bar{y}^t - y^t$, where $\bar{x}^t$ and $\bar{y}^t$ are approximate solutions to $\operatorname{argmin}_{x \in \mathbb{X}} f(x, y^t)$ and $\operatorname{argmax}_{y \in \mathbb{Y}} f(x^t, y)$, respectively. Unlike ZOPGDA, which should set the learning rate according to the characteristics of the objective function, such as the maximum and minimum eigenvalues of the Hessian of the objective function, Adversarial-CMA-ES is designed to adapt the learning rate $\eta = \eta_x = \eta_y$ in (2) during the optimization. Therefore, Adversarial-CMA-ES should be more practical for black-box min–max optimization. In numerical experiments, Adversarial-CMA-ES was compared with co-evolutionary approaches [5, 14, 20], which are also derivative-free approaches for black-box min–max optimization. Adversarial-CMA-ES outperformed the existing co-evolutionary approaches in the worst-case scenario. It was observed for co-evolutionary approaches to fail to converge to the optimal solution, even on a strongly convex–concave and Lipschitz smooth (gradient is Lipschitz continuous) function.

Despite their promising results on some problems, ZOPGDA and Adversarial-CMA-ES have limitations [3]. **Difficulty (I)**: When the objective function is Lipschitz smooth and strongly convex–concave, the number of f-calls that Adversarial-CMA-ES performs until it locates an ϵ-optimal solution (a solution x around the optimum solution x^* satisfying $f(x) \leqslant f(x^*) + \epsilon$ for some $\epsilon > 0$) scales as $O\left(1 + \sigma_{\max}^2(H_{x,x}^{-1/2} H_{x,y}(-H_{y,y})^{-1/2})\right)$, where $H_{x,x} = \nabla_x \nabla_x f(x^*, y^*)$, $H_{x,y} = \nabla_x \nabla_y f(x^*, y^*)$, and $H_{y,y} = \nabla_y \nabla_y f(x^*, y^*)$ are the blocks of the Hessian matrix of f at the global min–max saddle point (x^*, y^*); $\sigma_{\max}(\cdot)$ represents the maximum singular value. In other words, the convergence slows down as the influence of the interaction term between x and y, $H_{x,y}$ grows. **Difficulty (II)**: Adversarial-CMA-ES fails to converge to a min–max saddle point if the objective function is not a strongly convex–concave and Lipschitz smooth function. Although these issues have only been reported for Adversarial-CMA-ES; similar limitations have been reported for the first-order approach [16] on which ZOPGDA is based. In our experiments with ZOPGDA, these limitations were observed. These situations occur frequently, and therefore are important issues that must be addressed.

In this study, we consider a black-box min–max optimization approach that can address the aforementioned issues.

Contribution. The study makes the following contributions:

(1) A black-box min–max optimization approach, covariance matrix adaptation evolution strategy (CMA-ES) with the worst-case ranking approximation (CMA-ES with WRA), is proposed. CMA-ES with WRA aims to optimize the worst-case objective function F using CMA-ES [2, 10, 12], while the rankings of the worst-case objective function values of the solution candidates are approximated by the WRA mechanism. The WRA mechanism approximates the rankings of the solution candidates by solving the internal maximization problems approximately, $\max_y f(x, y)$, for each solution candidate using CMA-ES with a warm starting strategy and an early stopping strategy.

(2) To verify that CMA-ES with WRA can address Difficulty (I), we compared CMA-ES with WRA with the existing approaches, ZOPGDA and Adversarial-CMA-ES, on a Lipschitz smooth and strongly convex–concave function. We empirically observed that CMA-ES with WRA could locate an ϵ-optimal solution within $O\left(\log(\sigma_{\max}(H_{x,x}^{-1/2} H_{x,y}(-H_{y,y})^{-1/2}))\right)$. We provide a theoretical but not rigorous reasoning for this scaling.

(3) To verify that CMA-ES with WRA can address Difficulty (II), we conducted numerical experiments on functions that were not strongly convex–concave and Lipschitz smooth around the optimal solution. We compared CMA-ES with WRA with the existing approaches for these test problems. CMA-ES with WRA could locate an ϵ-optimal solution, whereas the existing approaches failed.

(4) Additionally, we investigated how the number of f-calls performed by CMA-ES with WRA changed if the coefficient of the interaction term, $H_{x,y}$, changed on functions that were not strongly convex–concave and Lipschitz smooth around the optimal solution. We observed a similar scaling of the number of f-calls to that of a Lipschitz smooth and strongly convex–concave function.

Our implementation of CMA-ES with WRA is publicly available.[1]

2 PROBLEM DESCRIPTION

Our objective is to find the optimal solution x^* that minimizes the *worst-case objective function* $F : \mathbb{X} \to \mathbb{R}$, defined as followings:

$$F(x) = \max_{y \in \mathbb{Y}} f(x, y), \tag{3}$$

where $f : \mathbb{X} \times \mathbb{Y} \to \mathbb{R}$ is the objective function, and $\mathbb{X} \subseteq \mathbb{R}^m$ and $\mathbb{Y} \subseteq \mathbb{R}^n$ are the search domains for the design and scenario variables, respectively. As mentioned in the introduction, we consider derivative-free and black-box situations. Therefore, the worst-case objective function F is not explicitly available.

We introduce the definition of a *min–max saddle point* of f and the *strong convexity–concavity*. A neighborhood of $z^* \in \mathbb{W} \subseteq \mathbb{R}^\ell$ is defined as a subset $\mathcal{E}_z \subseteq \mathbb{W}$, such that there exists an open ball $\mathbb{B}(z^*, r) = \{z \in \mathbb{W} : \|z - z^*\| < r\}$ included in $\mathcal{E}_z$. A critical point of f is a point (x, y), such that $\nabla f(x, y) = (\nabla f_x(x, y), \nabla f_y(x, y)) = 0$.

Definition 2.1 (min–max saddle point [3]). A point $(\tilde{x}, \tilde{y}) \in \mathbb{X} \times \mathbb{Y}$ is a local min–max saddle point of a function $f : \mathbb{X} \times \mathbb{Y} \to \mathbb{R}$, if there exists a neighborhood $\mathcal{E}_x \times \mathcal{E}_y \subseteq \mathbb{X} \times \mathbb{Y}$, including $(\tilde{x}, \tilde{y})$, such that for any $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$, the condition $f(\tilde{x}, y) \leqslant f(\tilde{x}, \tilde{y}) \leqslant f(x, \tilde{y})$

---

[1]https://gist.github.com/a2hi6/ac511f101a494197b5fab56a407aa094

holds. If $\mathcal{E}_x = \mathbb{X}$ and $\mathcal{E}_y = \mathbb{Y}$, the point $(\tilde{x}, \tilde{y})$ is called the global min–max saddle point. A strict min–max saddle point is one where the equality only holds if $(x, y) = (\tilde{x}, \tilde{y})$. A saddle point that is not a strict min–max saddle point is called a weak min–max saddle point.

