Saddle Point Optimization with Approximate Minimization Oracle and Its Application to Robust Berthing Control

YOUHEI AKIMOTO, Faculty of Engineering, Information and Systems, University of Tsukuba & RIKEN Center for Advanced Intelligence Project, Japan
YOSHIKI MIYAUCHI and **ATSUO MAKI**, Department of Naval Architecture and Ocean Engineering, Graduate School of Engineering, Osaka University, Japan

2

We propose an approach to saddle point optimization relying only on oracles that solve minimization problems approximately. We analyze its convergence property on a strongly convex-concave problem and show its linear convergence toward the global min-max saddle point. Based on the convergence analysis, we develop a heuristic approach to adapt the learning rate. An implementation of the developed approach using the (1+1)-CMA-ES as the minimization oracle, namely, Adversarial-CMA-ES, is shown to outperform several existing approaches on test problems. Numerical evaluation confirms the tightness of the theoretical convergence rate bound as well as the efficiency of the learning rate adaptation mechanism. As an example of real-world problems, the suggested optimization method is applied to automatic berthing control problems under model uncertainties, showing its usefulness in obtaining solutions robust to uncertainty.

CCS Concepts: • **Mathematics of computing** → **Continuous optimization**; • **Theory of computation** → **Convergence and learning in games**; **Theory of randomized search heuristics**;

Additional Key Words and Phrases: Minimax optimization, saddle point optimization, robust optimization, robust control, reliability, zero-order approach, convergence analysis, automatic berthing

ACM Reference format:

Youhei Akimoto, Yoshiki Miyauchi, and Atsuo Maki. 2022. Saddle Point Optimization with Approximate Minimization Oracle and Its Application to Robust Berthing Control. *ACM Trans. Evol. Learn. Optim* 2, 1, Article 2 (April 2022), 32 pages.

<https://doi.org/10.1145/3510425>

1 INTRODUCTION

Simulation-based optimization has recently received increasing attention from researchers. Here, the objective function $h : \mathbb{X} \rightarrow \mathbb{R}$, where $\mathbb{X} \subseteq \mathbb{R}^m$, is not explicitly given, but its value for each $x \in \mathbb{X}$ can be computed through computational simulation. Solvers for simulation-based optimization problems have been widely developed. While some are domain-specific, others are

This work is partially supported by JSPS KAKENHI Grants No. 19H04179 and No. 19K04858.

Authors' addresses: Y. Akimoto, Faculty of Engineering, Information and Systems, University of Tsukuba & RIKEN Center for Advanced Intelligence Project, 1-1-1 Tennodai, Tsukuba, Ibaraki, Japan, 305-8573; email: akimoto@cs.tsukuba.ac.jp; Y. Miyauchi and A. Maki, Department of Naval Architecture and Ocean Engineering, Graduate School of Engineering, Osaka University, 2-1 Yamadaoka, Suita, Osaka, 565-0971, Japan; emails: yoshiki_miyauchi@naoe.eng.osaka-u.ac.jp, maki@naoe.eng.osaka-u.ac.jp.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

2688-3007/2022/04-ART2 \$15.00

<https://doi.org/10.1145/3510425>

general-purpose solvers. For a case where simulation-based optimization is required, we first need to design a simulator that models reality, for example, a physical equation, and computes the objective function value for each solution. Then, we apply a numerical solver to solve $\operatorname{argmin}_{x \in \mathbb{X}} h(x)$. However, owing to modeling errors and uncertainties, the optimal solution to $\operatorname{argmin}_{x \in \mathbb{X}} h(x)$ computed through a simulator is not necessarily optimal in real environments in which the obtained solution is used. This issue threatens the reliability of solutions obtained through simulation-based optimization.

An approach to obtain a solution that is robust against modeling errors and uncertainty is to formulate the problem as a min–max optimization,

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

where $y \in \mathbb{Y}$ represents the model parameters and the uncertain parameters. In the following, y is referred to as the *uncertainty parameter*. Assume that the real environment is represented by $y_{\text{real}} \in \mathbb{Y}$. The original objective $h(x)$ is equivalent to $f(x, y_{\text{est}})$ with an estimated parameter $y_{\text{est}} \in \mathbb{Y}$. Then, the solution $x_{y_{\text{est}}} = \operatorname{argmin}_{x \in \mathbb{X}} f(x, y_{\text{est}})$ obtained via simulation does not guarantee good performance in the real environment. That is, $f(x_{y_{\text{est}}}, y_{\text{real}})$ may be arbitrarily greater than $f(x_{y_{\text{est}}}, y_{\text{est}})$. In contrast, the solution $x_{\mathbb{Y}} = \operatorname{argmin}_{x \in \mathbb{X}} \max_{y \in \mathbb{Y}} f(x, y)$ to (1) guarantees that $f(x_{\mathbb{Y}}, y_{\text{real}}) \leq \max_{y \in \mathbb{Y}} f(x_{\mathbb{Y}}, y)$. That is, by minimizing the worst-case objective value, one can guarantee performance in the real environment provided that $y_{\text{real}} \in \mathbb{Y}$.

Robust Berthing Control. As an important real-world application of the min–max optimization (1), we consider an automatic ship berthing task [Maki et al. 2021, 2020], which can be formulated as an optimization of the feedback controller of a ship. Currently, the domestic shipping industry in Japan is facing a shortage of experienced on-board officers. Moreover, the existing workforce of officers is aging [Ministry of Land and Tourism 2020]. This has generated considerable interest in autonomous ship operation to improve maritime safety, shipboard working environments, and productivity, and the technology is being actively developed. Automatic berthing/docking requires fine control so that the ship can reach the target position located near the berth but avoid colliding with it. Therefore, automatic berthing is central to the realization of automatic ship operations. Because it is difficult to train the controller in a real environment owing to cost and safety issues, a typical approach first models the state equation of a ship, for example, using system identification techniques [Abkowitz 1980; Araki et al. 2012; Miyauchi et al. 2021; Wakita et al. 2022] and then optimizes the feedback controller in a simulator. However, such an approach always suffers from modeling errors and uncertainties. For instance, the coefficients of a state equation model are often estimated based on captive model tests in towing tanks and regressions; hence, they may include errors. Moreover, the weather conditions at the time of operation could differ from those considered in the model. Optimization of the feedback controller on a simulator with an estimated model may result in a catastrophic accident, such as collision with the berth. Thus, to design a berthing control solution robust against modeling errors and uncertainties, we formulate the problem as a min–max optimization (1), where x is the parameter of the feedback controller and y is the parameter representing the coefficients of the state equation model and weather conditions.

Saddle Point Optimization. Here, we consider min–max continuous optimization (1), where $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ is the objective function and $\mathbb{X} \times \mathbb{Y} \subseteq \mathbb{R}^m \times \mathbb{R}^n$ is the search domain. In addition to the abovementioned situation, min–max optimization can be applied in many fields of engineering, including robust design [Conn and Vicente 2012; Qiu et al. 2018], robust control [Pinto et al. 2017; Shioya et al. 2018], constrained optimization [Cherukuri et al. 2017], and generative adversarial networks (GANs) [Goodfellow et al. 2014; Salimans et al. 2016]. In particular, we are interested in the min–max optimization of a *derivative-free* and *black-box*

objective f , where the gradient or higher-order information is unavailable (derivative-free) and no special structures such as convexity or the Lipschitz constant are available in advance (black-box) [Frazier 2018].¹

We aim to locate a local min–max saddle point of f , that is, a point (x^*, y^*) satisfying $f(x, y^*) \geq f(x^*, y^*) \geq f(x^*, y)$ in a neighborhood of (x^*, y^*) . Generally, it is difficult to locate the global minimum of the worst-case objective $F(x) := \max_{y \in \mathbb{Y}} f(x, y)$. In a non-convex optimization context, the goal is often to locate a local minimum of an objective rather than the global minimum as a realistic target. However, in the min–max optimization context, it is still difficult to locate a local minimum of the worst-case objective $F(x)$, because doing so requires the maximization itself and there may exist local maxima of $f(x, y)$ unless $f(x, y)$ is concave in y for all x . A local min–max saddle point is considered as a local optimal solution in the min–max optimization context, because it is a local minimum in x and a local maximum in y . Therefore, as a practical target, we focus on locating the local min–max saddle point of (1).

Related Works. First-order approaches are often employed for (1) if gradients are available. A simultaneous gradient descent-ascent (GDA) approach,

$$(x_{t+1}, y_{t+1}) = (x_t, y_t) + \eta(-\nabla_x f(x_t, y_t), \nabla_y f(x_t, y_t)), \quad (2)$$

has often been analyzed for its local and global convergence properties on twice continuously differentiable functions owing to its simplicity and popularity. A condition on the learning rate $\eta > 0$ for the dynamics (2) to be asymptotically stable at a local min–max saddle point has been studied [Mescheder et al. 2017; Nagarajan and Kolter 2017]. Subsequently, Adolphs et al. [2019] showed the existence of asymptotically stable points of (2) that are not local min–max saddle points. Liang and Stokes [2019] have derived a sufficient condition on η for (2) to converge toward the global min–max saddle point on a locally strongly convex–concave function. Frank-Wolfe type approaches have also been analyzed for constrained situations [Gidel et al. 2017; Nouiehed et al. 2019]. Although a convergence guarantee was not provided, [Bertsimas et al. 2010a, b] have proposed a first-order approach targeting on f that is non-concave in y .

Zero-order approaches for (1) include coevolutionary approaches [Al-Dujaili et al. 2019; Branke and Rosenbusch 2008; Jensen 2004; Qiu et al. 2018; Zhou and Zhang 2010], surrogate-model-based approaches [Bogunovic et al. 2018; Conn and Vicente 2012; Picheny et al. 2019], and gradient approximation approaches [Liu et al. 2020]. Compared to first-order approaches, zero-order approaches have not been thoroughly analyzed in terms of their convergence guarantees and convergence rates. In particular, coevolutionary approaches are often designed heuristically and without convergence guarantees. Indeed, they fail to converge toward a min–max saddle point even on strongly convex–concave problems, as has been reported by Akimoto [2021] and as noted below in the experimental results. Recently, Bogunovic et al. [2018] showed regret bounds for a Bayesian optimization approach and Liu et al. [2020] showed an error bound for a gradient approximation approach, where the error is measured by the square norm of the gradient. Both analyses show sublinear rates under possibly stochastic (i.e., noisy) versions of (1). However, compared to the first-order approach, which exhibits linear convergence, they show slower convergence.

Contributions. We propose an approach to saddle point optimization (1) that relies solely on numerical solvers that approximately solve $\operatorname{argmin}_{x' \in \mathbb{X}} f(x', y)$ for each $y \in \mathbb{Y}$ and $\operatorname{argmin}_{y' \in \mathbb{Y}} -f(x, y')$ for each $x \in \mathbb{X}$. Given an initial solution $(x_0, y_0) \in \mathbb{X} \times \mathbb{Y}$, our approach

¹In this article, *simulation-based optimization* is used to refer to the problems described in the first paragraph above. The terms *derivative-free* and *black-box* are used to refer to the characteristics of the objective function.

iteratively locates the approximate solutions $\tilde{x}_t \approx \operatorname{argmin}_{x' \in \mathbb{X}} f(x', y_t)$ and $\tilde{y}_t \approx \operatorname{argmin}_{y' \in \mathbb{Y}} -f(x_t, y')$ and updates the solution as

$$(x_{t+1}, y_{t+1}) = (x_t, y_t) + \eta \cdot (\tilde{x}_t - x_t, \tilde{y}_t - y_t), \quad (3)$$

where $\eta > 0$ is the learning rate. This approach takes inspiration from the GDA method (2), where we replace $-\nabla_x f(x_t, y_t)$ and $\nabla_y f(x_t, y_t)$ with $\tilde{x}_t - x_t$ and $\tilde{y}_t - y_t$. However, unlike the GDA approach, the solvers need not be gradient-based. This is advantageous in the following situations: (1) there exists a well-developed numerical solver suitable for $\operatorname{argmin}_{x' \in \mathbb{X}} f(x', y)$ and/or $\operatorname{argmin}_{y' \in \mathbb{Y}} -f(x, y')$; (2) derivative-free approaches such as the covariance matrix adaptation evolution strategy (CMA-ES) [Akimoto and Hansen 2020; Hansen and Auger 2014; Hansen et al. 2003; Hansen and Ostermeier 2001] are sought, because gradient information is not available or gradient-based approaches are known to be sensitive to their initial search points.

We analyze the proposed approach on strongly convex-concave problems, and prove its linear convergence in terms of the number of numerical solver calls. In particular, we provide an upper bound on η to guarantee linear convergence toward the global min-max saddle point and the convergence rate bound. This corresponds to the known result for the GDA approach (2). Compared to existing derivative-free approaches for saddle point optimization, this result is unique in that our convergence is linear, while the existing results show sublinear convergence [Bogunovic et al. 2018; Liu et al. 2020]. Although our motivational application is not necessarily a strongly convex-concave problem, the quantitative analysis helps to understand limitations of the approach—need for η adaptation—and provide inspiration on how to improve the approach.

Moreover, we developed a heuristic adaptation mechanism for the learning rate in a black-box optimization setting. In the black-box setting, we do not know in advance the characteristic constants of a problem that determines the upper bound for the learning rate to guarantee convergence. Therefore, a learning rate adaptation mechanism is highly desired to avoid trial and error in tuning the learning rate. We implemented two variants of the proposed approach, one using (1+1)-CMA-ES [Arnold and Hansen 2010], a zero-order approach, as the minimization solver, and another using sequential least squares quadratic programming (SLSQP) [Kraft 1988], a first-order approach. Empirical studies on test problems show that the learning rate adaptation achieved performance competitive to the proposed approach with the optimal static learning rate, while obviating the need for time-consuming parameter tuning. We also demonstrate the limitations of existing coevolutionary approaches as well as the proposed approach.

We apply our approach to robust berthing control optimization, as an example of a real-world application with a non-convex-concave objective. We consider the wind conditions and the coefficients of the state equation for the wind force as the uncertainty parameter y . Some related works address the wind force as an external disturbance when planning the trajectories [Miyauchi et al. 2022]; however, they treat the wind condition as an observable disturbance, and the control signal is selected according to the observed wind condition. In contrast, we optimize the on-line feedback controller under wind disturbance without considering the wind condition as an input to the controller. Moreover, among other studies on automatic berthing control, to the best of our knowledge, the present work is the first to address model uncertainty. Compared to a naive baseline approach, the proposed approach located solutions with better worst-case performance.

This article is an extension of a previous work [Akimoto 2021]. We have improved on the previous work in the following respects. First, we improved the convergence analysis in Section 3.2. We have removed unnecessary assumptions on the problem by refining the proof. Second, we have incorporated the covariance matrix adaptation into our proposed approach in Section 4.3. Third, we have implemented a restart strategy and other ingenuity for practical use, summarized in

Section 4.2. Fourth, we have extended the comparison with existing approaches in Section 5.2. Finally, we have evaluated the usefulness of the proposed approach in a real-world application in Section 6.

Our implementation of the proposed approach in the Python programming language, Adversarial-CMA-ES, is publicly available at GitHub Gist.²

Notation. For a twice continuously differentiable function $f : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}$, that is, $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$, let $H_{x,x}(x, y)$, $H_{x,y}(x, y)$, $H_{y,x}(x, y)$, and $H_{y,y}(x, y)$ be the blocks of the Hessian matrix $\nabla^2 f(x, y) = \begin{bmatrix} H_{x,x}(x, y) & H_{x,y}(x, y) \\ H_{y,x}(x, y) & H_{y,y}(x, y) \end{bmatrix}$ of f , whose (i, j) th components are $\frac{\partial^2 f}{\partial x_i \partial x_j}$, $\frac{\partial^2 f}{\partial x_i \partial y_j}$, $\frac{\partial^2 f}{\partial y_i \partial x_j}$, and $\frac{\partial^2 f}{\partial y_i \partial y_j}$, respectively, evaluated at a given point (x, y) .

