Revisiting norm optimization for multi-objective black-box problems : a finite-time analysis

Al-Dujaili, Abdullah; Suresh, Sundaram

2018

Al-Dujaili, A., & Suresh, S. (2019). Revisiting norm optimization for multi-objective black-box problems : a finite-time analysis. *Journal of Global Optimization*, 73(3), 659-673.
doi:10.1007/s10898-018-0709-z

<https://hdl.handle.net/10356/143267>

<https://doi.org/10.1007/s10898-018-0709-z>

© 2018 Springer. This is a post-peer-review, pre-copyedit version of an article published in *Journal of Global Optimization*. The final authenticated version is available online at:
<http://dx.doi.org/10.1007/s10898-018-0709-z>.

Downloaded on 26 Apr 2023 23:36:31 SGT

Revisiting Norm Optimization for Multi-Objective Black-Box Problems: A Finite-Time Analysis

Abdullah Al-Dujaili · S. Suresh

Received: date / Accepted: date

Abstract The complexity of Pareto fronts imposes a great challenge on the convergence analysis of multi-objective optimization methods. While most theoretical convergence studies have addressed finite-set and/or discrete problems, others have provided probabilistic guarantees, assumed a total order on the solutions, or studied their asymptotic behaviour. In this paper, we revisit the Tchebycheff weighted method in a hierarchical bandits setting and provide a finite-time bound on the Pareto-compliant additive ϵ -indicator. To the best of our knowledge, this paper is one of few that establish a link between weighted sum methods and quality indicators in finite time.

Keywords Multi-objective optimization · Black-box optimization · Derivative-free optimization · Finite-time analysis

Problem. This paper is concerned with the Multi-Objective Black-Box Optimization (MOBBO) problem given a finite number of function evaluations. With n decision variables and m objectives, the problem has the mathematical form:

$$(1) \quad \begin{aligned} \text{minimize} \quad & \mathbf{y} = \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_m(\mathbf{x})) \\ \text{where} \quad & \mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X} \subset \mathbb{R}^n \\ & \mathbf{y} = (y_1, \dots, y_m) \in \mathcal{Y} \subset \mathbb{R}^m, \end{aligned}$$

where $\mathbf{x}$ is called the *decision vector (solution)*, $\mathbf{y}$ is called the *objective vector*, $\mathcal{X}$ is the *feasible decision space*, and $\mathcal{Y} = \times_{1 \leq j \leq m} \mathcal{Y}_j$ is the reachable *objective space*, where $\mathcal{Y}_j$ is the j th-objective space. By black-box, we mean to say that there is no closed-form expression of $\mathbf{f}$ and that its derivatives are neither symbolically nor numerically available. However, $\mathbf{f}$ can be evaluated point-wise, but each evaluation is typically expensive in terms of computational resources (e.g., time, power, money).

In practice, the objective functions $\{f_j\}_{1 \leq j \leq m}$ are conflictual. Therefore, the problem may have a set of incomparable optimal solutions: each is inferior to the other in some objectives and superior in other objectives, inducing a partial order on the set of feasible solutions.

Abdullah Al-Dujaili
CSAIL, MIT, Cambridge, MA 02139, USA
E-mail: aldujail@mit.edu

S. Suresh
School of Computer Engineering, Nanyang Technological University, Singapore 639798
E-mail: ssundaram@ntu.edu.sg

Related Work. The notion of Pareto-optimality was first introduced in the field of engineering in 1970 [17]. The MOBBO solvers can be broadly classified into generative or preference methods based on the decision maker’s role. The former does not require any inputs from the decision maker in solving the problem (only selects solution at the end), whereas preference methods require input from decision maker at the beginning. The knowledge of decision maker may affect the solution quality in the preference method with regard to some objectives, but it reduces the complexity in MOBBO solver. Conventionally, most commonly used MOBBO solvers convert multiple objectives into a single (or a series of) objective optimization problem [11, 19]. The adaptive scalarization [8] overcomes the problem of non-uniform distribution of optimal solution and the non-convex Pareto fronts in scalarizing approaches [6].