Definition 2.2 (strongly convex concave function [3]). A twice continuously differential function $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$ is locally μ-strongly convex–concave around a critical point $(\check{x}, \check{y})$ for some $\mu > 0$ if there exist open sets $\mathcal{E}_x \subseteq \mathbb{R}^m$ including $\check{x}$ and $\mathcal{E}_y \subseteq \mathbb{R}^n$ including $\check{y}$, such that $H_{x,x}(x, y) - \mu \cdot I$ and $-H_{y,y}(x, y) - \mu \cdot I$ are non-negative definite for all $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$. The function f is a globally μ-strongly convex–concave function if $\mathcal{E}_x = \mathbb{X}$ and $\mathcal{E}_y = \mathbb{Y}$. f is locally or globally strongly convex–concave if it is locally or globally μ-strongly convex–concave for some $\mu > 0$.

Importantly, if the objective function f is twice-continuously differentiable and globally strongly convex–concave, there exists a unique critical point $(x^*, y^*) \in \mathbb{R}^m \times \mathbb{R}^n$, which is a unique min–max saddle point of f. Moreover, x^* is a unique global minimum point of the worst-case objective function F. Let the *worst-case scenario variables* $\hat{y}(x)$ for x be defined as $\hat{y}(x) = \operatorname{argmax}_y f(x, y)$, i.e., $F(x) = f(x, \hat{y}(x))$. Then, it is known that $\hat{y}$ is uniquely determined, continuously differentiable, and $\hat{y}(x^*) = y^*$ [3].

3 LIMITATIONS OF EXISTING APPROACHES

The existing approaches for the derivative-free min–max optimization problems, ZOPGDA [17] and Adversarial-CMA-ES [3], are designed to converge to a local min–max saddle point of f under the assumption that f is at least locally strongly convex–concave. If the objective is locally strongly convex–concave, the simultaneous update of x and y of the form (2) is intuitive. The reason for this is as follows. The worst-case scenario $\hat{y}(x^{t+1})$ is supposed to be close to $\hat{y}(x^t)$ if x^{t+1} and x^t are close. If y^t approximates $\hat{y}(x^t)$ well, the next scenario y^{t+1} that is close to y^t is expected to approximate $\hat{y}(x^{t+1}) \approx \hat{y}(x^t) \approx y^t$. It has been demonstrated in [3] that this type of approach can converge linearly toward the global min–max saddle point if the learning rates η_x and η_y are sufficiently small.

However, as mentioned in the introduction, several limitations are highlighted in [3]. Among them, we focus on Difficulty (I) and (II), which have been introduced in Section 1. Here, we elaborate on them with some examples.

Difficulty (I). The learning rate must be sufficiently small for convergence, depending on the *interaction term* between x and y of the objective function. For example, consider $f(x, y) = (a/2)x^2 + bxy - (c/2)y^2$. Then, [3] shows that the learning rate needs to be set proportional to $O(ac/(ac + b^2))$. As the coefficient of the interaction term, b, increases, compared with the coefficients of the non-interaction terms a and c, the learning rate needs to be smaller. This results in a slower convergence, where the number of f-calls scales as $1 + b^2/(ac)$. A similar negative result was shown in [16] for the first-order simultaneous update of x and y. ZOPGDA is an approximation of the first-order counterpart; thus, the same limitation is expected and was observed in our experiments.

Difficulty (II). Adversarial-CMA-ES reportedly fails to converge to a local min–max saddle point if f is not strongly convex–concave and Lipschitz smooth (that is, the gradient is Lipschitz continuous) [3]. One such example is the bi-linear function $f(x, y) = xy$

on a bounded domain $[-1, 1] \times [-1, 1]$. The failure of convergence of the first-order simultaneous update is also reported in [16]. Therefore, ZOPGDA is also considered to fail to converge. Another example is $f(x, y) = (1/4)x^4 + bxy - (1/4)y^4$. Here, the situation is similar to that of $f(x, y) = (a/2)x^2 + bxy - (c/2)y^2$ but $a = x^2/2$ and $c = y^2/2$, i.e., $ac/(ac + b^2)$ becomes smaller as x and y approach 0. Therefore, we expect the learning rate to be smaller as the algorithm approaches the global min–max saddle point, jeopardizing the linear convergence. For both examples, we observe that Adversarial-CMA-ES and ZOPGDA fail to converge in our experiments.

A possible cause of these limitations is the sensitivity of the worst-case scenarios $\hat{y}(x)$ against small changes in x. In the case of $f(x, y) = (a/2)x^2 + bxy - (c/2)y^2$, we have $\hat{y}(x) = (b/c)x$, i.e., the change in the worst-case scenario is proportional to b/c. If $b \gg c$, a small change in x may lead to a great change in $\hat{y}(x)$. Then, the simultaneous update (2) may fail to keep track of the worst-case scenario. To prevent this, the learning rate η must be set sufficiently small, resulting in a slow convergence. In the case of $f(x, y) = xy$ on $[-1, 1] \times [-1, 1]$, the worst-case scenario is $\hat{y}(x) = \operatorname{sign}(x)$, which is not a continuous function of x around $x^* = \operatorname{argmin}_x F(x) = 0$. A small change in x near $x^* = 0$ could result in a sign flip for x, changing the worst-case scenario between -1 and 1. Consequently, the simultaneous update (2) may fail to keep track of the worst-case scenario. Thus, in this case, a small learning rate is ineffective.

4 PROPOSED APPROACH

We propose a novel approach to black-box min–max optimization problems (1), named the CMA-ES with the worst-case ranking approximation (CMA-ES with WRA). This approach attempts to minimize the worst-case objective function F directly using CMA-ES [2, 10, 12] to mitigate the aforementioned limitations of the existing approaches. The worst-case objective function value $F(x)$ for each solution candidate x is then approximated by solving the maximization problem $\max_y f(x, y)$. The proposed worst-case ranking approximation (WRA) mechanism uses a warm starting strategy and an early stopping strategy to reduce the number of f-calls for the internal maximization problems.

4.1 Addressing Difficulty (I) and (II)

Our main strategy to address Difficulty (I) and (II) described in the previous section is to minimize the worst-case objective function F directly. Here, we explain why minimizing F is expected to mitigate these difficulties.

First, we explain why we expect that minimizing F will not suffer from a large interaction term (**Difficulty (I)**). An arbitrary strongly convex–concave and Lipschitz smooth function can be approximated by a convex–concave quadratic function around the global min–max saddle point. Therefore, for simplicity, consider a convex–concave quadratic function $f(x, y) = \frac{1}{2}x^\mathsf{T}Ax + x^\mathsf{T}By - \frac{1}{2}y^\mathsf{T}Cy$, where $A \in \mathbb{R}^{m \times m}$ and $C \in \mathbb{R}^{n \times n}$ are symmetric positive definite, and $B \in \mathbb{R}^{m \times n}$ is an arbitrary matrix. The worst-case scenario is $\hat{y}(x) = C^{-1}B^\mathsf{T}x$, and the worst-case objective function is $F(x) = f(x, \hat{y}(x)) = \frac{1}{2}x^\mathsf{T}(A + BC^{-1}B^\mathsf{T})x$. Irrespective of the coefficients, this is a convex quadratic function. An approach exploiting the second-order information, such as CMA-ES, empirically shows

linear convergence on an arbitrary convex quadratic function with a convergence rate independent of its Hessian matrix [13]. Therefore, minimizing F by CMA-ES is expected to show a linear convergence with a convergence rate independent of the interaction term.