For symmetric matrices A and B , by $A \geq B$ and $A > B$, we mean that $A - B$ is non-negative and positive definite, respectively. For simplicity, we write $A \geq a$ and $A > a$ for $a \in \mathbb{R}$ to mean $A \geq a \cdot I$ and $A > a \cdot I$, respectively. For a positive definite symmetric matrix A , let $\sqrt{A}$ denote the matrix square root, that is, $\sqrt{A}$ is a positive definite symmetric matrix such that $A = \sqrt{A} \cdot \sqrt{A}$. Let $\|z\|_A = [z^T A z]^{1/2}$ for a positive definite symmetric A .

Let $J_g(z)$ denote the Jacobian of a differentiable $g = (g_1, \dots, g_k) : \mathbb{R}^\ell \rightarrow \mathbb{R}^k$, where the (i, j) th element is $\partial g_i / \partial z_j$ evaluated at $z = (z_1, \dots, z_\ell) \in \mathbb{R}^\ell$. If $k = 1$, then we write $J_g(z) = \nabla g(z)^T$.

2 SADDLE POINT OPTIMIZATION

Our objective is to locate the global or local min-max saddle point of the min-max optimization problem (1). In the following, we first define the min-max saddle point. We introduce the notion of the suboptimality error to measure the progress toward the global min-max saddle point. Finally, we introduce a strongly convex-concave function as an important class of the objective function, on which we performed convergence analysis, which is described in the next section.

2.1 Min-Max Saddle Point

The min-max saddle point of a function $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ is defined as follows. Here, $\mathcal{E}_z \subseteq \mathbb{R}^\ell$ is called a *neighborhood* of $z^* \in \mathbb{R}^\ell$ if there exists an open ball $B(z^*, r) = \{z \in \mathbb{R}^\ell : \|z - z^*\| < r\}$ such that $B(z^*, r) \subseteq \mathcal{E}_z$. A point (x, y) is called a *critical point* if $\nabla f(x, y) = (\nabla_x f(x, y), \nabla_y f(x, y)) = 0$.

Definition 2.1 (Min-Max Saddle Point). A point $(x^*, y^*) \in \mathbb{X} \times \mathbb{Y}$ is a *local min-max saddle point* of a function $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ if there exists a neighborhood $\mathcal{E}_x \times \mathcal{E}_y \subseteq \mathbb{X} \times \mathbb{Y}$ including (x^*, y^*) such that for any $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$, the condition $f(x, y^*) \geq f(x^*, y^*) \geq f(x^*, y)$ holds. If $\mathcal{E}_x = \mathbb{X}$ and $\mathcal{E}_y = \mathbb{Y}$, then the point (x^*, y^*) is called the *global min-max saddle point*. If the equality holds only if $(x, y) = (x^*, y^*)$, then it is called a *strict min-max saddle point*.

For twice continuously differentiable function $f \in C^2(\mathbb{X} \times \mathbb{Y}, \mathbb{R})$, a point (x^*, y^*) is a strict local min-max saddle point if it is a critical point and $H_{x,x}(x^*, y^*) > 0$ and $H_{y,y}(x^*, y^*) < 0$ hold. In general, the opposite does not hold. For example, a local min-max saddle point can be a boundary point of $\mathbb{X} \times \mathbb{Y}$ that is not a critical point.

We comment on the relation between the min-max saddle point and the solutions to the worst-case objective function $F(x) := \max_{y \in \mathbb{Y}} f(x, y)$. If there exists a global min-max saddle point (x^*, y^*) of f , then x^* is one of the global minimal point of the worst-case objective function F , and we have $F(x^*) = f(x^*, y^*)$. However, a global minimal point $\bar{x} \in \mathbb{X}$ of $F(x)$ and its corresponding worst uncertainty parameter $\bar{y} \in \arg\max_{y \in \mathbb{Y}} f(\bar{x}, y)$ do not necessarily form a global min-max saddle point in general. An example case is $f(x, y) = (x + \sin(\pi y))^2$, where the worst-case objective

²<https://gist.github.com/youheiakimoto/ab51e88c73baf68effd95b750100aad0>.

function is $F(x) = (|x| + 1)^2$ and its global minimal point is $\bar{x} = 0$. The corresponding worst uncertainty parameters are $\frac{1}{2} + i$ for all integer i , but $(\bar{x}, \bar{y})$ is not even a local min-max saddle point for any $\bar{y} = \frac{1}{2} + i$. Moreover, if (x^*, y^*) is a local min-max saddle point of f , the point x^* is not necessarily a local minimal point of F .

2.2 Suboptimality Error

The suboptimality error [Gidel et al. 2017] is a quantity that measures the progress toward the global min-max saddle point, defined as follows.

Definition 2.2 (Suboptimality Error). For function $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$, the suboptimality error $G_x : \mathbb{X} \times \mathbb{Y} \rightarrow [0, \infty)$ in x and the suboptimality error $G_y : \mathbb{X} \times \mathbb{Y} \rightarrow [0, \infty)$ in y are defined as

$$G_x(x, y) := f(x, y) - \min_{x' \in \mathbb{X}} f(x', y), \quad (4)$$

$$G_y(x, y) := \max_{y' \in \mathbb{Y}} f(x, y') - f(x, y), \quad (5)$$

and the *suboptimality error* is

$$G(x, y) := G_x(x, y) + G_y(x, y) = \max_{y' \in \mathbb{Y}} f(x, y') - \min_{x' \in \mathbb{X}} f(x', y). \quad (6)$$

The suboptimality error $G(x, y)$ is zero if and only if (x, y) is the global min-max saddle point of f . Moreover, the local min-max saddle points of f are characterized by suboptimality errors. This is summarized in the following proposition, whose proof is given in Appendix A.

PROPOSITION 2.3. *The point (x^*, y^*) is the global min-max saddle point of f if and only if it is the global minimal point of G , that is, $G(x, y) \geq 0$ for any $(x, y) \in \mathbb{X} \times \mathbb{Y}$. The point (x^*, y^*) is a local min-max saddle point of f if and only if x^* and y^* are local minimal points of $G_x(\cdot, y^*)$ and $G_y(x^*, \cdot)$, respectively, that is, there exists a neighborhood $\mathcal{E}_x \times \mathcal{E}_y$ of (x^*, y^*) such that $G_x(x, y^*) \geq G_x(x^*, y^*)$ and $G_y(x^*, y) \geq G_y(x^*, y^*)$ for any $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$.*

2.3 Strongly Convex-Concave Function

A strongly convex-concave function is often considered for the theoretical investigation of first-order min-max optimization approaches.

Definition 2.4. A twice continuously differentiable function $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$ is *locally μ -strongly convex-concave* around a critical point (x^*, y^*) for some $\mu > 0$ if there exist open sets $\mathcal{E}_x \subseteq \mathbb{R}^m$ including x^* and $\mathcal{E}_y \subseteq \mathbb{R}^n$ including y^* such that $H_{x,x}(x, y) \geq \mu$ and $-H_{y,y}(x, y) \geq \mu$ for all $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$. The function f is *globally μ -strongly convex-concave* if $\mathcal{E}_x = \mathbb{R}^m$ and $\mathcal{E}_y = \mathbb{R}^n$. We say that f is *locally or globally strongly convex-concave* if f is locally or globally μ -strongly convex-concave for some $\mu > 0$.

If the objective function f is a globally strongly convex-concave, then the global minimal point of the worst-case objective function $F(x)$ is the global min-max saddle point, and it is the only local min-max saddle point.

The implicit function theorem, for example, Theorem 5 of de Oliveira [2013], provides important characteristics of strongly convex-concave functions.

PROPOSITION 2.5 (IMPLICIT FUNCTION THEOREM). *Let (x^*, y^*) be a min-max saddle point of $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$ and f be (at least) locally strongly convex-concave around (x^*, y^*) in $\mathcal{E}_x \times \mathcal{E}_y \subseteq \mathbb{R}^m \times \mathbb{R}^n$.*

There exist open sets $\mathcal{E}_{x,x} \subseteq \mathcal{E}_x$ including x^ and $\mathcal{E}_{x,y} \subseteq \mathcal{E}_y$ including y^* , such that there exists a unique $\hat{y} : \mathcal{E}_{x,x} \rightarrow \mathcal{E}_{x,y}$ such that $\nabla_y f(x, \hat{y}(x)) = 0$. Moreover, $y^* = \hat{y}(x^*)$ and $J_{\hat{y}}(x) = -(H_{y,y}(x, \hat{y}(x)))^{-1} H_{y,x}(x, \hat{y}(x))$ for all $x \in \mathcal{E}_{x,x}$.*

Analogously, there exist open sets $\mathcal{E}_{y,y} \subseteq \mathcal{E}_y$ including y^* and $\mathcal{E}_{y,x} \subseteq \mathcal{E}_x$ including x^* , such that there exists a unique $\hat{x} : \mathcal{E}_{y,y} \rightarrow \mathcal{E}_{y,x}$ such that $\nabla_x f(\hat{x}(y), y) = 0$. Moreover, $x^* = \hat{x}(y^*)$ and $J_{\hat{x}}(y) = -(H_{x,x}(\hat{x}(y), y))^{-1} H_{x,y}(\hat{x}(y), y)$ for all $y \in \mathcal{E}_{y,y}$.

If f is globally strongly convex-concave, then one can take $\mathcal{E}_{x,x} = \mathcal{E}_{y,x} = \mathbb{R}^m$ and $\mathcal{E}_{y,y} = \mathcal{E}_{x,y} = \mathbb{R}^n$ in the above statements.

Proposition 2.5 states that for a globally strongly convex-concave $f \in C^2(\mathbb{X} \times \mathbb{Y}, \mathbb{R})$, for each $x \in \mathbb{R}^m$ there exists a unique global maximal point $\hat{y}(x) \in \mathbb{R}^n$ such that $\hat{y}(x) = \operatorname{argmax}_{y \in \mathbb{R}^n} f(x, y)$, and for each $y \in \mathbb{R}^n$ there exists a unique global minimal point $\hat{x}(y) \in \mathbb{R}^m$ such that $\hat{x}(y) = \operatorname{argmin}_{x \in \mathbb{R}^m} f(x, y)$.

The following lemma shows the positivity of the Hessian of the suboptimality error G , which implies that the suboptimality error G is a globally strongly convex function. The proof is provided in Appendix A.

LEMMA 2.6. Suppose that $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$ is globally μ -strongly convex-concave for some $\mu > 0$. The Hessian matrix of the suboptimality error G is $\nabla^2 G(x, y) = \operatorname{diag}(G_{x,x}(x, \hat{y}(x)), G_{y,y}(\hat{x}(y), y))$, where

$$\begin{aligned} G_{x,x}(x, y) &= H_{x,x}(x, y) + H_{x,y}(x, y)(-H_{y,y}(x, y))^{-1}H_{y,x}(x, y), \\ G_{y,y}(x, y) &= -H_{y,y}(x, y) + H_{y,x}(x, y)(H_{x,x}(x, y))^{-1}H_{x,y}(x, y), \end{aligned}$$

and they are symmetric, and $G_{x,x}(x, y) \geq \mu$ and $G_{y,y}(x, y) \geq \mu$.

3 ORACLE-BASED SADDLE POINT OPTIMIZATION

We now analyze saddle point optimization based on the approximate minimization oracle outlined in (3). In the following, we formally state the condition for the approximate minimization oracle. Then, we show the global convergence of (3) on strongly convex-concave functions.

3.1 Approximate Minimization Oracle

First, we formally define the requirement for the minimization problem solvers.

Definition 3.1 (Approximate Minimization Oracle). Given an objective function $h : \mathbb{Z} \rightarrow \mathbb{R}$ to be minimized and a reference solution $\bar{z} \in \mathbb{Z}$, an *approximate minimization oracle* $\mathcal{M}$ outputs a solution $\tilde{z} = \mathcal{M}(h, \bar{z})$ satisfying $h(\tilde{z}) < h(\bar{z})$ unless $\bar{z}$ is a local minimal point of h .

We now reformulate the saddle point optimization with approximate minimization oracles. Suppose that we have an approximate minimization oracle $\mathcal{M}_x$ solving $\operatorname{argmin}_{x' \in \mathbb{X}} f(x', y)$ for any $y \in \mathbb{Y}$ and an approximate minimization oracle $\mathcal{M}_y$ solving $\operatorname{argmin}_{y' \in \mathbb{Y}} -f(x, y')$ for any $x \in \mathbb{X}$. At each iteration, the algorithm asks the approximate minimization oracles to output the approximate solutions to $\operatorname{argmin}_{x' \in \mathbb{X}} f(x', y_t)$ and $\operatorname{argmin}_{y' \in \mathbb{Y}} f(x_t, y')$ with the current solution (x_t, y_t) as their reference point. Let $\tilde{x}_t = \mathcal{M}_x(f(\cdot, y_t), x_t)$ and $\tilde{y}_t = \mathcal{M}_y(-f(x_t, \cdot), y_t)$. The update follows

$$\begin{aligned} x_{t+1} &= x_t + \eta \cdot (\tilde{x}_t - x_t), \\ y_{t+1} &= y_t + \eta \cdot (\tilde{y}_t - y_t). \end{aligned} \tag{7}$$

A point $(x, y) \in \mathbb{X} \times \mathbb{Y}$ is a stationary point of the dynamics of (7) only if it is a local min-max saddle point of f . Moreover, if (x, y) is a strict local min-max saddle point of f , it is a stationary point of the dynamics of (7). Therefore, if it converges, then the final solution is guaranteed to be a local min-max saddle point of f . To guarantee its convergence, we further assume the following requirement.

ASSUMPTION 3.2. Given an objective function $h : \mathbb{Z} \rightarrow \mathbb{R}$ to be minimized and a reference solution $\bar{z} \in \mathbb{Z}$, an approximate minimization oracle $\mathcal{M}$ with an approximation precision parameter $\epsilon \in [0, 1)$ outputs a solution $\tilde{z} = \mathcal{M}(h, \bar{z})$, satisfying

$$h(\tilde{z}) - \min_{z \in \mathbb{Z}} h(z) \leq \epsilon \cdot (h(\bar{z}) - \min_{z \in \mathbb{Z}} h(z)). \quad (8)$$

We are particularly interested in algorithms that decrease the objective function value at a geometric rate on (at least locally) strongly convex objective h as instances of the approximate minimization oracle $\mathcal{M}$. That is, the runtime—number of h calls or ∇h calls—to decrease the objective function difference $h(z) - h(z^*)$ from a local minimum by the factor ϵ is $O(\log(1/\epsilon))$. For example, the gradient descent method is well known to exhibit a geometric decrease in the objective function value on strongly convex functions with Lipschitz continuous gradients [Boyd and Vandenberghe 2004; Karimi et al. 2016]. The (1+1)-ES also exhibits a geometric decrease on strongly convex functions with Lipschitz continuous gradients [Morinaga and Akimoto 2019]. We can satisfy the oracle requirement (8) by performing $O(\log(1/\epsilon))$ iterations of such algorithms. The condition can also be satisfied by algorithms that exhibit slower convergence, that is, sublinear convergence. However, for such algorithms, the runtime increases as a candidate solution becomes closer to a local optimum. Therefore, the stopping condition for the internal algorithm to satisfy (8) needs to be carefully designed.

3.2 Analysis on Strongly Convex-Concave Functions

Next, we investigate the convergence property of the oracle-based saddle point optimization (7) on strongly convex-concave functions. In particular, we are interested in knowing how small the learning rate η needs to be to guarantee convergence and how fast it converges. The following theorem provides an upper bound of the suboptimality error at iteration $t + 1$ given the solution at iteration t . The proof is provided in Appendix A.