With regard to black-box optimization, optimistic algorithms—whose foundations come from the multi-armed bandit theory [1, 3, 4, 14]—consider partitions of the search space at multiple scales in search of optimal solutions. These methods enjoy provable finite-time performance and asymptotic convergence. On the other hand, research works with the Tchebycheff metric have lacked theoretical guarantees [18]. To this end, this paper aims to bridge the gap in understanding the theoretical underpinnings of the weighted Tchebycheff method and link the convergence of scalarization methods to Pareto-compliant quality indicators in a multi-arm bandits setting.

Our Contributions. This paper addresses a class of weighted sum methods for MOBBO problems. While most of the literature work has established the asymptotic optimality of weighted sum methods to a single Pareto-optimal solution under certain conditions, our theoretical contributions here are of two-fold. First, we show that the weighted Tchebycheff problem for Lipschitz MOBBO is as well Lipschitz. Second, we present a finite-time upper bound on the Pareto-compliant quality indicator of the approximation set obtained from solving the weighted Tchebycheff problem capturing its convergence to the whole Pareto front. All of this is motivated by the success of the optimism in the face of uncertainty principle that helps us employ an optimistic method, which we refer to as the Weighted Optimistic Optimization (WOO) algorithm. The sequential decision-making approach in WOO formulates the weighted decision space $\mathcal{X}$ as a hierarchy of simple bandit problems over subspaces of $\mathcal{X}$ and looks for the optimal solution through $\mathcal{X}$ -partitioning search trees. At each step, WOO expands the subspace which may contain the optimum. Based on smoothness assumption, the convergence analysis is presented in finite time and validated empirically using a set of synthetic problems.

Paper Organization. The rest of the paper is organized as follows. First, a brief introduction to basic concepts in multi-objective optimization is provided, along with the notations and terminology used through out the paper. Then, we motivate treating the weighted Tchebycheff problem in a multi-arm bandits setting by proving its smoothness. This is followed by introducing the weighted optimistic optimization algorithm that exploits the smoothness without the need for its knowledge. Furthermore, theoretical and empirical analysis of the proposed algorithm is presented. Towards the end, we conclude with a discussion on potential future research investigations.

1 Formal Background

This section presents basic concepts in multi-objective optimizations. First, the notion of Pareto dominance is described. Second, approaches to assess the performance of multi-objective solvers are discussed. Third, we formally define the weighted Tchebycheff problem.

1.1 Pareto Dominance

An objective vector $\mathbf{y}^1 \in \mathcal{Y}$ is more preferable than another vector $\mathbf{y}^2 \in \mathcal{Y}$, if $\mathbf{y}^1$ is at least as good as $\mathbf{y}^2$ in all objectives *and* better with respect to at least one objective. $\mathbf{y}^1$ is then said to be *dominating* $\mathbf{y}^2$. This notion of dominance is commonly known as *Pareto dominance* [16], which leads to a *partial order* on the objective space, where we can define a Pareto optimal vector to be one that is non-dominated by any other vector in $\mathcal{Y}$. Nevertheless, $\mathbf{y}^1$ and $\mathbf{y}^2$ may be incomparable to each other, because each is inferior to the other in some objectives and superior in other objectives. Hence, there can be several Pareto optimal vectors. This concept is presented in the following definitions [10, 20].

Definition 1 (Pareto dominance) The vector $\mathbf{y}^1$ dominates the vector $\mathbf{y}^2$, that is to say, $\mathbf{y}^1 \prec \mathbf{y}^2 \iff y_j^1 \leq y_j^2$ for all $j \in \{1, \dots, m\}$ and $y_k^1 < y_k^2$ for at least one $k \in \{1, \dots, m\}$.

Definition 2 (Strict Pareto dominance) The vector $\mathbf{y}^1$ strictly dominates the vector $\mathbf{y}^2$ if $\mathbf{y}^1$ is better than $\mathbf{y}^2$ in all the objectives, that is to say, $\mathbf{y}^1 \prec \mathbf{y}^2 \iff y_j^1 < y_j^2$ for all $j \in \{1, \dots, m\}$.