Second, we show how the proposed approach can mitigate the issue of convergence on a function f that is not strongly convex–concave and Lipschitz smooth (**Difficulty (II)**). This is, to some extent, intuitive. The proposed approach directly minimizes the worst-case objective F; thus, it will converge toward a local optimal solution of F if F is solvable by the search algorithm. For example, in the case of a bi-linear objective function $f(x, y) = xy$ on $[-1, 1] \times [-1, 1]$, the worst-case objective function is $F(x) = |x|$. This is a monotonic transformation of a quadratic function x^2. A comparison-based search algorithm, such as CMA-ES, is known to be invariant to the monotonic transformation of the objective function. Therefore, if a comparison-based search algorithm that can solve a quadratic function efficiently is used, the worst-case objective function F can also be efficiently solved.

4.2 Worst-case Ranking Approximation

The difficulty in directly solving the worst-case objective function F is that we must evaluate $F(x)$ for each solution candidate x by solving the maximization problem $\max_y f(x, y)$. The maximization problem cannot be solved precisely because f is a black-box function. Hence, one must rely on a numerical solver. However, this is time–consuming because each $F(x)$ evaluation requires a single optimization process, which necessitates several f-calls.

For each maximization problem, we use (a) warm starting and (b) early stopping of the numerical solver to reduce the number of f-calls. The proposed approach uses CMA-ES as the numerical solver for the worst-case objective function F. At each iteration t, it samples λ solution candidates, $x_1, \ldots, x_\lambda$, from the Gaussian distribution $\mathcal{N}(m^t, \Sigma^t)$. Their rankings, denoted as $\text{Rank}_F(\{x_i\}_{i=1}^\lambda)$, are then computed based on F, which is now approximated by solving $\max_y f(x, y)$. The distribution parameters, mean vector m, covariance matrix Σ, and any other dynamic parameters θ, are updated based on the solution candidates and their rankings. We have two important remarks. (1) CMA-ES [2, 10, 12] is comparison-based; thus, it behaves identically on F and its approximation $\hat{F}$ if $\text{Rank}_F(\{x_i\}_{i=1}^\lambda)$ and $\text{Rank}_{\hat{F}}(\{x_i\}_{i=1}^\lambda)$ are equivalent. That is, $\hat{F}$ does not need to approximate F more accurately than that required to approximate the rankings. This point is important in designing a stopping condition for the maximization problem. (2) The search distribution $\mathcal{N}(m^t, \Sigma^t)$ does not significantly change in one iteration; therefore, the solution candidates generated in the current and last iteration are similarly distributed. Therefore, the worst-case scenario for the solution candidates generated in this iteration are expected to be distributed similarly to the solution candidates. This suggests the necessity of looking for the worst-case scenarios based on the result of previous iterations.

Therefore, we design the worst-case ranking approximation (WRA) mechanism. It takes λ solution candidates as the input and returns their approximate rankings. To reduce f-calls inside WRA, the warm starting and early stopping strategies are implemented. The algorithm of WRA is summarized in Algorithm 1. It internally maintains

λ CMA-ES instances for the worst-case scenario search. The following sections cover the warm starting and early stopping strategies for these internal CMA-ES instances.

4.2.1 Warm Starting Strategy.
For each solution candidate x_i, we select and run one of the internal CMA-ES to approximate $F(x_i)$. The purpose of the warm starting strategy is to help us choose an internal CMA-ES that will reduce the number of f-calls.

The warm starting part is described in Line 1–Line 6 of Algorithm 1. Let y_k be the worst scenario parameter obtained by the kth CMA-ES instance in the last iteration. Then, for each solution candidate x_i ($i = 1, \ldots, \lambda$), we compute the objective function values $f(x_i, y_k)$ for $k = 1, \ldots, \lambda$. The worst-case scenario is then selected. Let $k_i^{\text{worst}} = \text{argmax}_k f(x_i, y_k)$ be the index of the worst-case scenario. Then, we select the k_i^{worst}th CMA-ES instance for the worst-case scenario search for x_i. Starting with the CMA-ES instance that generates the worst-case scenario for x_i, we expect that the number of f-calls for the adaptation of the distribution parameters to be significantly lower, as compared to using a new CMA-ES with initial distribution parameters.[2] If the same CMA-ES instance is selected for a different solution candidate, a clone is created.

4.2.2 Early Stopping Strategy.
In a double loop strategy, determining the best time to stop the internal maximization process is difficult. However, as mentioned earlier, we can stop the worst-case scenario search without any compromise if the rankings, $\text{Rank}_F(\{x_i\}_{i=1}^\lambda)$, of the worst-case objective function values of solution candidates are correctly estimated. Accordingly, this condition is eased. We stop the worst-case scenario search if Kendall's rank correlation coefficient τ [15] between the worst-case objective function values $\{F(x_i)\}_{i=1}^\lambda$ and its approximated values $\{\hat{F}(x_i)\}_{i=1}^\lambda$ is sufficiently high, for example, $\tau \geq \tau_{\text{threshold}}$. Two CMA-ES with a high τ value should behave similar in each other [1]; therefore, τ is frequently used to measure the quality of a surrogate function [4, 9, 18]. However, because we cannot obtain $F(x_i)$, τ cannot be computed. Therefore, τ is approximated using the worst-case objective function values estimated in the current iteration and those estimated in previous iterations.

The worst-case ranking approximation with an early stopping strategy is described in Line 7–Line 22 of Algorithm 1. Let the first estimate of the worst-case objective function value for each solution candidate x_i be denoted by F_i^0; then, all the CMA-ES instances are run for a certain number of iterations, which will be described later. We call it a *round*, and the round is counted by $j \geq 1$. The worst-case objective function value for each solution candidate x_i estimated after the round j is denoted by F_i^j. Then, τ between $\{F_i^j\}_{i=1}^\lambda$ and $\{F_i^{j-1}\}_{i=1}^\lambda$ is computed as τ between the ground truth worst-case objective function values $\{F(x_i)\}_{i=1}^\lambda$ and their estimates $\{F_i^{j-1}\}_{i=1}^\lambda$. A round is repeated until $\tau > \tau_{\text{threshold}}$.