THEOREM 3.3. Suppose that $f \in C^2(\mathbb{R}^m \times \mathbb{R}^n, \mathbb{R})$ is globally strongly convex-concave, and there exist $\beta_G \geq \alpha_G > 0$ and $\beta_H \geq \alpha_H > 0$ such that

- (1) $\beta_H \geq \sqrt{H_{x,x}^*}^{-1} H_{x,x}(x, y) \sqrt{H_{x,x}^*}^{-1} \geq \alpha_H$,
- (2) $\beta_H \geq \sqrt{-H_{y,y}^*}^{-1} (-H_{y,y}(x, y)) \sqrt{-H_{y,y}^*}^{-1} \geq \alpha_H$,
- (3) $\beta_G \geq \sqrt{H_{x,x}^*}^{-1} G_{x,x}(x, y) \sqrt{H_{x,x}^*}^{-1} \geq \alpha_G$,
- (4) $\beta_G \geq \sqrt{-H_{y,y}^*}^{-1} G_{y,y}(x, y) \sqrt{-H_{y,y}^*}^{-1} \geq \alpha_G$,

where $H_{x,x}^* = H_{x,x}(x^*, y^*)$, $H_{y,y}^* = H_{y,y}(x^*, y^*)$ and (x^*, y^*) is the global min-max saddle point of f . Consider approach (7) with approximate minimization oracles $\mathcal{M}_x$ and $\mathcal{M}_y$ satisfying Assumption 3.2 with approximate precision $\epsilon < \frac{\alpha_H^5}{\beta_H^4 \beta_G}$. Let

$$\eta^* = \frac{\alpha_H}{\beta_H} \frac{\alpha_H}{\beta_G} \frac{1 - (\beta_H/\alpha_H)^2 \sqrt{(\beta_G/\alpha_H) \cdot \epsilon}}{(1 + \sqrt{\epsilon})^2}, \quad (9)$$

$$\gamma = -2\eta \frac{\alpha_H}{\beta_H} \left(1 - \frac{\beta_H^2}{\alpha_H^2} \sqrt{\frac{\beta_G}{\alpha_H} \cdot \epsilon} \right) + \eta^2 \cdot (1 + \sqrt{\epsilon})^2 \cdot \frac{\beta_G}{\alpha_H}. \quad (10)$$

Then, for any $\eta < 2 \cdot \eta^*$, we have $\gamma < 0$ and $\log(G(x_{t+1}, y_{t+1})) - \log(G(x_t, y_t)) < \gamma$. In other words, the runtime T_ζ to reach $\{(x, y) \in \mathbb{R}^m \times \mathbb{R}^n : G(x, y) \leq \zeta \cdot G(x_0, y_0)\}$ for $\zeta \in (0, 1)$ is $T_\zeta \leq \lceil \frac{1}{|\gamma|} \log(\frac{1}{\zeta}) \rceil$

for any initial point $(x_0, y_0) \in \mathbb{R}^m \times \mathbb{R}^n$. Moreover, $G(x_{t+1}, y_{t+1}) > G(x_t, y_t)$ for all (x_t, y_t) if $\eta > 2 \cdot \bar{\eta}$, where

$$\bar{\eta} = \frac{\beta_H}{\alpha_G} \frac{\beta_H}{\alpha_H} \frac{1 + \sqrt{(\beta_G/\alpha_H) \cdot \epsilon}}{(1 - \sqrt{(\beta_H/\alpha_H) \cdot \epsilon})^2}. \quad (11)$$

Linear Convergence. The proposed approach (7) satisfying Assumption 3.2 converges linearly toward the global min-max saddle point on a strongly convex-concave objective function if $\eta < 2\eta^*$. If $\mathcal{M}_x$ and $\mathcal{M}_y$ are implemented with algorithms that exhibit linear convergence, then we can conclude that the runtime in terms of f -calls and/or ∇f -calls is

$$O\left(\frac{1}{|\gamma|} \log\left(\frac{1}{\zeta}\right) \log\left(\frac{1}{\epsilon}\right)\right). \quad (12)$$

Necessary Condition. To exhibit convergence, shrinking the learning rate η is not only sufficient but also necessary. To determine the closeness of the upper bound $2 \cdot \eta^*$ in the sufficient condition and the lower bound $2 \cdot \bar{\eta}$ in the necessary condition, consider a convex-concave quadratic function $f(x, y) = \frac{a}{2}x^2 + bxy - \frac{c}{2}y^2$ for instance, where $a > 0, b \in \mathbb{R}$ and $c > 0$. Then, we have $\alpha_H = \beta_H = 1$ and $\alpha_G = \beta_G = \frac{ac}{ac+b^2}$. Ignoring the effect of ϵ , we have $\eta^* = \bar{\eta} = \frac{ac}{ac+b^2}$. This implies that the sufficient condition for linear convergence, $\eta < 2 \cdot \eta^*$, is indeed the necessary condition for the convergence itself in this example situation. This reveals a limitation of existing approaches [Al-Dujaili et al. 2019; Pinto et al. 2017], which corresponds to (7) with $\eta = 1$.

Runtime Bound. The runtime bound T_ζ is proportional to $\frac{1}{|\gamma|}$ in (10). The log-progress bound $|\gamma|$ is roughly proportional to $2 \cdot \eta$ if $\eta \ll 1$. That is, the runtime is proportional to $\frac{1}{2 \cdot \eta}$. The minimal runtime bound is obtained when $\eta = \eta^*$, where

$$\gamma = \gamma^* := -\frac{\alpha_H}{\beta_G} \left(\frac{\alpha_H}{\beta_H}\right)^2 \left(\frac{1 - (\beta_H/\alpha_H)^2 \sqrt{(\beta_G/\alpha_H) \cdot \epsilon}}{1 + \sqrt{\epsilon}}\right)^2. \quad (13)$$

Provided that $\epsilon \ll 1$, we have $\eta^* \approx \frac{\alpha_H}{\beta_G} \frac{\alpha_H}{\beta_H}$ and $\gamma^* \approx -\frac{\alpha_H}{\beta_G} \left(\frac{\alpha_H}{\beta_H}\right)^2$. The main factor that limits η^* and γ^* is $\frac{\alpha_H}{\beta_G}$. As noted above in the above example of a convex-concave quadratic function, the ratio $\frac{\alpha_H}{\beta_G}$ is smaller as the influence of the interaction term between x and y on the objective function value is greater than that to the other terms, that is, as b^2/ac is greater. The other factor, $\frac{\alpha_H}{\beta_H}$, is smaller as the condition number $\text{Cond}(H_{x,x}(x, y)(H_{x,x}^*)^{-1})$ or $\text{Cond}(H_{y,y}(x, y)(H_{y,y}^*)^{-1})$ is higher. This depends on the change in the Hessian matrix over the search space $\mathbb{R}^m \times \mathbb{R}^n$. If the objective function is convex-concave quadratic, that is, $f(x, y) = \frac{1}{2}x^T H_{x,x} x + x^T H_{x,y} y + \frac{1}{2}y^T H_{y,y} y$, then the Hessian matrix is constant over the search space, and we have $\alpha_H/\beta_H = 1$, whereas $\beta_G = 1 + \sigma_{\max}^2(\sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1})$ and $\alpha_G = 1 + \sigma_{\min}^2(\sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1})$, where $\sigma_{\min}$ and $\sigma_{\max}$ denote the smallest and greatest singular values. Therefore, we have

$$\frac{1}{|\gamma^*|} = \left(1 + \sigma_{\max}^2 \left(\sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1}\right)\right) \left(\frac{1 + \sqrt{\epsilon}}{1 - \sqrt{\epsilon}}\right)^2. \quad (14)$$

Comparison with GDA. Theorem 1 of Liang and Stokes [2019] shows that the runtime T_ζ of the GDA (2) is $O(\frac{1}{\gamma_{\text{GDA}}} \log(\frac{1}{\zeta}))$, where

$$\frac{1}{\gamma_{\text{GDA}}} = \frac{\max_{(x,y) \in \mathbb{R}^m \times \mathbb{R}^n} \lambda_{\max}(K(x, y))}{\min_{(x,y) \in \mathbb{R}^m \times \mathbb{R}^n} \lambda_{\min}(\text{diag}(H_{x,x}(x, y)^2, (-H_{y,y}(x, y))^2))}, \quad (15)$$

$\lambda_{\min}$ and $\lambda_{\max}$ denote the smallest and greatest eigenvalues, and

$$K(x, y) = \begin{bmatrix} H_{x,x}^2 + H_{x,y}H_{y,x} & -H_{x,x}H_{x,y} + H_{x,y}(-H_{y,y}) \\ -H_{y,x}H_{x,x} + (-H_{y,y})H_{y,x} & (-H_{y,y})^2 + H_{y,x}H_{x,y} \end{bmatrix}, \quad (16)$$

where we drop (x, y) from $H_{x,x}(x, y)$, $H_{y,y}(x, y)$, $H_{x,y}(x, y)$, and $H_{y,x}(x, y)$ for compact expression. To compare this with our result, consider the pre-conditioned convex-concave quadratic function $f(x, y) = \frac{1}{2}x^T Ix + x^T \sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1} y + \frac{1}{2}y^T (-I)y$. Then, it may be easily observed that $\lambda_{\min}(\text{diag}(H_{x,x}(x, y)^2, (-H_{y,y}(x, y))^2)) = 1 = \alpha_H = \beta_H$ and $\lambda_{\max}(K(x, y)) = 1 + \sigma_{\max}^2(\sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1}) = \beta_G$. Therefore, $\frac{1}{\gamma_{\text{GDA}}} = 1 + \sigma_{\max}^2(\sqrt{H_{x,x}}^{-1} H_{x,y} \sqrt{-H_{y,y}}^{-1})$. Because the GDA requires one ∇f call per iteration, T_{ζ} is the runtime w.r.t. ∇f calls as well. It indicates that the runtime bound w.r.t. f and/or ∇f calls are the same for the GDA and the oracle-based saddle point optimization (7) with linearly convergent oracles, whose runtime is obtained by substituting (14) into (12), by ignoring the effect of ϵ . Note, however, that the runtime of the GDA depends on the pre-conditioning, as it is a first-order approach. The number of oracle calls of the oracle-based saddle point optimization is independent of the pre-conditioning, but the number of f and/or ∇f calls in each oracle call may depend on the pre-conditioning.

4 SADDLE POINT OPTIMIZATION WITH LEARNING RATE ADAPTATION

In this section, we propose practical implementations of the saddle point optimization approach (7) with a heuristic mechanism to adapt the learning rate η . We implement the proposed approach using two minimization routines. The first is (1+1)-CMA-ES, which is a zero-order randomized hill-climbing approach. The second is SLSQP, which is a first-order deterministic hill-climbing approach.

4.1 Learning Rate Adaptation

The main limitation of oracle-based saddle point optimization when it is applied to a simulation-based optimization task is that we rarely know the right η value in advance. As we see in Theorem 3.3, $\eta < 2 \cdot \eta^*$ must be selected to guarantee convergence on a convex-concave function. However, the optimal value, η^* , depends on the problem characteristics and is unknown in advance when considering a black-box setting. In practice, it is a tedious task to find a reasonable η .

To address this issue, we propose adapting η during the optimization process. The overall framework is presented in Algorithm 1, where we assume $f_{\mathcal{Y}} = f$ for the moment, that is, $\mathcal{Y} = \emptyset$ to simplify the main idea.

The main idea is to estimate the convergence speed in terms of the suboptimality error by running N_{step} iterations of algorithm (7) with a candidate learning rate η_c (lines 6–21). If the estimated convergence speed $\tilde{\gamma}_c$ associated with η_c is better (greater absolute value with a negative sign) than the estimated convergence speed $\tilde{\gamma}$ associated with the base learning rate η , then we replace η with η_c (lines 22–27). The next candidate learning rate is chosen randomly from $\min(\eta \cdot c_{\eta}, 1)$ (greater learning rate), η (current learning rate), and $\min(\eta/c_{\eta}, \eta_{\min})$ (smaller learning rate) with equal probability, where $\eta_{\min}$ is the minimal learning rate value and $c_{\eta} > 1$ is the hyperparameter that determines the granularity of the η update. A smaller c_{η} results in a smoother η change, but it may require more time to adapt η . It is advised to set $c_{\eta} < 2$, because Theorem 3.3 indicates that the upper bound on η for convergence is $2 \cdot \eta^*$, where η^* is the optimal value.

We estimate the convergence speed by running the algorithm for N_{step} iterations. The suboptimality error $G(x, y)$ is approximated by F_s in line 17. Because of the oracle condition (8), if there exists a unique (hence, global) min-max saddle point, then we have

ALGORITHM 1: Saddle Point Optimization with Learning Rate Adaptation

Require: $x \in \mathbb{X}, y \in \mathbb{Y}, \theta^x, \theta^y, a_\eta > 0, b_\eta \geq 0, c_\eta > 1$
Require: (optional) $P_x, P_y, \eta_{\min} \geq 0, G_{\text{tol}} \geq 0, d_{\min}^y \geq 0$

```

1:  $\eta \leftarrow 1, \tilde{\gamma} \leftarrow 0, \mathcal{Y} \leftarrow \emptyset, \mathcal{X} \leftarrow \emptyset$ 
2: for  $t = 1, \dots, T$  do
3:    $(x_t, \tilde{x}, \theta_t^x, y_t, \tilde{y}, \theta_t^y) \leftarrow (x, x, \theta^x, y, y, \theta^y)$ 
4:    $\eta_c \leftarrow \{\min(\eta \cdot c_\eta, 1), \eta, \max(\eta/c_\eta, \eta_{\min})\}$  w.p. 1/3 for each
5:    $N_{\text{step}} \leftarrow \lfloor b_\eta + a_\eta/\eta_c \rfloor$ 
6:   for  $s = 1, \dots, N_{\text{step}}$  do
7:     Let  $f_{\mathcal{Y}}(x, y) := \max_{y' \in \mathcal{Y} \cup \{y\}} f(x, y')$ 
8:      $(\tilde{x}, \theta^x) \leftarrow (x, \theta_t^x)$  if  $f_{\mathcal{Y}}(\tilde{x}, y) > f_{\mathcal{Y}}(x, y)$ 
9:      $(\tilde{y}, \theta^y) \leftarrow (y, \theta_t^y)$  if  $f(x, \tilde{y}) < f(x, y)$ 
10:     $\tilde{x}, \theta^x \leftarrow \mathcal{M}_x(f_{\mathcal{Y}}(\cdot, y), \tilde{x}; \theta^x)$ 
11:     $\tilde{y}, \theta^y \leftarrow \mathcal{M}_y(-f(x, \cdot), \tilde{y}; \theta^y)$ 
12:     $(\tilde{x}, \theta^x) \leftarrow (x', \theta_t^x)$  if  $f_{\mathcal{Y}}(x', y) < f_{\mathcal{Y}}(\tilde{x}, y)$  for  $x' \sim P_x$ 
13:    if  $f(x, y') > f(x, \tilde{y})$  for  $y' \sim P_y$  then
14:       $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{\tilde{y}\}$  if  $f(x, \tilde{y}) \geq f_{\mathcal{Y}}(x, y)$  and  $\|\bar{y} - \tilde{y}\| > d_{\min}^y$  for all  $\bar{y} \in \mathcal{Y}$ 
15:       $(\tilde{y}, \theta^y) \leftarrow (y', \theta_t^y)$ 
16:    end if
17:     $F_s \leftarrow f_{\mathcal{Y}}(x, \tilde{y}) - f_{\mathcal{Y}}(\tilde{x}, y)$ 
18:     $(x, y) \leftarrow (x, y) + \eta_c(\tilde{x} - x, \tilde{y} - y)$ 
19:    break if  $s \geq b_\eta$  and  $F_s > \dots > F_{s-b_\eta+1}$ 
20:  end for
21:   $\tilde{\gamma}_c, \sigma_{\tilde{\gamma}_c} \leftarrow \text{SLOPE}(\log(F_1), \dots, \log(F_s))$ 
22:  if  $\tilde{\gamma} \geq 0$  and  $\tilde{\gamma}_c \geq 0$  then
23:     $\eta \leftarrow \eta/c_\eta^3$ 
24:    else if  $\tilde{\gamma}_c \leq \tilde{\gamma}$  or  $\eta = \eta_c$  then
25:       $\eta \leftarrow \eta_c, \tilde{\gamma} \leftarrow \tilde{\gamma}_c$ 
26:    end if
27:     $(x, y, \theta^x, \theta^y) \leftarrow (x_t, y_t, \theta_t^x, \theta_t^y)$  if  $\tilde{\gamma}_c - 2\sigma_{\tilde{\gamma}_c} > 0$ 
28:    if  $F_s \leq G_{\text{tol}}$  then
29:       $\mathcal{X} \leftarrow \mathcal{X} \cup \{x\}$ 
30:       $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{y\}$  if  $\|\bar{y} - \tilde{y}\| > d_{\min}^y$  for all  $\bar{y} \in \mathcal{Y}$ 
31:      Re-initialize  $x, y, \theta^x, \theta^y$  and reset  $\eta \leftarrow 1$  and  $\tilde{\gamma} \leftarrow 0$ 
32:    end if
33:  end for
34: return  $\text{argmin}_{x' \in \mathcal{X} \cup \{x\}} f_{\mathcal{Y}}(x', y)$ 

```

$(1 - \max(\epsilon_x, \epsilon_y)) \cdot G(x, y) \leq F_s \leq G(x, y)$. Then, we have

$$\frac{1}{N_{\text{step}} - 1} \left| \log \left(\frac{G(x_{N_{\text{step}}}, y_{N_{\text{step}}})}{G(x_1, y_1)} \right) - \log \left(\frac{F_{N_{\text{step}}}}{F_1} \right) \right| \leq \frac{|\log(1 - \max(\epsilon_x, \epsilon_y))|}{N_{\text{step}} - 1}.$$

Based on Theorem 3.3, if the objective function is strongly convex-concave, then the convergence speed will be proportional to $1/\eta$. Then, to approximate the convergence speed in line 17, $N_{\text{step}} \in \Omega(1/\eta)$ must be set to alleviate the estimation error, that is, the right-hand side of the above inequality. Therefore, we set $N_{\text{step}} = \lfloor b_\eta + a_\eta/\eta_c \rfloor$, where $a_\eta > 0$ and $b_\eta \geq 0$ are the

hyperparameters. The greater they are, the more accurate the estimated convergence speed, but the slower the speed of adaptation of η . If the objective function is not strongly convex-concave, then the above argument may not hold, yet we optimistically expect it to reflect the convergence speed of the algorithm toward a local min-max saddle point.