Definition 3 (Weak Pareto dominance) The vector $\mathbf{y}^1$ weakly dominates the vector $\mathbf{y}^2$ if $\mathbf{y}^1$ is not worse than $\mathbf{y}^2$ in all the objectives, that is to say, $\mathbf{y}^1 \preceq \mathbf{y}^2 \iff y_j^1 \leq y_j^2$ for all $j \in \{1, \dots, m\}$.

Definition 4 (Pareto optimality of vectors) Let $\hat{\mathbf{y}} \in \mathcal{Y}$ be a vector. $\hat{\mathbf{y}}$ is Pareto optimal $\iff \nexists \mathbf{y} \in \mathcal{Y}$ such that $\mathbf{y} \prec \hat{\mathbf{y}}$. The set of all Pareto optimal vectors is referred to as the Pareto front and denoted as $\mathcal{Y}^*$. The corresponding decision vectors (solutions) are referred to as the Pareto optimal solutions or the Pareto set and denoted by $\mathcal{X}^*$.

Thus, the solution to the MOBBO problem (1) is its Pareto optimal solutions (Pareto front in the objective space). Practically, MOBBO solvers aim to identify a set of objective vectors that represent the Pareto front (or a good approximation of it). We refer to this set as the *approximation set*.

Definition 5 (Approximation set) Let $A \subseteq \mathcal{Y}$ be a set of objective vectors. A is called an approximation set if any element of A does not dominate or is not equal to any other objective vector in A . The set of all approximation sets is denoted as Ω . Note that $\mathcal{Y}^* \in \Omega$.

Furthermore, denote the *ideal point (utopian vector)* by $\mathbf{y}^* \stackrel{\text{def}}{=} (\min_{\mathbf{y} \in \mathcal{Y}^*} y_1, \dots, \min_{\mathbf{y} \in \mathcal{Y}^*} y_m)$. Likewise, let us denote the (or one of the) global optimizer(s) of the j th objective function by $\mathbf{x}_j^*$, i.e., $y_j^* = f_j(\mathbf{x}_j^*)$. Note that $\mathbf{x}_j^* \in \mathcal{X}^*$. Without loss of generality, we assume that $\mathbf{y}^*$ is the zero vector.

1.2 Performance Assessment for Multi-Objective Optimization Methods

Given two approximation sets $A, B \in \Omega$, it is not that easy to tell which set is better, particularly if their elements are incomparable [20]. In general, two aspects are considered in an approximation set: i). its distance (the closer the better) to the optimal Pareto front and ii). its diversity (the higher the better) within the optimal Pareto front. To this end, several *quality indicators* have been proposed [9]. The quality of an approximation set is measured by a so-called (unary) quality indicator $I : \Omega \rightarrow \mathbb{R}$, assessing a specific property of the approximation set. Likewise, an l -ary quality indicator $I : \Omega^l \rightarrow \mathbb{R}$ quantifies quality differences between l approximation sets [5, 20]. A quality indicator is not Pareto-compliant if it contradicts the order induced by the Pareto-dominance relations described in Section 1. One commonly-used quality indicators is the Pareto-compliant additive ϵ -indicator, which is defined formally next.

Definition 6 (Additive ϵ -indicator [20]) For any two approximation sets $A, B \in \Omega$, the additive ϵ -indicator $I_{\epsilon+}$ is defined as:

$$(2) \quad I_{\epsilon+}(A, B) = \inf_{\epsilon \in \mathbb{R}} \{ \forall \mathbf{y}^2 \in B, \exists \mathbf{y}^1 \in A : \mathbf{y}^1 \preceq_{\epsilon+} \mathbf{y}^2 \}$$

where $\mathbf{y}^1 \preceq_{\epsilon+} \mathbf{y}^2 \iff y_j^1 \leq \epsilon + y_j^2$ for all $j \in \{1, \dots, m\}$. If B is the Pareto front $\mathcal{Y}^*$ (or a good—in terms of diversity and closeness to the Pareto front—approximation reference set $R \in \Omega$ if $\mathcal{Y}^*$ is unknown) then $I_{\epsilon+}(A, B)$ is referred to as the unary additive epsilon indicator and is denoted by $I_{\epsilon+}^1(A)$, i.e., $I_{\epsilon+}^1(A) \stackrel{\text{def}}{=} I_{\epsilon+}(A, \mathcal{Y}^*)$.