---

[2]CMA-ES has dynamic parameters, θ, other than the distribution parameters, such as the so-called evolution paths. After each WRA call, we only keep the solution y_k and the distribution parameters (m_k, Σ_k) and initialize all the other parameters, θ, of each internal CMA-ES instance. That is, we avoid sharing the dynamic parameters θ for worst-case search for different solution candidates. This is to avoid a systematic bias caused by the change in the objective function because of the change in the solution candidate. The phenomenon is explained in [3].

In each round, each CMA-ES instance run is terminated if the worst-case scenario improves $c_{\max}$ times. Here, we assume that the worst-case scenario has been significantly improved over the last round. Additionally, the run is terminated if all the coordinate-wise standard deviation, $\sqrt{\Sigma_{\ell,\ell}}$, become smaller than the threshold $V_{\min}$. In this case, we assume that the search distribution has already converged, and the worst-case scenario will not be updated significantly anymore. However, if this condition is satisfied in the last call of WRA (that is, in the iteration of the CMA-ES solving F), this condition is immediately satisfied in the first round of the current call of WRA. To prevent this, we force the internal CMA-ES instance to run at least $T_{\min}$ iterations.

4.2.3 Post processing. After computing the worst-case ranking approximation, we perform post-processing (Line 23–Line 33 in Algorithm 1) for the next WRA call. First, we prevent the coordinate-wise standard deviation $\sqrt{[\Sigma_k]_{\ell,\ell}}$ from becoming smaller than $V_{\min}$. Otherwise, the termination condition in each round of the worst-case scenario search will be satisfied immediately after $T_{\min}$ iterations. Second, we prevent the Gaussian distributions of λ CMA-ES instances from converging to the same point. It is important to distance the CMA-ES instances from each other because the worst-case scenarios can be distinct, even for close solution candidates, for example, on a bi-linear function. If the output worst-case scenarios of two CMA-ES instances are close to each other (the distance is smaller than $V_{\min} \cdot \sqrt{n}$), we reset the distribution parameters of one of these instances.

5 NUMERICAL EXPERIMENTS

To test the following hypotheses, we compare CMA-ES with WRA with the existing approaches: Adversarial-CMA-ES[3] and ZOPGDA[4] in numerical experiments. (a) CMA-ES with WRA is more efficient in terms of the number of f-calls if the objective function is Lipschitz smooth and strongly convex–concave, but the influence of the interaction term between x and y is large (Section 5.2). (b) CMA-ES with WRA can converge to the optimal solution x^* even if the objective function is not locally Lipschitz smooth and strongly convex–concave around x^*, whereas the existing approaches fail to converge (Section 5.3). Additionally, we investigate how much the number of f-calls spent by CMA-ES with WRA scales when a coefficient b of the interaction term $bx^T y$ is changed on objective functions that are not necessarily Lipschitz smooth and strongly convex–concave around x^* (Section 5.4).

5.1 Common Settings

We designed eight test problems, summarized in Table 1. They are designed to have different characteristics (smoothness, convexity, and concavity) around the optimal solution of the objective function. The search domains of the design variables and scenario variables are $\mathbb{X} = [-3, 3]^m$ and $\mathbb{Y} = [-3, 3]^n$, respectively. The dimension of the design variables is $m = 20$, and the dimension of the scenario variables is $n = 20$.

The configuration of CMA-ES with WRA is as follows: The hyperparameters for WRA are set as follow: $\tau_{\text{threshold}} = 0.7$, $c_{\max} = 2$,

---

Algorithm 1 Worst-case ranking approximation

Require: solution candidates to be ranked: $x_1, \ldots, x_\lambda$
Require: stopping conditions: $\tau_{\text{threshold}} \in (0, 1]$, $c_{\max} \geqslant 1$, $V_{\min} \geqslant 0$
Require: $(m_1, \Sigma_1), \ldots, (m_\lambda, \Sigma_\lambda)$ and $y_i \sim N(m_i, \Sigma_i)$ for $i = 1, \ldots, \lambda$
1: // initialization part
2: **for** $i = 1, \ldots, \lambda$ **do**
3:      Evaluate $f(x_i, y_k)$ for all $k = 1, \ldots, \lambda$
4:      Select the worst index $k_i^{\text{worst}} = \text{argmax}_{k \in \{1, \ldots, \lambda\}} f(x_i, y_k)$ and let $\tilde{y}_i = y_{k_i^{\text{worst}}}, F_i^0 = f(x_i, \tilde{y}_i), \tilde{m}_i = m_{k_i^{\text{worst}}}, \tilde{\Sigma}_i = \Sigma_{k_i^{\text{worst}}}$
5:      Initialize internal parameter θ_i of a CMA-ES instance
6: **end for**
7: // worst-case ranking approximation part
8: $\tau = -1$, $j = 0$, $t_1, \ldots, t_\lambda = 0$, $h_1, \ldots, h_\lambda = 0$
9: **while** $\tau \leqslant \tau_{\text{threshold}}$ **do**
10:      $j = j + 1$
11:      **for** $i = 1, \ldots, \lambda$ **do**
12:          $c = 0$
13:          **while** $c < c_{\max}$ or $h_i = 0$ **do**
14:              Perform one iteration of CMA-ES with $(\tilde{m}_i, \tilde{\Sigma}_i, \theta_i)$ and obtain the worst candidate y_i' and the updated parameters $(\tilde{m}_i, \tilde{\Sigma}_i, \theta_i)$
15:              **if** $F_i^j < f(x_i^t, \tilde{y}_i')$ **then**
16:                  $F_i^j = f(x_i^t, \tilde{y}_i')$, $\tilde{y}_i = \tilde{y}_i'$, and $c = c + 1$
17:              **end if**
18:              $h_i = 1$ **if** $\max_\ell \left\{ \sqrt{[\tilde{\Sigma}_i]_{\ell,\ell}} \right\} < V_{\min}$ and $t_i \geqslant T_{\min}$
19:          **end while**
20:      **end for**
21:      $\tau = \text{Kendall}(\{F_i^{j-1}\}_{i=1}^\lambda, \{F_i^j\}_{i=1}^\lambda)$
22: **end while**
23: // post-process part
24: **for** $i = 1, \ldots, \lambda$ **do**
25:      $(y_i, m_i, \Sigma_i) = (\tilde{y}_i, \tilde{m}_i, \tilde{\Sigma}_i)$
26:      $D_i = \text{diag}\left(\max\left(1, \frac{V_{\min}}{\sqrt{[\Sigma_i]_{1,1}}} \right), \ldots, \max\left(1, \frac{V_{\min}}{\sqrt{[\Sigma_i]_{n,n}}} \right) \right)$
27:      $\Sigma_i = D_i \Sigma_i D_i$
28:      **for** $k = i + 1, \ldots, \lambda$ **do**
29:          **if** distance$(y_i, y_k) < V_{\min} \cdot \sqrt{n}$ **then**
30:              Reset (m_k, Σ_k) and sample $y_k \sim N(m_k, \Sigma_k)$
31:          **end if**
32:      **end for**
33: **end for**
34: **return** the rankings of $F_1^j, \ldots, F_\lambda^j$

$V_{\min} = 10^{-4}$, and $T_{\min} = 10$. The initial mean vectors and the covariance matrices of the internal CMA-ES instances are $m_i \sim \mathcal{U}(-3, 3)^n$ and $\Sigma_i = \text{diag}(1.5, \ldots, 1.5)^2$. When the distribution parameters of the internal CMA-ES instances are reset during the post processing phase of WRA, we use the same initialization. In the CMA-ES solving F, the initial mean vector is drawn from $m^0 \sim \mathcal{U}(-3, 3)^m$, and the initial covariance matrix is set to $\Sigma^0 = \text{diag}(1.5, \ldots, 1.5)^2$. The hyperparameters and the initial values of the dynamic parameters θ for the CMA-ES instances in WRA and the CMA-ES solving F are set to their default values, as proposed in [2].