After estimating the convergence speed $\tilde{\gamma}_c$ for the candidate learning rate η_c , we replace η and $\tilde{\gamma}$ with η_c and $\tilde{\gamma}_c$ if $\tilde{\gamma}_c$ is equal to or better than the convergence speed $\tilde{\gamma}$ for the current learning rate η (line 25). We also update $\tilde{\gamma}$ when $\eta = \eta_c$. If both $\tilde{\gamma}$ and $\tilde{\gamma}_c$ are non-negative, then the learning rate is too high, and we reduce η by multiplying $1/c_\eta^3$. If $\tilde{\gamma}_c - 2\sigma_{\tilde{\gamma}} > 0$, where $\sigma_{\tilde{\gamma}}$ is the estimated standard deviation of $\tilde{\gamma}$, then we revert the solutions and other strategy parameters θ^x and θ^y .

Based on our preliminary experiments and the above argument, we set $a_\eta = 1$, $b_\eta = 5$, and $c_\eta = 1.1$ as the default values.

4.2 Ingenuity for Practical Use

Our approach is designed to locate a min-max saddle point. However, in practice, we often cannot guarantee the existence of min-max saddle points. In such a situation, x and y may not converge and oscillate. For example, consider $f(x, y) = x^T y$ on $[-1, 1] \times [-1, 1]$. The worst y is -1 if $x < 0$ and 1 if $x > 0$, and the best x is 1 if $y < 0$ and -1 if $y > 0$. This causes a cyclic behavior: x is positive, then y becomes positive, then x becomes negative, then y becomes negative, and so on. To stabilize the algorithm in such situations, we maintain a set $\mathcal{Y}$ of $y \in \mathbb{Y}$ and replace f with $f_y(x, y) := \max_{y' \in \mathcal{Y} \cup \{y\}} f(x, y')$ using the approach described in Section 4.1. In the above example, provided that there are points $y_1 < 0$ and $y_2 > 0$ in $\mathcal{Y}$, the optimal x of f_y is zero regardless of y . This is the optimal solution for $\min_{-1 \leq x \leq 1} \max_{-1 \leq y \leq 1} f(x, y)$. However, if we replace $-f$ with $-f_y$ for the objective function of $\mathcal{M}_y$, then the optimization is likely to fail, because $f_y(x, y)$ is constant with respect to y over $\{y \in \mathbb{Y} : f(x, y) \geq f(x, y') \text{ for some } y' \in \mathcal{Y}\}$. Therefore, we replace f with f_y only for the parts regarding x optimization. We initialize $\mathcal{Y}$ with the empty set; hence, $f_y = f$ at the beginning. The output $\tilde{y}$ of $\mathcal{M}_y$ is registered to $\mathcal{Y}$ if a random sample $y' \sim P_y$ provides a worse objective value than $\tilde{y}$, $f(x, \tilde{y}) \geq f_y(x, y)$, and none of the registered points $\bar{y} \in \mathcal{Y}$ is in the closed ball centered at $\tilde{y}$ with radius $d_{\min}^y$, which is a hyperparameter.

The existence of multiple local min-max saddle points is another difficulty that is often encountered in practice. For such problems, we would like to locate a local min-max saddle point whose worst-case objective value is as small as possible. To tackle this difficulty, we implement a restart strategy in lines 28–32 of Algorithm 1. First, we check whether the current solution is nearly a local min-max saddle point by checking $F_s \leq G_{\text{tol}}$, where G_{tol} is a user-defined threshold parameter. Note that F_s can be close to zero at a local min-max saddle point even if the true suboptimality error is nonzero, because F_s is computed using the outputs of $\mathcal{M}_x$ and $\mathcal{M}_y$. Therefore, a small F_s value is a sign of a local min-max saddle point. If this restart criterion is satisfied, then we register the current solution x as a candidate for the final solution and append the current y to $\mathcal{Y}$ unless it is sufficiently close to the already registered points in $\mathcal{Y}$. Then, we re-initialize the solutions x and y and the internal parameters θ^x and θ^y , and restart the search with $\eta = 1$.

The other details are described as follows. First, we allow the sharing of the internal parameters θ^x and θ^y of $\mathcal{M}_x$ and $\mathcal{M}_y$ over oracle calls. Second, we feed the last outputs $\tilde{x}$ and $\tilde{y}$ to $\mathcal{M}_x$ and $\mathcal{M}_y$ as the reference points instead of the current solutions x and y if the former is better. This contributes to realizing smaller approximation errors ϵ . Third, we optionally try random samples $x' \sim P_x$ and $y' \sim P_y$ and check if they are better than the outputs of the oracles if P_x and P_y are given. A typical choice for P_x and P_y is the uniform distribution on $\mathbb{X}$ and $\mathbb{Y}$ if they are bounded. Fourth, we optionally introduce the minimal learning rate $\eta_{\min}$. Because a small η slows down the optimization speed, it is not practical to set an extremely small η , even though it is necessary for convergence.

4.3 Adversarial-CMA-ES

We implemented the proposed approach with (1+1)-CMA-ES as $\mathcal{M}_x$ and $\mathcal{M}_y$. The (1+1)-CMA-ES is a derivative-free randomized hill-climbing approach with step-size adaptation and covariance matrix adaptation. It samples a candidate solution $z' \sim \mathcal{N}(z, (\sigma A)(\sigma A)^T)$, where σ is the step size and $A \cdot A^T$ is the covariance matrix. The step size is adapted with the so-called 1/5-success rule [Devroye 1972; Rechenberg 1973; Schumer and Steiglitz 1968], which maintains σ such that the probability of generating a better solution is approximately 1/5. We implemented a simplified 1/5-success rule proposed by Kern et al. [2004]. The covariance matrix was adapted with the active covariance matrix update [Arnold and Hansen 2010]. The results show empirically that the covariance matrix learned the inverse Hessian matrix on a convex quadratic function. The algorithm is summarized in Algorithm 2. We call Algorithm 1 with Algorithm 2 as $\mathcal{M}_x$ and $\mathcal{M}_y$ *Adversarial-CMA-ES*.

We shared the strategy parameter $\theta = (\sigma, A)$ over oracle calls. Here, we implicitly assumed that the objective function h of the current oracle call and that of the last oracle call are similar, because the changes in x_t and y_t are small if η is small. Then, reusing the strategy parameter of the last oracle call reduced the need for its adaptation time.

We ran (1+1)-CMA-ES until it improved the solution $\tau_{\text{es}} \cdot \ell + \tau'_{\text{es}}$ times. The reason for this procedure is described below. Because the step size is maintained such that the probability of generating a successful solution is approximately 1/5, the algorithm runs approximately $T = 5 \cdot (\tau_{\text{es}} \cdot \ell + \tau'_{\text{es}})$ iterations. It was shown by Morinaga and Akimoto [2019] that the expected runtime $\mathbb{E}[T_\epsilon]$ of (1+1)-ES with the simplified 1/5-success rule is $\Theta(\log(1/\epsilon))$ on strongly convex functions with Lipschitz continuous gradients and their strictly increasing transformations. Moreover, the scaling of the runtime with respect to dimension ℓ is $\Theta(\ell)$ on general convex quadratic functions [Morinaga et al. 2021]. Therefore, we expect that T iterations of (1+1)-CMA-ES approximates $\mathcal{M}$ with $\epsilon \in \exp(-\Theta(T/\ell)) = \exp(-\Theta(1))$. The reason that we count the number of successful iterations instead of the number of total iterations is to avoid producing no progress because of a bad initialization of each oracle call.

Another optional stopping condition is $\sigma < \bar{\sigma}_{\min}$ for a given minimal step size $\bar{\sigma}_{\min} \geq 0$. Once σ reaches $\bar{\sigma}_{\min}$, Algorithm 2 returns $\sigma = \bar{\sigma}_{\min}$. Then, the next $\mathcal{M}$ call starts with $\sigma = \bar{\sigma}_{\min}$ and it is expected to stop after a few iterations. That is, if σ for $\mathcal{M}_x$ reaches $\bar{\sigma}_{\min}$ while $\sigma > \bar{\sigma}_{\min}$ for $\mathcal{M}_y$, Algorithm 1 spends more f -calls for $\mathcal{M}_y$ than for $\mathcal{M}_x$, and vice versa.

Based on our preliminary experiments, we set $\tau_{\text{es}} = \tau'_{\text{es}} = 5$ as their default values. If they are set greater, then we expect that (1+1)-CMA-ES approximates condition (8) with a smaller ϵ .

4.4 Adversarial-SLSQP

We also implemented the algorithm with a SLSQP subroutine [Kraft 1988] to demonstrate the applicability of the proposed η adaptation mechanism. This was a first-order approach, which required access to ∇f . Unlike Adversarial-CMA-ES, no strategy parameter for SLSQP is shared over oracle calls. The maximum number of iterations is set to $\tau_{\text{slsqp}} = 5$. We used the `scipy` implementation of SLSQP as $\mathcal{M}$ in Algorithm 1. We call this first-order approach *Adversarial-SLSQP*.

5 NUMERICAL EVALUATION

Through experiments on test problems as described below, we confirmed the following hypotheses: (A) Our implementations of the proposed approach, Adversarial-CMA-ES and Adversarial-SLSQP, performed as well as the theory implies. (B) Our learning rate adaptation located a nearly optimal learning rate with little compromise of the objective function calls. (C) Local strong convexity-concavity of the objective function is necessary for good min-max performance of the proposed approach. (D) Existing coevolutionary approaches fail to converge even on a convex-concave quadratic problem.

ALGORITHM 2: (1+1)-CMA-ES as Minimization Oracle

Require: $h : \mathbb{R}^\ell \rightarrow \mathbb{R}, z \in \mathbb{R}^\ell, \sigma > 0, A \in \mathbb{R}^{\ell \times \ell}, h_z = h(z), \tau_{\text{es}}, \tau'_{\text{es}} \in \mathbb{N}$

Require: (optional) $\bar{\sigma}_{\min} \geq 0$

- 1: $c = e^{\frac{2}{2+\ell}}, c_p = \frac{1}{12}, c_c = \frac{2}{\ell+2}, c_{\text{cov}+} = \frac{2}{\ell^2+6}, c_{\text{cov}-} = \frac{0.4}{\ell^{1.6+1}}, p_{\text{thre}} = 0.44$
- 2: $p \leftarrow 0 \in \mathbb{R}^\ell, p_{\text{succ}} \leftarrow 0.5 \in [0, 1], n_{\text{succ}} = 0$
- 3: Initialize $H \in \mathbb{R}^5$ with $H_1 = h_z$ and $H_2 = H_3 = H_4 = H_5 = \infty$
- 4: **while** $n_{\text{succ}} < \tau_{\text{es}} \cdot \ell + \tau'_{\text{es}}$ **do**
- 5: $z' \leftarrow z + \sigma AN(0, I)$
- 6: $h_{z'} = h(z')$
- 7: **if** $h_{z'} \leq H_1$ **then**
- 8: $H \leftarrow (h_{z'}, H_1, H_2, H_3, H_4)$
- 9: $p_{\text{succ}} \leftarrow (1 - c_p) \cdot p_{\text{succ}} + c_p$
- 10: **if** $p_{\text{succ}} > p_{\text{thre}}$ **then**
- 11: $p \leftarrow (1 - c_c) \cdot p, c_{\text{cov}} = c_{\text{cov}+}(1 - c_c \cdot (2 - c_c))$
- 12: **else**
- 13: $p \leftarrow (1 - c_c) \cdot p + \sqrt{c_c \cdot (2 - c_c)} \frac{z' - z}{\sigma}, c_{\text{cov}} = c_{\text{cov}+}$
- 14: **end if**
- 15: $w = A_{\text{inv}} \cdot p$
- 16: $a = \sqrt{1 - c_{\text{cov}}}, b = \frac{\sqrt{1 - c_{\text{cov}}}}{\|w\|^2} \left(\sqrt{1 + \frac{c_{\text{cov}}}{1 - c_{\text{cov}}}} \|w\|^2 - 1 \right)$
- 17: $A \leftarrow a \cdot A + b \cdot (A \cdot w) \cdot w^T, A_{\text{inv}} \leftarrow \frac{1}{a} \cdot A_{\text{inv}} - \frac{b}{a^2 + a \cdot b \cdot \|w\|^2} \cdot w \cdot (w^T A_{\text{inv}})$
- 18: $\sigma \leftarrow \sigma \cdot c, z \leftarrow z', n_{\text{succ}} \leftarrow n_{\text{succ}} + 1$
- 19: **else**
- 20: $p_{\text{succ}} \leftarrow (1 - c_p) \cdot p_{\text{succ}}$
- 21: **if** $h_{z'} > H_5$ and $p_{\text{succ}} \leq p_{\text{thre}}$ **then**
- 22: $w = A_{\text{inv}} \cdot \frac{z' - z}{\sigma}$
- 23: $c_{\text{cov}} = c_{\text{cov}-}$ **if** $c_{\text{cov}-}(2 \cdot \|w\|^2 - 1) \leq 1$ **else** $c_{\text{cov}} = \frac{1}{2 \cdot \|w\|^2 - 1}$
- 24: $a = \sqrt{1 + c_{\text{cov}}}, b = \frac{\sqrt{1 + c_{\text{cov}}}}{\|w\|^2} \left(\sqrt{1 - \frac{c_{\text{cov}}}{1 + c_{\text{cov}}}} \|w\|^2 - 1 \right)$
- 25: $A \leftarrow a \cdot A + b \cdot (A \cdot w) \cdot w^T, A_{\text{inv}} \leftarrow \frac{1}{a} \cdot A_{\text{inv}} - \frac{b}{a^2 + a \cdot b \cdot \|w\|^2} \cdot w \cdot (w^T A_{\text{inv}})$
- 26: **end if**
- 27: $\sigma \leftarrow \sigma \cdot c^{-1/4}$
- 28: **end if**
- 29: $\sigma \leftarrow \sigma \cdot \frac{\|A\|_F}{\sqrt{\ell}}, A \leftarrow A \cdot \frac{\sqrt{\ell}}{\|A\|_F}, A_{\text{inv}} \leftarrow A_{\text{inv}} \cdot \frac{\|A\|_F}{\sqrt{\ell}}, p \leftarrow p \cdot \frac{\sqrt{\ell}}{\|A\|_F}$ every ℓ iterations
- 30: **break if** $\sigma < \bar{\sigma}_{\min}$
- 31: **end while**
- 32: **return** $z, \max(\sigma, \bar{\sigma}_{\min}), A$

5.1 On Convex-Concave Quadratic Functions

To confirm hypotheses (A) and (B), we ran Adversarial-CMA-ES and Adversarial-SLSQP on the following convex-concave quadratic function $f_1 : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}$ with $n = m$:

$$f_1(x, y) = \frac{a}{2} \|x\|^2 + b \langle x, y \rangle - \frac{c}{2} \|y\|^2, \quad (17)$$

where $a, c > 0$ and $b \in \mathbb{R}$. The global min-max saddle point was located at $(x^*, y^*) = (0, 0)$. The suboptimality error is $G_1(x, y) = \frac{a \cdot c + b^2}{2a \cdot c} (\|x\|^2 + \|y\|^2)$. In this problem, we have $\alpha_H = \beta_H = 1$ and

Fig. 1. The number of f -calls until $G(x, y) \leq 10^{-5}$ is reached on f_1 with $n = m = 10$ and $a = b = c = 1$. (a) Adversarial-CMA-ES with η -adaptation (adapt) and fixed $\eta = \eta^* \times 2^{\frac{3-k}{3}}$ for $k = 1, \dots, 12$. (b) Adversarial-SLSQP with η -adaptation (adapt) and fixed $\eta = \eta^* \times 2^{\frac{3-k}{3}}$ for $k = 1, \dots, 12$. The dashed lines are proportional to $\frac{1}{(\eta/\eta^*)(2-\eta/\eta^*)}$.