In essence, $I_{\epsilon+}^1(A)$ measures the smallest amount ϵ needed to translate each element in the Pareto front $\mathcal{Y}^*$ such that it is *weakly* dominated by at least one element in the approximation set A . Note that $I_{\epsilon+}^1(A) \stackrel{\text{def}}{=} I_{\epsilon+}(A, \mathcal{Y}^*) \geq 0$ as no element in A strictly dominates any element in $\mathcal{Y}^*$. Thus, the closer $I_{\epsilon+}^1(A)$ to 0, the better the quality of A .

1.3 Weighted Sum Methods

In weighted sum methods, the idea is to assign a non-negative weight value w for each objective and minimize the weighted sum of the objectives. Denote the *element-wise product* or *Hadamard product* of two vectors $\mathbf{a}$ and $\mathbf{b}$ of the same dimensionality by $\mathbf{a} \odot \mathbf{b}$, and the l_p -norm of a vector $\mathbf{a}$ by $\|\mathbf{a}\|_p$, then weighted sum methods have the mathematical form $\|\mathbf{w} \odot |\mathbf{f}(\mathbf{x}) - \mathbf{z}^*|\|_p$. As the focus of this paper is the weighted Tchebycheff problem, we define it formally next.

Definition 7 (Weighted Tchebycheff problem) Let $\mathbf{w} \in \mathbb{R}^m$ be a non-negative vector, and $\mathbf{z}^* \in \mathbb{R}^m$ be a reference point. Then, the weighted Tchebycheff formulation of problem (1) is defined as:

$$(3) \quad \begin{aligned} \text{minimize} \quad & g_{\mathbf{w}}(\mathbf{x}) = \max_{1 \leq j \leq m} w_j |f_j(\mathbf{x}) - z_j^*| \\ & = \|\mathbf{w} \odot |\mathbf{f}(\mathbf{x}) - \mathbf{z}^*|\|_{\infty} \\ \text{where} \quad & \mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X} \subset \mathbb{R}^n. \end{aligned}$$

Under certain conditions, the solution to problem (3) corresponds to a Pareto-optimal solution of problem (1). This is stated in the following theorem.

Theorem 1 [13, Theorem 3.4.5]. Let $\mathbf{x}$ be a Pareto-optimal solution of problem (1), then there exists a positive weighting vector $\mathbf{w}$ such that $\mathbf{x}$ is a solution of the weighted Tchebycheff problem (3) where the reference point is the utopian objective vector $\mathbf{y}^*$.

Proof. See [13, Page 98]. ■

Similar to problem (1), the weighted Tchebycheff problem (3) is black-box but with a single objective. Therefore, using a computational budget of $v(t)$ function evaluations, we would like to devise an algorithm that searches the decision space $\mathcal{X}$ over t iterations in a *sequential decision making framework*, where each sample $\mathbf{x}^k$ may depend on the previous sampled points and their corresponding function values $\{(\mathbf{x}^i, g_{\mathbf{w}}(\mathbf{x}^i))\}_{1 \leq i \leq k-1}$. After the final iteration t of the algorithm, we can obtain the sampled point with the best possible $g_{\mathbf{w}}$ value:

$$(4) \quad \mathbf{x}(t) \in \arg \min_{1 \leq i \leq v(t)} g_{\mathbf{w}}(\mathbf{x}^i).$$

Also, with regard to the corresponding multi-objective problem at hand, we can obtain an approximation set to the Pareto front as the set of the non-dominated sampled points:

$$(5) \quad \mathcal{Y}_*^t = \{\mathbf{f}(\mathbf{x}^i) \mid i = 1, \dots, v(t), \mathbf{f}(\mathbf{x}^k) \not\leq \mathbf{f}(\mathbf{x}^i), \\ \forall k = 1, \dots, v(t), k \neq i\} \in \Omega.$$

2 The Weighted Tchebycheff Problem: A Multi-Armed Bandit View

In the previous section, we formulated the weighted Tchebycheff problem as searching the decision space in a sequential decision making framework, where the present sample may depend on the previous observed samples. One popular principle of such framework is the *optimism in the face of uncertainty* principle, which suggests following the optimal strategy with respect to the most favorable scenario among all possible scenarios that are compatible with the obtained observations about the problem at hand [15]. It has been the main principle in multi-armed bandits settings and was later extended to many (including infinite) arms under probabilistic or structural (smoothness) assumptions about the arm rewards. In this section, we approach problem (3) in a multi-armed bandits setting assuming a structural smoothness. Our smoothness assumption is motivated by the next theorem, which shows that the Tchebycheff function is Lipschitz-continuous if the corresponding objectives are.