The hyperparameters for ZOPGDA and Adversarial-CMA-ES are set based on the original studies [17] and [3], respectively. For ZOPGDA, the learning rate parameters were set to $\eta_x = 0.02$ and $\eta_y = 0.05$. For Adversarial-CMA-ES, we set $G_{\text{tol}} = 10^{-7}$. The distribution parameters are initialized in the same way as CMA-ES with WRA.

---

[3]https://gist.github.com/youheiakimoto/ab51e88c73baf68effd95b750100aad0
[4]https://github.com/KaidiXu/ZO-minmax

Table 1: Test problem definitions and their characteristics. The domains are $\mathbb{X} = [-3, 3]^m$ and $\mathbb{Y} = [-3, 3]^n$. The worst-case scenario for each solution $x \in \mathbb{X}$ is denoted by $\hat{y}(x) = (\hat{y}_1(x), \ldots, \hat{y}_n(x))$. The characteristics are denoted as follows: strongly convex (st-cv), strongly concave (st-cc), convex (cv), concave (cc), smooth (sm), and non-smooth (non-sm). The scalar $b > 0$ is introduced to control the influence of the interaction term between x and y.

f	x	y	$\hat{y}_i(x)$	Optimum
$f_1 = x^T y$	linear	linear	$3\,\mathrm{sign}(x_i)$	$x^* = 0$
$f_2 = \frac{1}{2}\|x\|_2^2 + x^T y$	sm st-cv	linear	$3\,\mathrm{sign}(x_i)$	$x^* = 0$
$f_3 = \frac{1}{2}\|x+1\|_2^2 + \frac{1}{10}x^T y$	sm st-cv	linear	$3\,\mathrm{sign}(x_i)$	$x^* = -0.7$
$f_4 = \frac{1}{2}\|x\|_2^2 + x^T y + \frac{1}{2}\|y\|_2^2$	sm st-cv	sm st-cv	$\begin{cases} 3 \text{ or} -3 & x_i = 0 \\ 3\,\mathrm{sign}(x_i) & 0 < |x_i| \leqslant 3 \end{cases}$	$x^* = 0$
$f_5 = \frac{1}{2}\|x\|_2^2 + bx^T y - \frac{1}{2}\|y\|_2^2$	sm st-cv	sm st-cc	$\begin{cases} bx_i & |x_i| \leqslant 3/b \\ 3\,\mathrm{sign}(x_i) & |x_i| > 3/b \end{cases}$	$x^* = 0$
$f_6 = \frac{1}{2}\|x\|_2^2 + \|x\|_1 + bx^T y - \|y\|_1 - \frac{1}{2}\|y\|_2^2$	non-sm st-cv	non-sm st-cc	$\begin{cases} 0 & |x_i| \leqslant 1/b \\ bx_i - \mathrm{sign}(x_i) & 1/b < |x_i| \leqslant 4/b \\ 3\,\mathrm{sign}(x_i) & 4/b < |x_i| \end{cases}$	$x^* = 0$
$f_7 = \frac{1}{4}\|x\|_2^4 + bx^T y - \frac{1}{4}\|y\|_2^4$	cv	cc	$\begin{cases} (b/\|x\|_2^2)^{1/3}x_i & (b/\|x\|_2^2)^{1/3}|x_i| \leqslant 3 \\ 3\,\mathrm{sign}(x_i) & (b/\|x\|_2^2)^{1/3}|x_i| > 3 \end{cases}$	$x^* = 0$
$f_8 = \|x\|_1 + bx^T y - \|y\|_1$	non-sm cv	non-sm cc	$\begin{cases} 0 & |x_i| \leqslant 1/b \\ 3\,\mathrm{sign}(x_i) & |x_i| > 1/b \end{cases}$	$x^* = 0$

To deal with the box constraint, ZOPGDA by default uses the projected gradient. Adversarial-CMA-ES and CMA-ES with WRA use the mirroring technique along with upper-bounding of the coordinate-wise standard deviation [22].

We evaluate the performance of each optimization algorithm by running 20 independent runs. The termination criteria are as follows. The number of f-calls in each run is limited to $20,000,000$. Before the number of f-calls reaches $20,000,000$, the run is considered a success if $|F(m^t) - F(x^*)| \leqslant 10^{-6}$ is satisfied. For ZOPGDA, m^t is considered the estimate x^t of the solution at iteration t.

5.2 Experiment 1

To verify hypothesis (a), we applied three approaches to f_5 at $b \in \{1, 3, 10, 30, 100\}$.

The results of the experiment are shown in Figure 1. As shown in Figure 1, CMA-ES with WRA improved the scalability regarding the number of f-calls until convergence at coefficient b compared with ZOPGDA and Adversarial-CMA-ES. The increment for the number of f-calls with b was approximately proportional to $\log(b)$. This result will be discussed in Section 6. In this experiment, when $10 \geqslant b$, the number of f-calls performed by CMA-ES with WRA was less than that by the others.

We consider the results of Adversarial-CMA-ES and ZOPGDA. First, the number of f-calls increased proportional to b^2 when $b = \{1, 3\}$. At small b, the existing approaches converged with less f-calls than CMA-ES with WRA. For $b = 10$, we expected from the fitted curve in Figure 1 that Adversarial-CMA-ES converges successfully within the f-calls budget. However, it failed. This was probably because of the box constraint. The theoretical analysis in [3] assumes unbounded domains. Under the box constraint in this experiment, the character of f_5 at $b = 10$ resembled that of f_1. Concretely, when a design variable is in $|x_i| > 0.3 = 3/b$ for each i, the ith coordinate of the worst scenario is $\hat{y}_i(x) = 3\,\mathrm{sign}(x_i)$,
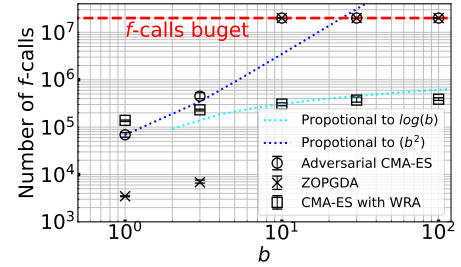

Figure 1: Comparison result among CMA-ES with WRA, ZOPGDA, and Adversarial-CMA-ES at $b \in \{1, 3, 10, 30, 100\}$ on f_5. Mean and standard deviation of the number of f-calls until successful convergence over 20 runs. ZOPGDA and Adversarial-CMA-ES failed to converge at $b \in \{10, 30, 100\}$.

which is the same as the worst scenario on f_1. As we will see in the next experiment, Adversarial-CMA-ES fails to converge to the optimal solution. Therefore, we believe that Adversarial-CMA-ES had difficulty converging toward the area with $|x_i| \leqslant 0.3$ for all $i = 1, \ldots, m$. Second, ZOPGDA could not converge to the optimal solution in any trials where $10 \geqslant b$ because the learning rate was not tuned.