$\alpha_G = \beta_G = 1 + \frac{b^2}{a \cdot c}$; hence, for $\epsilon \ll 1$, we have $\eta^* \approx \bar{\eta} \approx \frac{a \cdot c}{a \cdot c + b^2}$. Moreover, for $\eta = \delta \cdot \eta^*$ for $\delta \in (0, 2)$, $\gamma = -\frac{a \cdot c}{a \cdot c + b^2} \delta(2 - \delta)$. That is, Theorem 3.3 indicates that the runtime of the proposed approach with a fixed learning rate was proportional to $(1 + \frac{b^2}{a \cdot c}) \frac{1}{\delta(2 - \delta)}$.

The experimental settings were as follows. We draw the initial solution (x, y) uniform-randomly from $[-1, 5]^m \times [-1, 5]^n$. The strategy parameters for Adversarial-CMA-ES are $\theta^x = (\sigma^x, A^x)$ and $\theta^y = (\sigma^y, A^y)$. The step sizes σ^x and σ^y are initialized as one-fourth of the length of the initialization interval, that is, $\sigma^x = \sigma^y = 1.5$. The factors A^x and A^y are initialized by the identity matrix. We used the default hyperparameter values described in the previous section. We omitted lines 12 and 13 and lines 28–32 of Algorithm 1 (i.e., neither P_x nor P_y are given and $G_{\text{tol}} = 0$) in this experiment. The minimal learning rate was set to $\eta_{\min} = 10^{-4}$. The minimal step sizes are set to $\bar{\sigma}_{\min}^x = \bar{\sigma}_{\min}^y = 0$. We run 50 independent trials for each setting, with the maximum number of f -calls of 10^7 .

Figure 1 compares the proposed approaches with and without η -adaptation mechanism. For fixed η cases, we set η to $\delta \cdot \eta^*$ with $\delta \in \{2^{\frac{3-k}{3}} : k = 1, \dots, 12\}$. We remark that for both algorithms, all the trials with $\eta = 2 \times \eta^*$ fail to converge, as implied by Theorem 3.3. As expected, the runtimes of both algorithms with fixed η were nearly proportional to $\frac{1}{(\eta/\eta^*)(2-\eta/\eta^*)}$. The best η is approximately η^* . We conclude that our implementations closely approximate the oracle condition (8) and that the proposed approach works as the theory implies.

The proposed approach with the η -adaptation mechanism succeed in converging toward the global min-max saddle point. Comparing the runtime of the η -adaptation mechanism and the best fixed $\eta = \eta^*$, we compromise the number of f -calls at most three times in the median case for both Adversarial-CMA-ES and Adversarial-SLSQP to adapt η . There are also trials that required a few times more runtime than the median case. However, considering the difficulty in tuning η in advance, we conclude that this η -adaptation mechanism is promising to waive the need for η tuning in advance.

Figures 2(a) and 2(b) show the runtime of the proposed approaches with and without η -adaptation for varying b and for varying a/c . For the fixed η case, we set $\eta = \eta^*$. It may be observed that the runtimes in terms of the number of iterations are proportional to $1 + \frac{b^2}{a \cdot c}$, as expected from Theorem 3.3. Moreover, the number of iterations was largely the same for all algorithms, as they all approximate (8) with $\epsilon \ll 1$. In contrast, the number of f -calls was different for the two algorithms.

Fig. 2. The number of iterations and the number of f -calls until $G_1(x, y) \leq 10^{-5}$ is reached on f_1 . The solid lines indicate the median and the shaded areas indicate the 10-90 percentile ranges. Dashed lines are proportional to $1 + \frac{b^2}{a \cdot c}$.

This is because Adversarial-CMA-ES is expected to spend approximately $5(\tau_{\text{es}} \times \ell + \tau'_{\text{es}}) = 275$ f -calls per oracle call, whereas Adversarial-SLSQP spends $\tau_{\text{slsqp}} = 5$ f -calls. We remark that if one of the CMA-ES in Adversarial-CMA-ES (i.e., either $\mathcal{M}_x$ or $\mathcal{M}_y$) is replaced with SLSQP, the number of f -calls will be approximately halved. Therefore, it is advisable to use SLSQP, or another first-order approach, as an approximate minimization oracle if ∇f is available and cheap to compute. Figure 2(c) shows the scaling of the runtime with respect to the dimension $n = m$. The number of iterations did not depend on the search space dimension. The number of f -calls was also constant over varying $n = m$ for Adversarial-SLSQP. However, it was proportional to $n + m$ for Adversarial-CMA-ES.³ This is because the runtime of (1+1)-CMA-ES is proportional to the dimension, and iterations must be run proportional to the search space dimension to approximate (8).

5.2 Comparison with Baseline Approach

To confirm hypotheses (C) and (D), we ran Adversarial-CMA-ES on the six test problems summarized in Table 1. In all cases, the domain of the objective function is $\mathbb{X} \times \mathbb{Y} = [-1, 5]^m \times [-1, 5]^n$. The function f_2 is globally strongly convex-concave, while f_3 is locally strongly convex-concave. The function f_4 is globally convex-concave but not strongly convex-concave. These functions exhibited a global min-max saddle point at $(x^*, y^*) = (0, 0)$ and x^* was the global optimal solution to the worst-case objective $F(x) = f(x, \hat{y}(x))$. The function f_5 was not strongly convex-concave, but the worst case y is independent of x , and the optimal x is constant over y such that $\sum_{j=1}^n y_j > 0$.

³We comment on the computational complexity of the algorithm. The bottleneck of the execution time of each iteration of Algorithm 1 is an $\mathcal{M}_x$ -call and an $\mathcal{M}_y$ -call. The execution time for the η -adaptation was negligible. The time and space complexity of Algorithm 2 per f -call is $O(\ell^2)$, where ℓ is the search space dimension. Therefore, if the number of f -calls scales linearly in $n + m$, the execution time of Adversarial-CMA-ES scales as $O(n^3 + m^3)$.

Table 1. Definition of the Test Functions $f(x, y)$ and Their Worst-case Variable $\hat{y}(x) = \operatorname{argmax}_{y \in \mathbb{Y}} f(x, y)$

$f(x, y)$	$\hat{y}(x)$
$f_2 \quad \frac{1}{2} \ x\ ^2 + \frac{1}{m} \sum_{i=1}^m x_i \sum_{j=1}^n y_j - \frac{1}{2} \ y\ ^2$	$\left(\frac{1}{m} \sum_{i=1}^m x_i \right) \mathbf{1}$
$f_3 \quad \frac{1}{2} \min \left[\ x\ ^2, \ x - 4 \cdot \mathbf{1}\ ^2 \right] + \frac{1}{m} \sum_{i=1}^m x_i \sum_{j=1}^n y_j - \frac{1}{2} \ y\ ^2$	$\left(\frac{1}{m} \sum_{i=1}^m x_i \right) \mathbf{1}$
$f_4 \quad \frac{1}{2} \ x\ ^2 + \frac{1}{m} \sum_{i=1}^m x_i \sum_{j=1}^n y_j - \left(\frac{1}{2} \ y\ ^2 \right)^2$	$\left(\frac{1}{m \cdot n} \sum_{i=1}^m x_i \right)^{\frac{1}{3}} \mathbf{1}$
$f_5 \quad \frac{1}{m} \ x\ \sum_{j=1}^n y_j$	$5 \cdot \mathbf{1}$
$f_6 \quad \frac{1}{m} \sum_{i=1}^m x_i \sum_{j=1}^n y_j - \frac{1}{2} \ y\ ^2$	$\left(\frac{1}{m} \sum_{i=1}^m x_i \right) \mathbf{1}$
$f_7 \quad \frac{1}{2} \ x\ ^2 + \frac{1}{m} \sum_{i=1}^m x_i \sum_{j=1}^n y_j$	$\begin{cases} 5 \cdot \mathbf{1} & \sum_{i=1}^m x_i \geq 0 \\ -1 \cdot \mathbf{1} & \sum_{i=1}^m x_i < 0 \end{cases}$

The optimal solutions $x^* = 0$ to the worst-case objective functions for f_6 and f_7 were not min–max saddle points.

The experimental settings were as follows. We ran Adversarial-CMA-ES with and without sampling distributions P_x and P_y . For the distributions P_x and P_y , uniform distributions over $\mathbb{X}$ and $\mathbb{Y}$ are used. Moreover, we use the same initialization as in Section 5.1. The minimal learning rate is $\eta_{\min} = 10^{-4}$. The minimal step sizes were set to $\bar{\sigma}_{\min}^x = \bar{\sigma}_{\min}^y = 0$. The restart was not performed, that is, $G_{\text{tol}} = 0$. The boundary constraint was handled using the mirroring technique, that is, the domain was virtually extended to $\mathbb{R}^m \times \mathbb{R}^n$ by defining the function value $f(x, y)$ for $(x, y) \notin \mathbb{X} \times \mathbb{Y}$ by $f(T_{\mathbb{X}}(x), T_{\mathbb{Y}}(y))$, where $T_{\mathbb{X}}$ and $T_{\mathbb{Y}}$ map each coordinate to $U - |\text{mod}(x - L, 2(U - L)) - (U - L)|$, where $U = 5$ and $L = -1$ denote the upper and lower bounds of each coordinate. The output of (1+1)-CMA-ES ($\mathcal{M}_x$ and $\mathcal{M}_y$) is repaired into the feasible domain by applying $T_{\mathbb{X}}$ and $T_{\mathbb{Y}}$. We compare the results with those of the naive baseline approach, referred to as CMA-ES(N_y). We sampled $N_y = 10$ or 100 points uniform-randomly in $\mathbb{Y}$, and they were denoted as y^k for $k = 1, \dots, N_y$. The approximate worst-case objective was defined as $F_{N_y}(x) = \max_{1 \leq k \leq N_y} f(x, y^k)$. Then, we solve F_{N_y} with (1+1)-CMA-ES (Algorithm 2) using mirroring boundary constraint handling. These algorithms are run 10 times with different initial solutions. We also compared two coevolutionary approaches, MMDE [Qiu et al. 2018] and COEVA [Al-Dujaili et al. 2019]. These approaches were implemented based on the Python code provided by Al-Dujaili et al. [2019].

Figure 3 shows the results of 10 independent trials of Adversarial-CMA-ES, CMA-ES($N_y = 10$), CMA-ES($N_y = 100$), MMDE, and COEVA. Adversarial-CMA-ES succeeds in converging the global min–max saddle point on f_2 , f_3 , and f_6 . The functions f_2 and f_3 were locally strongly convex–concave functions, and Adversarial-CMA-ES performed well, as expected. The existing coevolutionary approaches, as well as CMA-ES(N_y), failed to converge on these problems. The benchmark problems used to evaluate the performance of existing coevolutionary approaches [Branke and Rosenbusch 2008; Qiu et al. 2018; Zhou and Zhang 2010] are rather low-dimensional problems ($m \leq 2$ and $n \leq 2$). These approaches do not work well on higher-dimensional problems and perform worse than the simple baseline, CMA-ES(N_y). CMA-ES(N_y) tends to the global optimal point on f_5 . This is because the optimal x^* is optimal for approximate worst-case functions provided that there exists y in N_y samples such that $\sum_{i=1}^n y_i > 0$ holds. In contrast, no approach succeeded in converging toward the global optimum of the worst-case function on f_4 , f_6 , and f_7 . From these results, we conclude that local strong convexity–concavity is an important factor for the convergence of Adversarial-CMA-ES. These results reveal the limitations of Adversarial-CMA-ES and the difficulty of locating the solution to the worst-case objective if it is not a min–max saddle point.

Fig. 3. Results of 10 independent runs of Adversarial-CMA-ES with and without sampling distribution (denoted as Adv-CMA-ES(P) and Adv-CMA-ES(no P), respectively), CMA-ES($N_y = 10$), CMA-ES($N_y = 100$), MMDE, and COEVA. The search space dimension is $m = 50$ and $n = 20$ for all cases.

6 APPLICATION TO ROBUST BERTHING CONTROL

In this section, we analyze the application of Adversarial-CMA-ES to robust berthing control tasks under model uncertainty.

6.1 Problem Description

The objective of our robust berthing control task is to obtain a controller that controls a subject ship toward the target position near a berth while avoiding collision to the berth under the worst situation in the predefined uncertainty set. The control target was a 3 m model ship representing the vessel MV ESSO OSAKA (Figure 4(a)). The state variables $s = (X, u, Y, v_m, \psi, r) \in \mathbb{R}^6$ and the control signal $a = (\delta, n_p, n_{BT}, n_{ST}) \in \mathbb{R}^4$ were as described in Figure 4(b). The controller $u_x : s \mapsto a$ was modeled by a neural network with $m = 99$ -dimensional parameter vector x . The objective function $f(x, y)$ measures the distance between the target position and the final position of the subject ship after a pre-defined control time using the controller u_x , penalized by the risk of the collision to the berth, under an uncertainty parameter $y \in \mathbb{Y}$ described below. The details of the controller and the objective function are explained in Appendix B.

The wind conditions $y^{(A)}$ and the model coefficients $y^{(B)}$ with respect to the wind forces are treated as the uncertain factors $y = (y^{(A)}, y^{(B)})$. The following three situations are considered. (A) The state equation model is accurately modeled, but the wind conditions are uncertain. In this

Fig. 4. Control target: 3 m model ship of ESSO OSAKA.

situation, the uncertainty parameters $y^{(A)} \in \mathbb{Y}_A$ represent the wind velocity U_T [m/s] and the wind direction γ_T [rad], and their feasible values are in $\mathbb{Y}_A = [0, 0.5] \times [0, 2\pi]$. The model coefficients $y^{(B)}$ are set to the same values as in Miyauchi et al. [2022], denoted by $y_{\text{est}}^{(B)}$. (B) Wind conditions are known, but the state equation model is uncertain. The coefficients in the state equation model for the effect of the wind force were derived by Fujiwara et al. [1998] using regression of wind tunnel experiment data, and we consider them to be relatively inaccurate. The uncertainty parameters $y^{(B)}$ consist of 10 coefficients for the wind force. The feasible domain $\mathbb{Y}_B$ is constructed to include the coefficient used by Miyauchi et al. [2022], that is, $y_{\text{est}}^{(B)} \in \mathbb{Y}_B$. For each variable, the feasible values are defined by the interval. The interval of the i th component of $y^{(B)}$, denoted by $[y_{\text{est}}^{(B)}]_i$, is set to $[0.9 \cdot [y_{\text{est}}^{(B)}]_i, 1.1 \cdot [y_{\text{est}}^{(B)}]_i]$ for all $i = 1, \dots, 10$. The other model coefficients are set to the same values as Miyauchi et al. [2022] and the wind condition is set to $y_{\text{est}}^{(A)} = (1.5\pi, 0.5)$, meaning that the velocity of wind blowing orthogonally from the sea to the berth is 0.5 [m/s]. (C) Wind conditions are unknown, and the model coefficients are uncertain. In this situation, y is composed of the uncertainty parameters $y^{(A)}$ and $y^{(B)}$, and the feasible values are set to $\mathbb{Y}_C = \mathbb{Y}_A \times \mathbb{Y}_B$.