Theorem 2 *The Tchebycheff function for an optimization problem with m Lipschitz-continuous objectives whose constants are $\{L_j\}_{1 \leq j \leq m}$, respectively, is also Lipschitz-continuous with the Lipschitz constant being less than or equal to $\sqrt{m} \cdot \max_j L_j$.*

Proof. As the objectives of the problem are Lipschitz continuous, for any $\mathbf{x}^1$ and $\mathbf{x}^2 \in \mathcal{X}$, we have

$$|f_j(\mathbf{x}^1) - f_j(\mathbf{x}^2)| \leq L_j \|\mathbf{x}^1 - \mathbf{x}^2\|_2, \quad j = 1, \dots, m.$$

Squaring both sides of the above inequality can be expressed in vector notation as follows.

$$\begin{aligned} \|\mathbf{f}(\mathbf{x}^1) - \mathbf{f}(\mathbf{x}^2)\|_2^2 &= \sum_{j=1}^m L_j^2 \|\mathbf{x}^1 - \mathbf{x}^2\|_2^2 \\ &\leq m \cdot \max_j L_j^2 \|\mathbf{x}^1 - \mathbf{x}^2\|_2^2. \end{aligned}$$

Since the l_∞ -norm of any vector is bounded by its l_2 -norm, it follows that

$$(6) \quad \|\mathbf{f}(\mathbf{x}^1) - \mathbf{f}(\mathbf{x}^2)\|_\infty \leq \sqrt{m} \cdot \max_j L_j \|\mathbf{x}^1 - \mathbf{x}^2\|_2.$$

Thus, the absolute difference in the weighted Tchebycheff function $g_{\mathbf{w}}(\mathbf{x})$ values at any two vectors $\mathbf{x}^1$ and $\mathbf{x}^2 \in \mathcal{X}$ for problem (1) can be bounded as below.

$$\begin{aligned} |g_{\mathbf{w}}(\mathbf{x}^1) - g_{\mathbf{w}}(\mathbf{x}^2)| &= |||\mathbf{w} \odot |\mathbf{f}(\mathbf{x}^1) - \mathbf{z}^*|||_\infty \\ &\quad - |||\mathbf{w} \odot |\mathbf{f}(\mathbf{x}^2) - \mathbf{z}^*|||_\infty| \\ &\leq |||\mathbf{w} \odot |\mathbf{f}(\mathbf{x}^1) - \mathbf{z}^*| \\ &\quad - \mathbf{w} \odot |\mathbf{f}(\mathbf{x}^2) - \mathbf{z}^*|||_\infty. \end{aligned}$$

To guarantee Pareto-optimal solutions to problem (1), Theorem 1 suggests that $\mathbf{z}^*$ should be the utopian vector $\mathbf{y}^*$, which is—without loss of generality—the zero vector based on Section 1.1. Therefore, we have

$$\begin{aligned} |g_{\mathbf{w}}(\mathbf{x}^1) - g_{\mathbf{w}}(\mathbf{x}^2)| &\leq \|\mathbf{w} \odot (\mathbf{f}(\mathbf{x}^1) - \mathbf{f}(\mathbf{x}^2))\|_{\infty} \\ &\leq \|\mathbf{w}\|_{\infty} \|\mathbf{f}(\mathbf{x}^1) - \mathbf{f}(\mathbf{x}^2)\|_{\infty} \\ &\leq \sqrt{m} \cdot \max_j w_j L_j \|\mathbf{x}^1 - \mathbf{x}^2\|_2. \end{aligned}$$

Therefore, $g_{\mathbf{w}}$'s Lipschitz constant is less than or equal $\sqrt{m} \cdot \max_j w_j L_j$. $\blacksquare$

Theorem 2 established the Lipschitz continuity of the weighted Tchebycheff problem provided that all objectives are Lipschitz-continuous. This nicely supports our smoothness assumption and justifies following the optimism in the face of uncertainty principle.