5.3 Experiment 2

To verify hypothesis (b), we applied three approaches to $f_1 - f_4$ and $f_6 - f_8$. We set $b = 1$ for $f_6 - f_8$.

The results of the experiment are shown in Figure 2. Except for the trials on f_4, CMA-ES with WRA achieved successful convergence in all trials. Nevertheless, the existing approaches failed to determine the optimal solution in all trials. This was because the objective function in the neighborhood of the optimal solution was not a Lipschitz smooth and strongly convex–concave function.
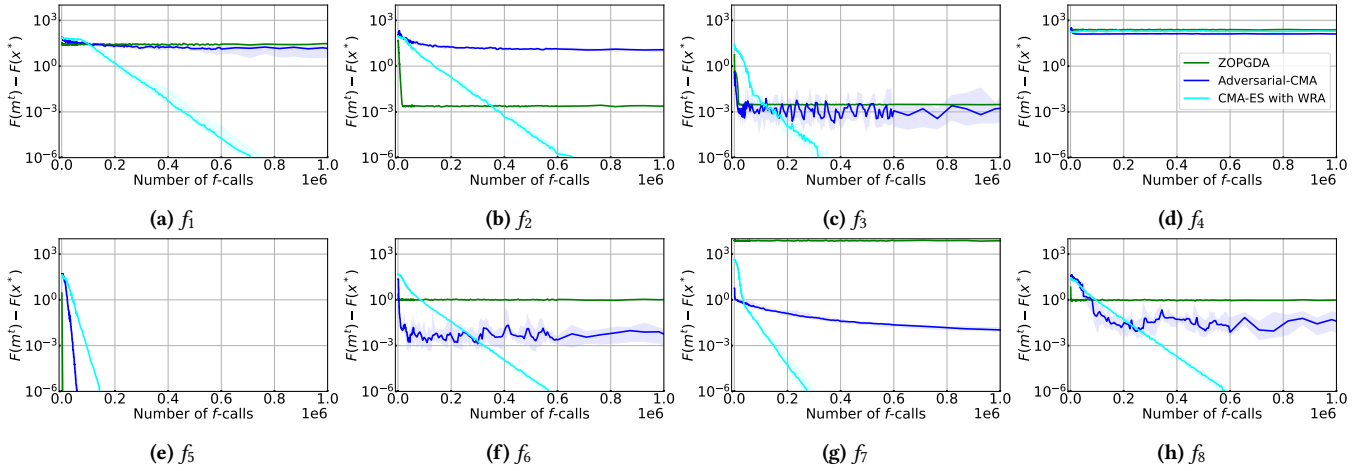
Figure 2: Gap $|F(m^t) - F(x^*)|$ with the number of f-calls at $b = 1$ on f_1–f_8. The solid line represents the median (50 percentile) over 20 runs. The shaded area represents the interquartile range (25–75 percentile) over 20 runs.
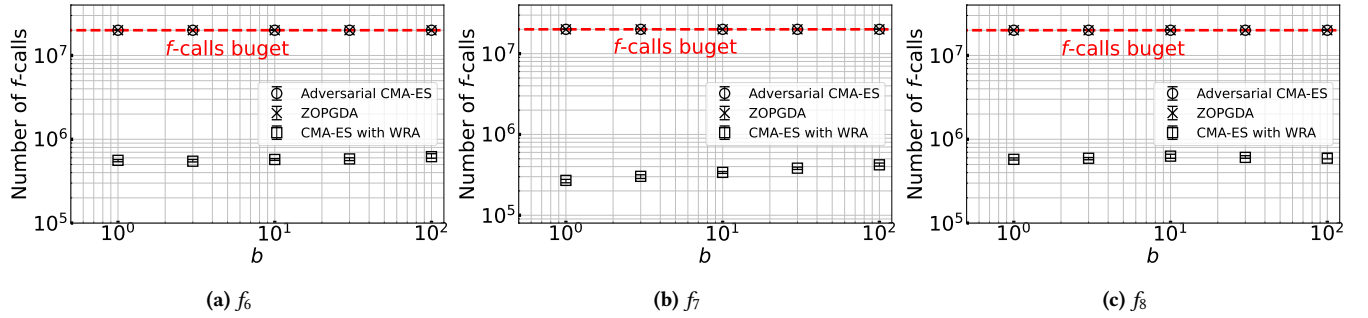

Figure 3: Comparison result among `CMA-ES with WRA`, `ZOPGDA` and `Adversarial-CMA-ES` at $b = \{1, 3, 10, 30, 100\}$ on f_6–f_8. The mean and standard deviation of the number of f-calls until successful convergence over 20 runs.

We discuss the results of `CMA-ES with WRA` on f_4. Let us consider the objective function $h(y) = f(x, y)$ for y on a solution candidate $x \in \mathbb{X}$. This objective function $h(y)$ has the local optimal solution on the boundary of the search domain. Therefore, the objective function $h(y)$ has 2^n local optimal solutions, and it is considered a multi-modal function with a weak structure. Such an objective function is difficult to optimize using any currently proposed algorithm [11]. Therefore, `WRA` failed to approximate the worst-case ranking. Thus, CMA-ES could not converge to x^* because it optimized for a function that differed significantly from F.

5.4 Experiment 3

To investigate the influence of the coefficient of the interaction term in the objective function, we applied three approaches to f_6–f_8 with $b \in \{1, 3, 10, 30, 100\}$.

Experimental results are shown in Figure 3. We can confirm that `CMA-ES with WRA` could achieve successful convergence in all trials. However, the existing approaches, `Adversarial-CMA-ES` and `ZOPGDA`, failed to converge to the optimal solution in any trials on f_6–f_8. This is as verified in Section 5.3.

For `CMA-ES with WRA` on f_6–f_8, the number of f-calls required for convergence did not significantly change at various b. Even for the results on f_7, where the number of f-calls changed the most, the ratio of the number of f-calls at $b = 1$ and $b = 100$ was approximately two. The results of the existing approaches in Figure 1 suggest that the number of f-calls increased proportionally to $1 + b^2 \approx 10^4$, implying that the factor of two can be considered as small.