6.2 Experiment Setting

We ran Adversarial-CMA-ES and CMA-ES($N_y = 100$) on $\mathbb{Y}_A$, $\mathbb{Y}_B$, and $\mathbb{Y}_C$.⁴ As baselines, we run (1+1)-CMA-ES on $f(x, y_{\text{fix}})$ under two different situations for $y_{\text{fix}} \in \mathbb{Y}$. The first situation was $y_{\text{fix}} = (y_{\text{no}}^{(A)}, y_{\text{est}}^{(B)})$, where $y_{\text{no}}^{(A)} = (0, 0)$ corresponds to no wind disturbance, and the second situation was $y_{\text{fix}} = (y_{\text{est}}^{(A)}, y_{\text{est}}^{(B)})$, where $y_{\text{est}}^{(A)} = (1.5\pi, 0.5)$ reflects our prior knowledge that such a wind is difficult to handle for avoiding collision with the berth. Note that (1+1)-CMA-ES on $f(x, y_{\text{fix}})$ corresponds to CMA-ES($N_y = 1$) with $y^1 = y_{\text{fix}}$. Each algorithm was run 20 times independently with random

⁴CMA-ES($N_y = 10$), MMDE and COEVA tested in Section 5.2 were omitted from the comparison based on our preliminary experiments. The worst-case performance of CMA-ES($N_y = 10$) were worse than the worst-case performance of CMA-ES($N_y = 100$) on our problems. The worst-case performance of MMDE and COEVA were not competitive to the other approaches as demonstrated in Figure 3.

initialization of x and y . The search space for x and y was scaled to $\mathbb{X} = [-1, 1]^m$ and $\mathbb{Y} = [-1, 1]^n$. The box constraint was treated using the mirroring technique described in Section 5.2. The initial solution (x, y) was drawn uniform-randomly from $\mathbb{X} \times \mathbb{Y}$. For CMA-ES($N_y = 100$), y^k for $k = 1, \dots, N_y$ were uniform-randomly generated. The step sizes σ^x and σ^y are initialized as one-fourth of the length of the initialization interval. The factors A^x and A^y are initialized by the identity matrix. The minimal step size is $\bar{\sigma}_{\min} = 10^{-8}$ for both σ^x and σ^y . We set $G_{\text{tol}} = 10^{-6}$ and $d_{\min}^y = \bar{\sigma}_{\min} \times \sqrt{n}$ for Adversarial-CMA-ES. The f -call budget was 10^6 .

For Adversarial-CMA-ES, we used the restart strategy proposed in Algorithm 1. The output of Adversarial-CMA-ES follows Algorithm 1. For CMA-ES(N_y), when the termination condition $\sigma < \bar{\sigma}_{\min}$ was satisfied, the candidate solution was recorded and the algorithm was re-started until it exhausted the f -call budget. Note that y^k ($k = 1, \dots, N_y$) were not resampled. The output of CMA-ES(N_y) is determined as follows: Let $\{x^1, \dots, x^r\}$ be the set of recorded candidate solutions and the solution obtained at the end of the run. We then selected $x = \operatorname{argmin}_{i=1, \dots, r} \max_{k=1, \dots, N_y} f(x^i, y^k)$ as the output of CMA-ES(N_y).

The obtained solutions were evaluated as follows. Because the ground truth worst-case objective function value $F(x) = \max_{y \in \mathbb{Y}} f(x, y)$ for a given x is unknown, we perform numerical optimization to approximate $F(x)$. We ran (1+1)-CMA-ES for $500 \times n$ iterations to obtain a local maximal point y of $f(x, y)$. As the objective is expected to have multiple local optima, we repeat it 100 times with different initial search points y . The initialization of (1+1)-CMA-ES is as described above.

6.3 Results and Discussion

Figure 5 shows the performance of the resulting controllers of 20 independent trials of each algorithm under different situations. Some of the trajectories observed for the obtained controllers are discussed in Appendix C.

The (1+1)-CMA-ES on $y = (y_{\text{no}}^{(A)}, y_{\text{est}}^{(B)})$ achieves the best performance under no wind disturbance (Figure 5(a)), while (1+1)-CMA-ES on $y = (y_{\text{est}}^{(A)}, y_{\text{est}}^{(B)})$ achieved the best performance under the certain wind condition, $y^{(A)} = y_{\text{est}}^{(A)}$ (Figure 5(c)). In all trials, they achieved a cost of $<10^{-3}$. However, their performances significantly were degraded under the worst case, particularly when the wind condition was unknown (Figures 5(b) and 5(e)), where the ship collided with the berth and the cost was >10 . The performance was less affected by the uncertainty in the model coefficients in this experiment, but the effect would be enhanced if we consider a wider uncertainty set $\mathbb{Y}_{(B)}$. Nonetheless, these results demonstrate the importance of considering model uncertainty to obtain robust berthing control.

The controllers obtained by Adversarial-CMA-ES and CMA-ES($N_y = 100$) on $\mathbb{Y}_A$ and $\mathbb{Y}_C$ achieved better performance under the worst situation in $\mathbb{Y}_A$ (Figure 5(b)) than those obtained by the other approaches. The numbers of trials that maxA, maxC, advA, and advC succeed in avoiding collision with the berth under the worst case in $\mathbb{Y}_A$ were 3, 1, 4, and 11 of 20. Interestingly, Adversarial-CMA-ES on $\mathbb{Y}_C$ (advC) achieved better worst-case performance in $\mathbb{Y}_A$ than Adversarial-CMA-ES on $\mathbb{Y}_A$ (advA), even though advC considered a wider uncertainty set and hence was expected to show more conservative performance. The reason for the superior performance of advC may be explained as follows. Both $y^{(A)}$ and $y^{(B)}$ represent the uncertainty regarding the wind force. Considering the worst case in $\mathbb{Y}_C$ results in being more conservative to the wind force. This may help to obtain a controller that can avoid the collision with the berth, while possibly losing the accuracy of the final position.

The advantage of Adversarial-CMA-ES over CMA-ES($N_y = 100$) was more pronounced in the worst-case performance on $\mathbb{Y}_B$ (Figure 5(d)). The median of advB and that of maxB were better than the median of the other results. All 20 trials of advB achieved berthing without collision with the

Fig. 5. Performance of the controllers obtained in 20 independent trials of (1+1)-CMA-ES on $y = (y_{no}^{(A)}, y_{est}^{(B)})$ and $y = (y_{est}^{(A)}, y_{est}^{(B)})$; CMA-ES($N_y = 100$) on $\mathbb{Y}_A, \mathbb{Y}_B, \mathbb{Y}_C$; and Adversarial-CMA-ES on $\mathbb{Y}_A, \mathbb{Y}_B, \mathbb{Y}_C$, denoted by cmaA, cmaB, maxA, maxB, maxC, advA, advB, and advC, respectively. Each box indicates the lower quartile $Q1$ and the upper quartile $Q3$, with the line indicating the median $Q2$. The lower and upper whiskers are the lowest datum above $Q1 - 1.5(Q3 - Q1)$ and the highest datum below $Q3 + 1.5(Q3 - Q1)$.

berth in the worst situation. In contrast, 9 of 20 trials failed in maxB. This may be because $N_y = 100$ was not sufficiently large to represent the uncertainty in the 10-dimensional space $\mathbb{Y}_B$.

In the worst-case performance on $\mathbb{Y}_C$, 5 of 20 trials of advC succeed in avoiding a collision with the berth, whereas all the trials of maxC failed in avoiding a collision with the berth. Because of the similarity between the results on the worst-case performances on $\mathbb{Y}_A$ and $\mathbb{Y}_C$, the difficulties in obtaining a robust controller under the worst-case in $\mathbb{Y}_C$ mainly comes from the difficulty in treating the worst case in $\mathbb{Y}_A$. The results may be improved by running the optimization process longer and performing more restarts to locate more local min-max saddle points.

7 CONCLUSION

We have proposed a framework for saddle point optimization with approximate minimization oracles. Our theoretical analysis has shown the conditions under which the learning rate for the approach converges linearly (i.e., geometrically) toward the min-max saddle point on strongly convex-concave functions. Numerical evaluation have shown the tightness of the theoretical results. We have also proposed a learning rate adaptation mechanism for practical use. Numerical evaluation on convex-concave quadratic problems has demonstrated that the proposed approach with the learning rate adaptation successfully converges linearly toward the min-max saddle point, with the compromise of f -calls being no more than three times that of f -calls with the best tuned fixed learning rate. Comparison with other baseline approaches on several test problems revealed the limitations of existing coevolutionary approaches as well as of the proposed approach on problems with the optimal solution that is not a min-max saddle point. The application of the proposed approach to a robust berthing control task demonstrated the usefulness of the proposed approach, and the results imply the importance of considering modeling errors to achieve a reliable and safe solution.

We conclude the present work with some suggestions for possible avenues for future research.

The main limitation of the proposed approach as a solver to (1) is that it may fail to converge to a local minimal solution of the worst-case objective $\max_{y \in \mathbb{Y}} f(x, y)$ if it is not a strict local min-max saddle point of f . Such failure cases were observed in Figure 3, not only for the proposed approach but also for existing coevolutionary approaches. Addressing this difficulty is an important direction for future work. For the GDA approach (2), Liang and Stokes [2019] have shown that the GDA failed to converge to the optimal solution on a bi-linear function $f(x, y) = x^T C y$ and some improved gradient-based approaches [Daskalakis et al. 2018; Mescheder et al. 2017; Yadav et al. 2018] successfully converged. We expect that these gradient-based approaches would be useful in improving the here-proposed approach. The other limitation is that the best possible runtime $\Omega(1/\gamma^*)$ in (13) scales as the interaction term; more precisely, β_G/α_H , increases. We intend to address this limitation in future work.

The main limitation of our theoretical result (Theorem 3.3) is that approximate minimization oracles are required to satisfy Assumption 3.2. In practice, it is often impossible to guarantee Assumption 3.2, as we do not know the global minimum/maximum of the objective functions. For the design of Adversarial-CMA-ES and Adversarial-SLSQP, we expect that it is approximately satisfied by running a linearly convergence approximate minimization oracle until a fixed number of improvements are observed. See Section 4.3 for details. In such cases, we have condition (8) not with a constant ϵ but rather with a possibly stochastic and time-dependent sequence $\{\epsilon_t\}$, which is not covered by Theorem 3.3. Dealing with such ϵ_t will enlarge the scope of Theorem 3.3 and bridge the gap between theory and practice.⁵

⁵Another possible approach is to replace condition $h(\tilde{z}) - \min_{z \in \mathbb{Z}} h(z) \leq \epsilon \cdot (h(\tilde{z}) - \min_{z \in \mathbb{Z}} h(z))$ in Assumption 3.2 with condition $\|\nabla h(\tilde{z})\| < \epsilon \cdot \|\nabla h(\tilde{z})\|$ under the additional assumption that the objective function is continuously differentiable.

The uncertainty parameter is typically constrained in a bounded set $\mathbb{Y} \subset \mathbb{R}^n$ in practice, however, the effect of $\mathbb{Y}$ on the convergence rate has not been investigated in this work. Our theoretical result was developed for unconstrained situation and the proof does not immediately generalize to the constrained situation. The effect of the bound $\mathbb{Y}$ on the convergence rate is to be investigated theoretically and empirically.

The experimental results on the robust berthing control task have demonstrated the usefulness of the proposed approach and the importance of considering model uncertainties. Simultaneously, they revealed the difficulty of obtaining a robust solution with satisfactory utility. Regarding the wind condition uncertainty, it is possible to decompose $\mathbb{Y}_A$ into disjoint subsets (e.g., based on the wind direction), train the robust feedback controller for each subset, and switch the controller based on the wind condition measured at the time of operation. Such an approach is not available for the uncertainty in the model coefficients. To improve the worst-case performance, it is important to reduce the set of uncertain parameter values $\mathbb{Y}$ as much as possible. In our experiments, we defined the interval for each uncertain coefficient to form $\mathbb{Y}$, but the corner case may be unrealistic and will degrade the worst-case performance unnecessarily. Designing more intelligent $\mathbb{Y}$ remains as a very important task for practical applications.

APPENDIX

A PROOFS

A.1 Proof of Proposition 2.3

PROOF. Assume that (x^*, y^*) is a local min-max saddle point of f . Then, by definition, there exists a neighborhood $\mathcal{E}_x \times \mathcal{E}_y$ of (x^*, y^*) such that $f(x, y^*) \geq f(x^*, y^*) \geq f(x^*, y)$ holds for any $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$. Let $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$. Then, $G_x(x, y^*) = f(x, y^*) - \min_{x' \in \mathbb{X}} f(x', y^*) \geq f(x^*, y^*) - \min_{x' \in \mathbb{X}} f(x', y^*) = G_x(x^*, y^*)$ and $G_y(x^*, y) = \max_{y' \in \mathbb{Y}} f(x^*, y') - f(x^*, y) \geq \max_{y' \in \mathbb{Y}} f(x^*, y') - f(x^*, y^*) = G_y(x^*, y^*)$. This implies that x^* and y^* are local minimal points of $G_x(x, y^*)$ and $G_y(x^*, y)$, respectively.

Conversely, assume that x^* and y^* are local minimal points of $G_x(x, y^*)$ and $G_y(x^*, y)$, respectively. Then, there exists a neighborhood $\mathcal{E}_x \times \mathcal{E}_y$ of (x^*, y^*) such that $G_x(x, y^*) \geq G_x(x^*, y^*)$ and $G_y(x^*, y) \geq G_y(x^*, y^*)$ for any $(x, y) \in \mathcal{E}_x \times \mathcal{E}_y$. They read $f(x, y^*) \geq f(x^*, y^*)$ and $f(x^*, y) \geq f(x^*, y^*)$, which implies that (x^*, y^*) is a local min-max saddle point of f .

If (x^*, y^*) is the global min-max saddle point of f , then (x^*, y^*) is a local minimal point of G . Moreover, we have $G_x(x^*, y^*) = G_y(x^*, y^*) = 0$, implying that it is the global minimal point of G . Conversely, if (x^*, y^*) is the global minimal point of G , then it is a local min-max saddle point. Moreover, because the global minimum of G is zero, we have $G_x(x^*, y^*) = G_y(x^*, y^*) = 0$. Then, we can take $\mathcal{E}_x = \mathbb{X}$ and $\mathcal{E}_y = \mathbb{Y}$ in the above proof, which implies that (x^*, y^*) is the global min-max saddle point. $\square$

A.2 Proof of Lemma 2.6

PROOF. Noting that $(\nabla_x f)(\hat{x}(y), y) = 0$ and $(\nabla_y f)(x, \hat{y}(x)) = 0$, we obtain

$$\begin{aligned} \nabla_x G_x(x, y) &= (\nabla_x f)(x, y), & \nabla_y G_x(x, y) &= (\nabla_y f)(x, y) - (\nabla_y f)(\hat{x}(y), y), \\ \nabla_x G_y(x, y) &= (\nabla_x f)(x, \hat{y}(x)) - (\nabla_x f)(x, y), & \nabla_y G_y(x, y) &= -(\nabla_y f)(x, y). \end{aligned} \quad (18)$$

An advantage of this approach is that this condition can be approximately verified by estimating the gradients of the objective function [Nesterov and Spokoiny 2017].

Moreover, we have

$$\begin{aligned}\nabla^2 G_x(x, y) &= \begin{bmatrix} H_{x,x}(x, y) & H_{x,y}(x, y) \\ H_{y,x}(x, y) & H_{y,y}(x, y) - H_{y,y}(\hat{x}(y), y) - [J_{\hat{x}}(\hat{x}(y), y)]^T H_{x,y}(\hat{x}(y), y) \end{bmatrix}, \\ \nabla^2 G_y(x, y) &= \begin{bmatrix} -H_{x,x}(x, y) + H_{x,x}(x, \hat{y}(x)) + [J_{\hat{y}}(x, \hat{y}(x))]^T H_{y,x}(x, \hat{y}(x)) & -H_{x,y}(x, y) \\ -H_{y,x}(x, y) & -H_{y,y}(x, y) \end{bmatrix}.\end{aligned}$$

In light of Proposition 2.5 and the symmetry $H_{x,y} = H_{y,x}^T$, we have $[J_{\hat{x}}(\hat{x}(y), y)]^T = -H_{y,x}(\hat{x}(y), y)(H_{x,x}(\hat{x}(y), y))^{-1}$ and $[J_{\hat{y}}(x, \hat{y}(x))]^T = -H_{x,y}(x, \hat{y}(x))(H_{y,y}(x, \hat{y}(x)))^{-1}$. Then, because $\nabla^2 G = \nabla^2(G_x + G_y) = \nabla^2 G_x + \nabla^2 G_y$, we obtain

$$\nabla^2 G(x, y) = \begin{bmatrix} G_{x,x}(x, \hat{y}(x)) & 0 \\ 0 & G_{y,y}(\hat{x}(y), y) \end{bmatrix}.$$

The symmetry of $G_{x,x}$ and $G_{y,y}$ are clear from their definitions. The positivity of $G_{x,x}$ and $G_{y,y}$ follows that $H_{x,x} > 0$, $-H_{y,y} > 0$, $H_{x,y}(-H_{y,y})^{-1}H_{y,x} \geq \sigma_{\min}(H_{x,y})^2/\sigma_{\max}(-H_{y,y}) > 0$ and $H_{y,x}(H_{x,x})^{-1}H_{x,y} \geq \sigma_{\min}(H_{x,y})^2/\sigma_{\max}(H_{x,x}) > 0$. This completes the proof. $\square$

A.3 Proof of Theorem 3.3

PROOF. Let $v_x = \tilde{x}_t - x_t$, $v_y = \tilde{y}_t - y_t$, and $v = (v_x, v_y)$. Define $\bar{x}(\tau) = x_t + \tau \cdot v_x$ and $\bar{y}(\tau) = y_t + \tau \cdot v_y$. Let $w_x = \hat{x}(y_t) - x_t$ and $w_y = \hat{y}(x_t) - y_t$. Define $\bar{\bar{x}}(\tau) = x_t + \tau \cdot w_x$ and $\bar{\bar{y}}(\tau) = y_t + \tau \cdot w_y$. Then, $\bar{x}(0) = \bar{\bar{x}}(0) = x_t$ and $\bar{y}(0) = \bar{\bar{y}}(0) = y_t$. Moreover, $\bar{x}(\eta) = x_{t+1}$ and $\bar{y}(\eta) = y_{t+1}$.

We first derive several inequalities on the norms of v and w . Note that $\nabla_x f(\hat{x}(y_t), y_t) = 0$ and $\nabla_y f(x_t, \hat{y}(x_t)) = 0$. In light of conditions (1) and (2) in the theorem statement, we have

$$(\alpha_H/2)\|w_x\|_{H_{x,x}^*}^2 \leq G_x(x_t, y_t) = f(x_t, y_t) - f(\hat{x}(y_t), y_t) \leq (\beta_H/2)\|w_x\|_{H_{x,x}^*}^2, \quad (19)$$

$$(\alpha_H/2)\|w_y\|_{-H_{y,y}^*}^2 \leq G_y(x_t, y_t) = f(x_t, \hat{y}(x_t)) - f(x_t, y_t) \leq (\beta_H/2)\|w_y\|_{-H_{y,y}^*}^2. \quad (20)$$

Because of condition (8), we have

$$(\alpha_H/2)\|w_x - v_x\|_{H_{x,x}^*}^2 \leq f(\tilde{x}_t, y_t) - f(\hat{x}(y_t), y_t) \leq \epsilon \cdot G_x(x_t, y_t), \quad (21)$$

$$(\alpha_H/2)\|w_y - v_y\|_{-H_{y,y}^*}^2 \leq f(x_t, \hat{y}(x_t)) - f(x_t, \tilde{y}_t) \leq \epsilon \cdot G_y(x_t, y_t). \quad (22)$$

Then, from Equations (19) to (22), we have

$$\|v_x\|_{H_{x,x}^*}^2 \leq (\|w_x\|_{H_{x,x}^*} + \|v_x - w_x\|_{H_{x,x}^*})^2 \leq 2[(1 + \sqrt{\epsilon})^2/\alpha_H] \cdot G_x(x_t, y_t), \quad (23)$$

$$\|v_y\|_{-H_{y,y}^*}^2 \leq (\|w_y\|_{-H_{y,y}^*} + \|v_y - w_y\|_{-H_{y,y}^*})^2 \leq 2[(1 + \sqrt{\epsilon})^2/\alpha_H] \cdot G_y(x_t, y_t). \quad (24)$$

Because $\epsilon < \alpha_H/\beta_H$, we also have

$$\|v_x\|_{H_{x,x}^*}^2 \geq (\|w_x\|_{H_{x,x}^*} - \|v_x - w_x\|_{H_{x,x}^*})^2 \geq (2/\beta_H)[1 - \sqrt{(\beta_H/\alpha_H) \cdot \epsilon}]^2 G_x(x_t, y_t), \quad (25)$$

$$\|v_y\|_{-H_{y,y}^*}^2 \geq (\|w_y\|_{-H_{y,y}^*} - \|v_y - w_y\|_{-H_{y,y}^*})^2 \geq (2/\beta_H)[1 - \sqrt{(\beta_H/\alpha_H) \cdot \epsilon}]^2 G_y(x_t, y_t). \quad (26)$$

By applying the mean value theorem repeatedly, we have

$$\begin{aligned}
 G(x_{t+1}, y_{t+1}) - G(x_t, y_t) &= G(\bar{x}(\eta), \bar{y}(\eta)) - G(\bar{x}(0), \bar{y}(0)) \\
 &= \int_0^\eta \nabla G(\bar{x}(\tau), \bar{y}(\tau))^T d\tau \cdot v \\
 &= \int_0^\eta \left[\nabla G(\bar{x}(0), \bar{y}(0)) + \int_0^\tau \nabla^2 G(\bar{x}(s), \bar{y}(s)) ds \cdot v \right]^T d\tau \cdot v \\
 &= \eta \cdot \nabla G(x_t, y_t)^T \cdot v + v^T \int_0^\eta \int_0^\tau \nabla^2 G(\bar{x}(s), \bar{y}(s))^T ds d\tau \cdot v.
 \end{aligned} \tag{27}$$

First, we evaluate the first term on the right-most side of Equation (27). Note first that we have $\nabla G(x_t, y_t)^T v = \nabla_x G(x_t, y_t)^T v_x + \nabla_y G(x_t, y_t)^T v_y$. In light of Equation (18), we have $\nabla_x G(x_t, y_t) = (\nabla_x f)(x_t, \hat{y}(x_t)) = (\nabla_x f)(\bar{x}(0), \bar{y}(1))$. Noting that $(\nabla_x f)(\bar{x}(1), \bar{y}(0)) = (\nabla_x f)(\hat{x}(y_t), y_t) = 0$, by the mean value theorem, we obtain

$$\nabla_x G(x_t, y_t)^T v_x = -w_x^T \cdot \int_0^1 H_{x,x}(\bar{x}(1-\tau), \bar{y}(\tau)) d\tau \cdot v_x + w_y^T \cdot \int_0^1 H_{y,x}(\bar{x}(1-\tau), \bar{y}(\tau)) d\tau \cdot v_x. \tag{28}$$

Analogously, we obtain

$$\nabla_y G(x_t, y_t)^T v_y = w_y^T \cdot \int_0^1 H_{y,y}(\bar{x}(1-\tau), \bar{y}(\tau)) d\tau \cdot v_y - w_x^T \cdot \int_0^1 H_{x,y}(\bar{x}(1-\tau), \bar{y}(\tau)) d\tau \cdot v_y. \tag{29}$$

Inserting $v_x = w_x + (v_x - w_x)$ and $v_y = w_y + (v_y - w_y)$ into Equations (28) and (29), using $H_{x,y} = H_{y,x}^T$, and rearranging them, we obtain

$$\nabla G(x_t, y_t)^T v = \int_0^1 \left(- \begin{bmatrix} w_x \\ w_y \end{bmatrix}^T \begin{bmatrix} H_{x,x} & 0 \\ 0 & -H_{y,y} \end{bmatrix} \begin{bmatrix} w_x \\ w_y \end{bmatrix} - \begin{bmatrix} v_x - w_x \\ v_y - w_y \end{bmatrix}^T \begin{bmatrix} H_{x,x} & -H_{x,y} \\ H_{y,x} & -H_{y,y} \end{bmatrix} \begin{bmatrix} w_x \\ w_y \end{bmatrix} \right) d\tau, \tag{30}$$

where we drop $(\bar{x}(1-\tau), \bar{y}(\tau))$ from $H_{x,x}$, $H_{x,y}$, $H_{y,x}$ and $H_{y,y}$ for compact expressions.

We aim to bound each term of Equations (27) and (30). The second term on Equation (27) is bounded by using conditions (3) and (4) of the theorem statement as well as the fact that $\nabla^2 G$ is block-diagonal (Lemma 2.6) as

$$\alpha_G \|v\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)}^2 \leq v^T \nabla^2 G(\bar{x}(s), \bar{y}(s))^T \cdot v \leq \beta_G \|v\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)}^2. \tag{31}$$

The first term on Equation (30) is bounded by using conditions (1) and (2) of the theorem statement as

$$-\beta_H \|w\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)}^2 \leq - \begin{bmatrix} w_x \\ w_y \end{bmatrix}^T \begin{bmatrix} H_{x,x} & 0 \\ 0 & -H_{y,y} \end{bmatrix} \begin{bmatrix} w_x \\ w_y \end{bmatrix} \leq -\alpha_H \|w\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)}^2. \tag{32}$$

The second term on Equation (30) is bounded as

$$\begin{aligned}
 &\left| \begin{bmatrix} v_x - w_x \\ v_y - w_y \end{bmatrix}^T \begin{bmatrix} H_{x,x} & -H_{x,y} \\ H_{y,x} & -H_{y,y} \end{bmatrix} \begin{bmatrix} w_x \\ w_y \end{bmatrix} \right| \leq \|w - v\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)} \cdot \|w\|_{\text{diag}(H_{x,x}^*, -H_{y,y}^*)} \\
 &\quad \cdot \sigma_{\max} \left(\begin{bmatrix} \sqrt{H_{x,x}^*}^{-1} H_{x,x} \sqrt{H_{x,x}^*}^{-1} & -\sqrt{H_{x,x}^*}^{-1} H_{x,y} \sqrt{-H_{y,y}^*}^{-1} \\ \sqrt{-H_{y,y}^*}^{-1} H_{y,x} \sqrt{H_{x,x}^*}^{-1} & \sqrt{-H_{y,y}^*}^{-1} (-H_{y,y}) \sqrt{-H_{y,y}^*}^{-1} \end{bmatrix} \right), \tag{33}
 \end{aligned}$$

where the greatest singular value is bounded by using conditions (1)–(4) in the theorem statement as

$$\sigma_{\max} \left(\begin{bmatrix} \sqrt{H_{x,x}^*}^{-1} H_{x,x} \sqrt{H_{x,x}^*}^{-1} & -\sqrt{H_{x,x}^*}^{-1} H_{x,y} \sqrt{H_{y,y}^*}^{-1} \\ \sqrt{H_{y,y}^*}^{-1} H_{y,x} \sqrt{H_{x,x}^*}^{-1} & \sqrt{H_{y,y}^*}^{-1} (-H_{y,y}) \sqrt{H_{y,y}^*}^{-1} \end{bmatrix} \right) \leq \beta_H \cdot \beta_G^{1/2} / \alpha_H^{1/2}. \quad (34)$$

Equations (19) to (24), (27) and (30) to (34) lead to

$$G(x_{t+1}, y_{t+1}) - G(x_t, y_t) \leq \left[-2\eta \frac{\alpha_H}{\beta_H} \left(1 - \frac{\beta_H^2}{\alpha_H^2} \sqrt{\frac{\beta_G}{\alpha_H}} \cdot \epsilon \right) + \eta^2 \cdot (1 + \sqrt{\epsilon})^2 \cdot \frac{\beta_G}{\alpha_H} \right] G(x_t, y_t). \quad (35)$$

Here, the right-hand side is $\gamma \cdot G(x_t, y_t)$ with γ defined in Equation (10). Hence, $G(x_{t+1}, y_{t+1}) \leq (1 + \gamma) \cdot G(x_t, y_t)$. Note that $\log(1 + \gamma) < \gamma$ for all $\gamma \in (-1, 0)$, we thus obtain $\log(G(x_{t+1}, y_{t+1})) - \log(G(x_t, y_t)) < \gamma$. Because $\log(G(x_t, y_t)) - \log(G(x_0, y_0)) < \gamma \cdot t$, the minimal t that $\log(G(x_t, y_t)) - \log(G(x_0, y_0)) \leq \log(\zeta)$ is no greater than $\lceil \frac{1}{\gamma} \log(\frac{1}{\zeta}) \rceil = T_\zeta$. Similarly, Equations (19) to (22), (25) to (27) and (30) to (34) lead to

$$G(x_{t+1}, y_{t+1}) - G(x_t, y_t) \geq \left[-2\eta \frac{\beta_H}{\alpha_H} \left(1 + \sqrt{\frac{\beta_G}{\alpha_H}} \cdot \epsilon \right) + \eta^2 \frac{\alpha_G}{\beta_H} \left(1 - \sqrt{\frac{\beta_H}{\alpha_H}} \cdot \epsilon \right)^2 \right] G(x_t, y_t). \quad (36)$$

The right-hand side of Equation (36) is positive if $\eta > 2 \cdot \bar{\eta}$. This completes the proof. $\square$

B DETAILS OF AUTOMATIC BERTHING CONTROL PROBLEM

Subject Ship. The control target is a 3 m model ship of MV ESSO OSAKA (Figure 4), following a related study [Maki et al. 2021, 2020]. The state variables $s = (X, u, Y, v_m, \psi, r) \in \mathbb{R}^6$ are the X [m] and Y [m] coordinates of the Earth-fixed coordinate system, the longitudinal velocity u [m/s] and the lateral velocity v_m [m/s] at the mid-ship, and the yaw direction ψ [rad] as seen from the X coordinates and the yaw angular velocity r [rad/s]. The control signal $a = (\delta, n_p, n_{BT}, n_{ST}) \in \mathbb{R}^4$ consists of the rudder angle δ [rad], propeller revolution number n_p [rps], the bow thruster revolution number n_{BT} [rps], and the stern thruster revolution number n_{ST} [rps]. Their feasible values are in $U = [-\frac{35}{180}\pi, \frac{35}{180}\pi] \times [-20, 20] \times [-20, 20] \times [-20, 20]$. We employ the state equation model $\dot{s} = \phi(s, a; y)$ proposed in Miyauchi et al. [2022], where $y \in \mathbb{Y}$ represents the model uncertainty described in Section 6.

Feedback Controller. The feedback controller $u_x : \mathbb{R}^6 \rightarrow U$ is modeled by the following neural network parameterized by $x = (B, W, V)$:

$$u_x(s) = V \cdot \text{softmax}(\alpha \cdot (B + W \cdot s)), \quad (37)$$

where $W \in [-1, 1]^{K \times 6}$ and $B \in [-1, 1]^K$ define a linear map $z = \alpha \cdot (B + W \cdot s)$ from the state vector s to the K -dimensional latent space, and $V \in U^K \subset \mathbb{R}^{m \times K}$ is a matrix consisting of K feasible control vectors as its columns. The softmax function

$$\text{softmax} : z = (z_1, \dots, z_K) \mapsto \frac{(\exp(z_1), \dots, \exp(z_K))}{\sum_{k=1}^K \exp(z_k)} \in \Delta^{K-1}$$

outputs a point in the $K - 1$ -dimensional standard simplex $\Delta^{K-1} = \{z \in \mathbb{R}^K : z_1 \geq 0, \dots, z_K \geq 0, \text{ and } \sum_{k=1}^K z_k = 1\}$. The output is a combination of the columns of V weighted by the softmax output. $\alpha > 0$ is a parameter that determines whether the output of softmax is close to the one-hot vector.

The architecture of this neural network is interpreted as follows. First, $z = \text{softmax}(\alpha \cdot (B + W \cdot s))$ on the first layer divides the state space into K regions. For example, if the greatest element of the

vector $B + W \cdot s$ is the k th coordinate, then z is approximated by the one-hot vector e_k with 1 on the k th coordinate and 0 on the other coordinates if α is sufficiently large. In such a situation, $u(s) = V \cdot z \approx V \cdot e_k = v_k$, where v_k is the k th column of V . In other words, this neural network approximates the control law that divides the state space using a Voronoi diagram with respect to the Euclidean metric and outputs the corresponding column of V as a control signal in each region. If we set α to be greater, then z is more likely to be close to a one-hot vector, which makes it easier to express the bang-bang type control. If we set α to be smaller, then z is more likely to take a value in the middle of Δ^{K-1} , which makes it easier to express a continuous control.