6 DISCUSSION ON THE EFFECT OF THE INTERACTION TERM

We discuss the effect of the interaction term of convex–concave problems on the number of required f-calls. As observed in Figure 1 and Figure 3, for $f_5, \ldots, f_8$, the number of f-calls until `CMA-ES with WRA` reaches the target threshold of the worst-case objective function F and scales with the coefficient b of the interaction term $x^T y$ in the order of $O(\log(b))$. Here, we provide an insightful but shallow analysis to describe this scaling.

We limit our attention to f_5 on an unbounded domain. The worst-case objective function is $F(x) = \frac{1+b^2}{2}\|x\|^2$.

Moreover, we assume that CMA-ES converges linearly on an arbitrary convex quadratic objective function, that is, the number of f-calls that CMA-ES performs to reach $\{x : \|x - x^*\| \leqslant \epsilon \cdot \|m^{(0)} - x^*\|\}$ is $O(\log(\|m^{(0)} - x^*\|/\epsilon))$, where x^* is the optimal solution of the objective function. Although no rigorous runtime analysis has been conducted for CMA-ES, a variant of CMA-ES, namely, the (1+1)-ES, exhibited a linear convergence on strongly convex functions with Lipschitz continuous gradients [19]. Moreover, we empirically observe that CMA-ES geometrically approaches the optimal solution.

Under this assumption, if CMA-ES is used to solve the worst-case objective function F directly, the number of the worst-case objective function calls to reach $\{x : \|x - x^*\| \leqslant \epsilon \cdot \|m^{(0)} - x^*\|\}$, where $x^* = \operatorname{argmin}_x F(x)$, is $O(\log(\|m^{(0)} - x^*\|/\epsilon))$. The worst-case objective function is approximated to ensure that Kendall's rank correlation coefficient between the true values $\{F(x_i)\}$ and their approximated value $\{\hat{F}(x_i)\}$ is sufficiently high. Therefore, we expect that CMA-ES with WRA behaves similarly to CMA-ES solving F directly [1]. Therefore, we anticipate that CMA-ES with WRA must approximate $O(\log(\|m^{(0)} - x^*\|/\epsilon))$ worst-case objective function values during the optimization.

We now estimate the number of f-calls for each F approximation. Because $F(x) \propto \|x\|^2$, the covariance matrix of the upper-level CMA-ES is expected to be adapted as $\Sigma \approx \sigma^2 \cdot I$, where $\sigma^2 > 0$ is a scalar. The $\mathcal{L}_2$ distance between the worst-case objective function values of two candidate solutions x_1 and x_2 generated independently from $\mathcal{N}(m, \Sigma)$ is $\mathbb{E}[(F(x_1) - F(x_2))^2]^{1/2} = (1 + b^2)(2\operatorname{Tr}(\Sigma^2))^{1/2} \approx (2d)^{1/2}(1 + b^2)\sigma^2$. The approximated worst-case objective function $\tilde{F}$ is required to have a sufficiently high Kendall's rank correlation coefficient with the true value F under the current search distribution $\mathcal{N}(m, \Sigma)$. Therefore, it is assumed that the comparison $\tilde{F}(x_1) \lesseqgtr \tilde{F}(x_2)$ of two solutions generated from the current distribution provides a true comparison $F(x_1) \lesseqgtr F(x_2)$ with high probability. Therefore, the worst-case objective function values need to be approximated with precision $|F(x) - \tilde{F}(x)| \leqslant c^2 \cdot (1 + b^2) \cdot \sigma^2$ for some constant $c > 0$. Observing that

$$
\begin{aligned}
F(x) - \tilde{F}(x) &= f(x, \hat{y}(x)) - f(x, \tilde{y}) \\
&= b x^{\mathrm{T}}(\hat{y}(x) - \tilde{y}) - (\|\hat{y}(x)\|^2 - \|\tilde{y}\|^2)/2 \\
&= \hat{y}(x)^{\mathrm{T}}(\hat{y}(x) - \tilde{y}) - (\|\hat{y}(x)\|^2 - \|\tilde{y}\|^2)/2 = \|\hat{y}(x) - \tilde{y}\|^2/2,
\end{aligned}
$$

the aforementioned condition reads $\|\hat{y}(x) - \tilde{y}\| \leqslant c \cdot (1 + b^2)^{1/2} \sigma$. The runtime to find such a $\tilde{y}$ using CMA-ES is $O\left(\log\left(\frac{\|y^{(0)} - \hat{y}(x)\|}{c \cdot (1 + b^2)^{1/2} \sigma}\right)\right)$. Notably, $\hat{y}(x) = bx$ and $\tilde{y}^{(0)}$ is a near-optimal solution to $\max_y f(x, y)$ for a solution x generated in a previous iteration. As the distribution parameters of the upper-level CMA-ES do not change rapidly, m and Σ remain from the last iteration. Thus, both $\hat{y}(x)$ and $\tilde{y}^{(0)}$ can be considered $\mathcal{N}(bm, b^2\Sigma)$-distributed. Their expected squared distance is then $b^2 \operatorname{Tr}(\Sigma) \approx db^2\sigma^2$. Therefore, we estimate $\|y^{(0)} - \hat{y}(x)\| \in O(b\sigma)$. Hence, the runtime to find such a $\tilde{y}$ is $O\left(\log\left(\frac{b}{c \cdot (1 + b^2)^{1/2}}\right)\right)$. That is, the number of f-calls required to approximate the worst-case objective function value for each x remains constant order over time.

Generally, we obtain the estimated number of f-calls until CMA-ES with WRA reaches the ϵ-optimal solution to the worst-case objective function as follows:

$$
O\left(\log\left(\frac{b}{c \cdot (1 + b^2)^{1/2}}\right) \cdot \log\left(\frac{\|m^{(0)} - x^*\|}{\epsilon}\right)\right). \tag{4}
$$

For $b \leqslant 1$, the first term scales as $\log(b)$. However, as $b \to \infty$, the first term approaches a constant.

7 CONCLUSION

This study focused on min–max continuous optimization problems whose objective function is a black-box. We addressed the following challenges of the existing approaches, Adversarial-CMA-ES and ZOPGDA. (I) The number of f-calls required to reach convergence depends largely on the interaction term $x^T y$ of the objective function. (II) The objective function in the neighborhood of the optimal solution needs to be a Lipschitz smooth and strongly convex–concave function for convergence.