Based on our preliminary experiments, we set $\alpha = 4$ and $K = 9$ in the following experiments. Then, x is of $m = 99$ dimension.

Objective Function. The objective is to find the parameter $x := (B, W, V) \in \mathbb{X}$ of the controller u_x that minimizes the cost C of the trajectory $(s_{t \in [0, t_{\max}]}, a_{t \in [0, t_{\max}]})$ in the worst environment $y \in \mathbb{Y}$ for u_x . It is modeled as

$$\begin{aligned} \min_{x \in \mathbb{X}} \max_{y \in \mathbb{Y}} f(x, y) &= \min_{x \in \mathbb{X}} \max_{y \in \mathbb{Y}} C(s_{t \in [0, t_{\max}]}, a_{t \in [0, t_{\max}]}, \\ \text{subject to } s_t &= s_0 + \int_0^t \phi(s_\tau, a_\tau; y) d\tau \quad \text{and} \quad a_t = u_x(s_{\lfloor t/dt \rfloor \cdot dt}), \end{aligned}$$

where dt [seconds] is the control time span, that is, the control signal a_t changes every dt , and s_0 is the initial state.

We define the cost of the trajectory as

$$C(s_{t \in [0, t_{\max}]}, a_{t \in [0, t_{\max}]}) = C_1 + w \cdot (C_2 + \mathbb{I}\{C_2 > 0\}),$$

where $w > 0$ is the hyperparameter that determines the trade-off between utility and safety,

$$C_1 = \frac{1}{6} \sum_{i=1}^6 (s_{t_{\max}, i} - s_{\text{fin}, i})^2,$$

evaluates the deviation of the final ship state from the target state s_{fin} , and

$$C_2 = \frac{1}{4} \sum_{i=1}^4 \int_0^{t_{\max}} \text{dist}(P_{\tau, i}, C_{\text{berth}}) d\tau$$

measures the collision risk, where $P_{\tau, 1}, \dots, P_{\tau, 4}$ represents the coordinates of the four vertices of the rectangle surrounding the ship at time τ and $\text{dist}(P, C_{\text{berth}})$ measures the distance from a point P to the closest point on the berth boundary. Refer to Maki et al. [2021, 2020] for the definitions of C_1 and C_2 .

Following Maki et al. [2021], we set $t_{\max} = 200$ [s] and $dt = 10$ [s]. The initial state is $s_0 = (15.0, 0.01, 6.0, 0.0, \pi, 0.0)$ and the target state is $s_{\text{fin}} = (3.0, 0.0, 9.5, 0.0, \pi, 0.0)$. The boundary of the berth is $C_{\text{berth}} = \{Y = 9.994625\}$. The trade-off coefficient is set to $w = 10$. That is, the cost $f(x, y) < 10$ implies that the controller u_x produces a trajectory without collision with the berth under the uncertainty parameter y .

Differences from Previous Works. Our problem formulation mostly follows previous studies [Maki et al. 2021, 2020] but with certain differences. First, we optimize the feedback controller, whereas the control signals for each time period as well as the total control time are directly optimized in Maki et al. [2021, 2020], which we believe is not suitable for obtaining robust control. Second, we modify the objective function. Previous studies include the term penalizing the control time as they formulate the problem as minimization of the control time. Because we did not optimize the control time, it is excluded from our objective function definition. Moreover, for better

Fig. 6. Trajectories of the best controller obtained by CMA-ES($y_{\text{no}}^{(A)}$). Fig. 7. Trajectories of the best controller obtained by Adversarial-CMA-ES on $\mathbb{Y}_A$.

collision avoidance, we replaced $w \cdot C_2$ with $w \cdot (C_2 + \mathbb{I}\{C_2 > 0\})$. Third, following Miyauchi et al. [2022], we implement thrusters to realize robust control under external disturbances and adopt the state equation model proposed in Miyauchi et al. [2022].

C ADDITIONAL RESULTS FOR AUTOMATIC BERTHING CONTROL PROBLEM

Figures 6 to 9 visualize the trajectories obtained in the experiments in Section 6. The route of the ship, that is, (X, Y, ψ) at each time, is displayed in the top figure. The X and Y axes are scaled

Fig. 8. Trajectories of the best controller obtained by CMA-ES($y_{\text{est}}^{(B)}$). Fig. 9. Trajectories of the best controller obtained by Adversarial-CMA-ES on $\mathbb{Y}_B$.

by $L_{pp} = 3$ [m]. The changes in the velocities, (u, v_m, r) , as well as the changes in the control signals, $(\delta, n_p, n_{BT}, n_{ST})$, are plotted at the bottom. Note that r and δ are plotted on a degree basis for better intuition. Figure 6 shows the trajectories observed for the best controller obtained by CMA-ES($y_{\text{no}}^{(A)}$), which is the controller optimized under $y = (y_{\text{no}}^{(A)}, y_{\text{est}}^{(B)})$, that is, no wind $y^{(A)} = y_{\text{no}}^{(A)}$ and model parameter $y^{(B)} = y_{\text{est}}^{(B)}$ used in the previous study. Figure 7 shows the trajectories observed for the best controller obtained by Adversarial-CMA-ES on $\mathbb{Y}_A$, which is the controller

optimized under the worst wind condition $y^{(A)} \in \mathbb{Y}_A$ with $y^{(B)} = y_{\text{est}}^{(B)}$. For Figures 6 and 7, the left figure is the trajectory under $y = (y_{\text{no}}^{(A)}, y_{\text{est}}^{(B)})$ and the right figure is the trajectory under the worst wind condition $y^{(A)} \in \mathbb{Y}_A$ with $y^{(B)} = y_{\text{est}}^{(B)}$. Figure 8 shows the trajectories observed for the best controller obtained by CMA-ES($y_{\text{est}}^{(B)}$), which is the controller optimized under $y = (y_{\text{est}}^{(A)}, y_{\text{est}}^{(B)})$. Figure 9 shows the trajectories observed for the best controller obtained by Adversarial-CMA-ES on $\mathbb{Y}_B$, which is the controller optimized under the worst model parameter $y^{(B)} \in \mathbb{Y}_B$ with wind condition $y^{(A)} = y_{\text{est}}^{(A)}$. For Figures 8 and 9, the left figure is the trajectory under $y = (y_{\text{est}}^{(A)}, y_{\text{est}}^{(B)})$ and the right figure is the trajectory under the worst model parameter $y^{(B)} \in \mathbb{Y}_B$ with $y^{(A)} = y_{\text{est}}^{(A)}$.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their valuable comments and suggestions.

REFERENCES

- Martin A. Abkowitz. 1980. Measurement of hydrodynamic characteristics from ship maneuvering trials by system identification. In *Transactions of Society of Naval Architects and Marine Engineers* 88, 283–318.
- Leonard Adolphs, Hadi Daneshmand, Aurelien Lucchi, and Thomas Hofmann. 2019. Local saddle point optimization: A curvature exploitation approach. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*. 486–495.
- Youhei Akimoto. 2021. Saddle point optimization with approximate minimization oracle. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'21)*. 493–501.
- Youhei Akimoto and Nikolaus Hansen. 2020. Diagonal acceleration for covariance matrix adaptation evolution strategies. *Evolution. Comput.* 28, 3 (2020), 405–435.
- Abdullah Al-Dujaili, Shashank Srikant, Erik Hemberg, and Una-May O'Reilly. 2019. On the application of danskin's theorem to derivative-free minimax problems. In *Proceedings of the AIP Conference*, Vol. 2070. 20–26.
- Motoki Araki, Hamid Sadat-Hosseini, Yugo Sanada, Kenji Tanimoto, Naoya Umeda, and Frederick Stern. 2012. Estimating maneuvering coefficients using system identification methods with experimental, system-based, and CFD free-running trial data. *Ocean Eng.* 51 (2012), 63–84.
- Dirk V. Arnold and Nikolaus Hansen. 2010. Active covariance matrix adaptation for the (1+1)-CMA-ES. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation (GECCO'10)*. 385–392.
- Dimitris Bertsimas, Omid Nohadani, and Kwong Meng Teo. 2010a. Nonconvex robust optimization for problems with constraints. *INFORMS J. Comput.* 22, 1 (2010), 44–58.
- Dimitris Bertsimas, Omid Nohadani, and Kwong Meng Teo. 2010b. Robust nonconvex optimization for simulation based problems. *Oper. Res.* 58, 1 (2010), 161–178.
- Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. 2018. Adversarially robust optimization with gaussian processes. In *Advances in Neural Information Processing Systems*. MIT Press, 5760–5770.
- Stephen Boyd and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press.
- Jürgen Branke and Johanna Rosenbusch. 2008. New approaches to coevolutionary worst-case optimization. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*. 144–153.
- Ashish Cherukuri, Bahman Ghahsifard, and Jorge Cortés. 2017. Saddle-point dynamics: Conditions for asymptotic stability of saddle points. *SIAM J. Control Optimiz.* 55, 1 (2017), 486–511.
- Andrew R. Conn and Luis N. Vicente. 2012. Bilevel derivative-free optimization and its application to robust optimization. *Optimiz. Methods Softw.* 27, 3 (2012), 561–577.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. 2018. Training GANs with optimism. In *Proceedings of the International Conference on Learning Representations*.
- Oswaldo de Oliveira. 2013. The implicit and inverse function theorems: Easy proofs. *Real Anal. Exchange* 39, 1 (2013), 207–218.
- Luc Devroye. 1972. The compound random search. In *Proceedings of the International Symposium on Systems Engineering and Analysis*. 195–110.
- Peter I. Frazier. 2018. A tutorial on bayesian optimization. Retrieved from <https://arXiv:1807.02811>.
- Toshifumi Fujiwara, Michio Ueno, and Tadashi Nimura. 1998. Estimation of wind forces and moments acting on ships. *J. Soc. Naval Arch. Japan* 1998 (1998), 77–90. Issue 183.
- Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. 2017. Frank-wolfe algorithms for saddle point problems. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*. 362–371.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. MIT Press, 2672–2680.
- Nikolaus Hansen and Anne Auger. 2014. Principled design of continuous stochastic search: From theory to practice. In *Theory and Principled Methods for the Design of Metaheuristics*. Springer, 145–180.
- Nikolaus Hansen, Sibylle D. Müller, and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolution. Comput.* 11, 1 (2003), 1–18.
- Nikolaus Hansen and Andreas Ostermeier. 2001. Completely derandomized self-adaptation in evolution strategies. *Evolution. Comput.* 9, 2 (2001), 159–195.
- Mikkel T. Jensen. 2004. *A New Look at Solving Minimax Problems with Coevolutionary Genetic Algorithms*. Kluwer Academic Publishers, 369–384.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. 2016. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. 795–811.
- Stefan Kern, Sibylle D. Müller, Nikolaus Hansen, Dirk Büche, Jiri Ocenasek, and Petros Koumoutsakos. 2004. Learning probability distributions in continuous evolutionary algorithms—A comparative review. *Natural Comput.* 3, 1 (2004), 77–112.
- Dieter Kraft. 1988. *A Software Package for Sequential Quadratic Programming*. Technical Report. DFVLR-FB 88-28, DLR German Aerospace Center—Institute for Flight Mechanics, Köln, Germany.
- Tengyuan Liang and James Stokes. 2019. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*. 907–915.
- Sijia Liu, Songtao Lu, Xiangyi Chen, Yao Feng, Kaidi Xu, Abdullah Al-Dujaili, Mingyi Hong, and Una-May O’Reilly. 2020. Min-max optimization without gradients: Convergence and applications to black-box evasion and poisoning attacks. In *Proceedings of the International Conference on Machine Learning*. 2307–2318.
- Atsuo Maki, Youhei Akimoto, and Naoya Umeda. 2021. Application of optimal control theory based on the evolution strategy (CMA-ES) to automatic berthing (part: 2). *J. Marine Sci. Technol.* 26 (2021), 835–845.
- Atsuo Maki, Naoki Sakamoto, Youhei Akimoto, Hiroyuki Nishikawa, and Naoya Umeda. 2020. Application of optimal control theory based on the evolution strategy (CMA-ES) to automatic berthing. *J. Marine Sci. Technol.* 25 (2020), 221–233.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. 2017. The numerics of GANs. In *Advances in Neural Information Processing Systems*. MIT Press, 1823–1833.
- Transport Ministry of Land, Infrastructure and Tourism. 2020. White paper on land, infrastructure, transport and tourism in Japan. Retrieved from <https://www.mlit.go.jp/en/statistics/white-paper-mlit-index.html>.
- Yoshiki Miyauchi, Atsuo Maki, Naoya Umeda, Dimas M. Rachman, and Youhei Akimoto. 2021. System parameter exploration of ship maneuvering model for automatic docking / berthing using CMA-ES. Retrieved from <https://arXiv:2111.06124>.
- Yoshiki Miyauchi, Ryohei Sawada, Youhei Akimoto, Naoya Umeda, and Atsuo Maki. 2022. Optimization on planning of trajectory and control of autonomous berthing and unberthing for the realistic port geometry. *Ocean Engineering* 245 (2022), 110390.
- Daiki Morinaga and Youhei Akimoto. 2019. Generalized drift analysis in continuous domain: Linear convergence of (1+1)-ES on strongly convex functions with lipschitz continuous gradients. In *Foundations of Genetic Algorithms*. 13–24.
- Daiki Morinaga, Kazuto Fukuchi, Jun Sakuma, and Youhei Akimoto. 2021. Convergence rate of the (1+1)-evolution strategy with success-based step-size adaptation on convex quadratic function. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO’21)*. 1169–1177.
- Vaishnavh Nagarajan and J. Zico Kolter. 2017. Gradient descent GAN optimization is locally stable. In *Advances in Neural Information Processing Systems*. MIT Press, 5591–5600.
- Yurii Nesterov and Vladimir Spokoiny. 2017. Random gradient-free minimization of convex functions. *Found. Comput. Math.* 17, 2 (2017), 527–566.
- Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. 2019. Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*. MIT Press, 14934–14942.
- Victor Picheny, Mickael Binois, and Abderrahmane Habbal. 2019. A bayesian optimization approach to find Nash equilibria. *J. Global Optimiz.* 73, 1 (2019), 171–192.
- Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust adversarial reinforcement learning. In *Proceedings of the International Conference on Machine Learning*. 2817–2826.
- Xin Qiu, Jian-Xin Xu, Yinghao Xu, and Kay Chen Tan. 2018. A new differential evolution algorithm for minimax optimization in robust design. *IEEE Trans. Cybernet.* 48, 5 (2018), 1355–1368.

- Ingo Rechenberg. 1973. *Evolutionsstrategie: Optimierung Technischer Systeme nach Prinzipien Der Biologischen Evolution*. Frommann-Holzboog.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. 2016. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*. MIT Press, 2234–2242.
- Michael A. Schumer and Kenneth Steiglitz. 1968. Adaptive step size random search. *IEEE Trans. Autom. Control* 13 (1968), 270–276.
- Hiroaki Shioya, Yusuke Iwasawa, and Yutaka Matsuo. 2018. Extending robust adversarial reinforcement learning considering adaptation and diversity. In *Proceedings of the International Conference on Learning Representations, Workshop Track*.
- Kouki Wakita, Atsuo Maki, Naoya Umeda, Yoshiki Miyauchi, Tohga Shimoji, Dimas M. Rachman, and Youhei Akimoto. 2022. On neural network identification for low-speed ship maneuvering model. *Journal of Marine Science of Technology* (2022).
- Abhay Yadav, Sohil Shah, Zheng Xu, David Jacobs, and Tom Goldstein. 2018. Stabilizing adversarial nets with prediction methods. In *Proceedings of the International Conference on Learning Representations*.
- Aimin Zhou and Qingfu Zhang. 2010. A surrogate-assisted evolutionary algorithm for minimax optimization. In *Proceedings of the IEEE Congress on Evolutionary Computation*. 1–7.

Received May 2021; revised November 2021; accepted January 2022