Our contributions are as follows. (A) We proposed a new approach (CMA-ES with a worst-case ranking approximation: CMA-ES with WRA) to address Difficulty (I) and (II). CMA-ES with WRA works because WRA estimates the ranking of the solution candidates on the worst-case function, and CMA-ES searches for the optimal solution using the estimated ranking information. (B) Numerical experiments on the strongly convex–concave function showed that CMA-ES with WRA improved the scalability of the number of f-calls against the coefficient b multiplied by $x^T y$. The number of f-calls resulting from CMA-ES with WRA scaled to approximately $\log(b)$, whereas that of the existing approaches was b^2. (C) To ensure that CMA-ES with WRA addresses Difficulty (II), we applied CMA-ES with WRA to test problems whose objective function was not limited to being Lipschitz smooth and strongly convex–concave in the neighborhood of the optimal solution. The experimental results showed that only CMA-ES with WRA, among the compared approaches, could converge to the optimal solution. However, CMA-ES with WRA could not converge to the optimal solution when the worst-case ranking could not be estimated properly, for example, when the objective function was a multi-modal function with a weak structure. (D) Additionally, we confirmed that the number of f-calls performed by CMA-ES with WRA was not significantly affected by changing the coefficient b multiplied by $x^T y$ on the objective functions that are not limited to being a Lipschitz smooth and strongly convex–concave function in the neighborhood of the optimal solution.

The limitations of this study are the lack of a theoretical analysis of the proposed approach and empirical evaluation on the scaling of the number of f-calls about the dimension m and n on broader class of functions. Moreover, the successful convergence of the proposed approach was not clearly identified. The sensitivity analysis to the hyper-parameters of WRA, $c_{\max}$ and $\tau_{\mathrm{threshold}}$, is yet to be performed. Future work on the proposed approach should include more theoretical and empirical analyses. Compared with the existing approaches, ZOPGDA and Adversarial-CMA-ES, CMA-ES with WRA requires significantly more f-calls if the objective function is strongly convex-concave and Lipschitz continuous

and the effect of the interaction term, $H_{x,y}$, is relatively small. This is, therefore, a limitation of the proposed approach.

ACKNOWLEDGMENTS

REFERENCES

[1] Y. Akimoto. 2022. Monotone Improvement of Information-Geometric Optimization Algorithms with a Surrogate Function. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '22)*. https://doi.org/10.1145/3512290.3528690

[2] Y. Akimoto and N. Hansen. 2020. Diagonal Acceleration for Covariance Matrix Adaptation Evolution Strategies. *Evolutionary Computation* 28, 3 (2020), 405–435. https://doi.org/10.1162/evco_a_00260

[3] Y. Akimoto, Y. Miyauchi, and A. Maki. 2022. Saddle Point Optimization with Approximate Minimization Oracle and Its Application to Robust Berthing Control. *ACM Trans. Evol. Learn. Optim.* (2022). https://doi.org/10.1145/3510425 Just Accepted.

[4] Y. Akimoto, T. Shimizu, and T. Yamaguchi. 2019. Adaptive Objective Selection for Multi-Fidelity Optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '19)*. 880–888. https://doi.org/10.1145/3321707.3321709

[5] H. J. C. Barbosa. 1999. A coevolutionary genetic algorithm for constrained optimization. In *Proceedings of the Congress on Evolutionary Computation (CEC '99)*, Vol. 3. 1605–161.

[6] D. Bertsimas, O. Nohadani, and K. M. Teo. 2010. Robust Optimization for Unconstrained Simulation-Based Problems. *Operations Research* 58, 1 (2010), 161–178. https://doi.org/10.1287/opre.1090.0715

[7] I. Bogunovic, J. Scarlett, S. Jegelka, and V Cevher. 2018. Adversarially Robust Optimization with Gaussian Processes. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS '18)*. 5765–5775.

[8] C. J. Freitas. 2002. The issue of numerical uncertainty. *Applied Mathematical Modelling* 26, 2 (2002), 237–248. https://doi.org/10.1016/S0307-904X(01)00058-0

[9] N. Hansen. 2019. A Global Surrogate Assisted CMA-ES. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '19)*. 664–672. https://doi.org/10.1145/3321707.3321842

[10] N. Hansen and A. Auger. 2014. *Principled Design of Continuous Stochastic Search: From Theory to Practice.* Springer Berlin Heidelberg, Berlin, Heidelberg, 145–180.

[11] N. Hansen, A. Auger, R. Ros, S. Finck, and P. Pošík. 2010. Comparing Results of 31 Algorithms from the Black-Box Optimization Benchmarking BBOB-2009. In *Proceedings of the 12th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '10)*. 1689–1696. https://doi.org/10.1145/1830761.1830790

[12] N. Hansen and A. Ostermeier. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evol. Comput.* 9, 2 (2001), 159–195. https://doi.org/10.1162/106365601750190398

[13] N. Hansen, R. Ros, N. Mauny, M. Schoenauer, and A. Auger. 2011. Impacts of invariance in search: When CMA-ES and PSO face ill-conditioned and non-separable problems. *Applied Soft Computing* 11, 8 (2011), 5755–5769. https://doi.org/10.1016/j.asoc.2011.03.001

[14] J. W. Herrmann. 1999. A genetic algorithm for minimax optimization problems. In *Proceedings of the Congress on Evolutionary Computation (CEC '99)*, Vol. 2. 1099–1103.

[15] M. Kendall and J. D. Gibbons. 1990. *Rank Correlation Methods* (5th ed.). Oxford University Press.

[16] T. Liang and J. Stokes. 2019. Interaction Matters: A Note on Non-asymptotic Local Convergence of Generative Adversarial Networks. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS '19)*. 907–915.

[17] S. Liu, S. Lu, X. Chen, Y. Feng, K. Xu, A. Al-Dujaili, M. Hong, and U. O'Reilly. 2020. Min-Max Optimization without Gradients: Convergence and Applications to Black-Box Evasion and Poisoning Attacks. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. 6282–6293.

[18] A. Miyagi, K. Fukuchi, J. Sakuma, and Y. Akimoto. 2021. Adaptive Scenario Subset Selection for Min–Max Black-Box Continuous Optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '21)*. 697–705. https://doi.org/10.1145/3449639.3459291

[19] D. Morinaga and Y. Akimoto. 2019. Generalized Drift Analysis in Continuous Domain: Linear Convergence of (1 + 1)-ES on Strongly Convex Functions with Lipschitz Continuous Gradients. In *Proceedings of the 15th ACM/SIGEVO Conference on Foundations of Genetic Algorithms (FOGA '19)*. 13–24. https://doi.org/10.1145/3299904.3340303

[20] X. Qiu, J. Xu, Y. Xu, and K. C. Tan. 2018. A New Differential Evolution Algorithm for Minimax Optimization in Robust Design. *IEEE Transactions on Cybernetics* 48, 5 (2018), 1355–1368. https://doi.org/10.1109/TCYB.2017.2692963

[21] Y. Wang and D. L. McDowell. 2020. 1 - Uncertainty quantification in materials modeling. In *Uncertainty Quantification in Multiscale Materials Modeling*, Y. Wang and D. L. McDowell (Eds.). Woodhead Publishing, 1–40. https://doi.org/10.1016/B978-0-08-102941-1.00001-8

[22] T. Yamaguchi and Y. Akimoto. 2018. A Note on the CMA-ES for Functions with Periodic Variables. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '18)*. 227–228. https://doi.org/10.1145/3205651.3205669