Nevertheless, the aforementioned theorem does not tell us about the quality of problem (7)'s solution $\mathbf{x}(t)$ and its relation to problem (1)'s approximation set $\mathcal{Y}_*^t$. Given $v(t)$ function evaluations, our next theorem upper bounds the quality of the approximation set $\mathcal{Y}_*^t$ obtained from all points sampled in t iterations to solve problem (3) by the best value $g_{\mathbf{w}}$ found of these samples.

Theorem 3 $I_{\epsilon+}^1(\mathcal{Y}_*^t) \leq \max_j \frac{1}{w_j} g_{\mathbf{w}}(\mathbf{x}(t))$.

Proof. From Definition 7 and (4), the l_{∞} -norms of the weighted objective vectors $\{\|\mathbf{w} \odot \mathbf{y}^i\|_{\infty}\}_{1 \leq i \leq v(t)}$ of the sampled points $\{\mathbf{x}^i\}_{1 \leq i \leq v(t)}$ are greater than or equal $g_{\mathbf{w}}(\mathbf{x}(t))$ (note that the ideal point $\mathbf{y}^*$ is the zero vector). This suggest that there exists one vector $\mathbf{y}^1 \in \mathcal{Y}_*^t$ such that for all vectors $\mathbf{y}^2 \in \mathcal{Y}_*$,

$$\mathbf{y}^1 \preceq_{g_{\mathbf{w}}(\mathbf{x}(t)) / \min_j w_j} \mathbf{y}^2.$$

Then, it follows from Definition 6 that $I_{\epsilon+}^1(\mathcal{Y}_*^t) \leq \max_j \frac{1}{w_j} g_{\mathbf{w}}(\mathbf{x}(t))$. See Figure 1 for pictorial proof of the above with a bi-objective problem. $\blacksquare$

While the former theorem justifies the smoothness assumption, the latter provides a link to analyze $\mathcal{Y}_*^t$ with respect to $\mathbf{x}(t)$. The next section—motivated by the two theoretical insights at hand—presents an optimistic algorithm

3 The Weighted Optimistic Optimization Algorithm

Single-objective continuous optimization problems (such as (3)) can be represented as a structured bandit problem where the objective value is a function of some arm parameters. To cope with the infinitely many arms (points in $\mathcal{X}$), arms can be generated in a hierarchical fashion transforming the problem from a many-arm bandit to a hierarchy of multi-armed bandits. As shown in Figure 2, one can use a space-partitioning procedure to iteratively construct finer and finer partitions of the search space $\mathcal{X}$ at multiple depths (scales) $h \in \mathbb{N}_0$. Formally, at depth $h \geq 0$, $\mathcal{X}$ can be partitioned into a set of P^h cells/subspaces $\mathcal{X}_{h,i}$ where $0 \leq i \leq P^h - 1$ such that $\bigcup_{0 \leq i \leq P^h - 1} \mathcal{X}_{h,i} = \mathcal{X}$. These cells are represented by nodes of a P -ary tree $\mathcal{T}$, where a node (h, i) represents the cell $\mathcal{X}_{h,i}$ —the root node $(0, 0)$ represents the entire search space $\mathcal{X}_{0,0} = \mathcal{X}$. A parent node possesses P child nodes $\{(h+1, i_p)\}_{1 \leq p \leq P}$, whose cells form a partition of the parent's cell $\mathcal{X}_{h,i}$. The set of leaves in $\mathcal{T}$ is denoted as $\mathcal{L}$. Likewise, the set of leaves at depth h are denoted by $\mathcal{L}_h$. Each cell is has a representative point (state) $\mathbf{x}_{h,i} \in \mathcal{X}_{h,i}$ at which the function is evaluated (observed). The sampled values $g_{\mathbf{w}}(\mathbf{x}_{h,i})$ are employed optimistically to guide the tree partitioning.

Fig. 1 A visual illustration of Theorem 3 in the particular case of $\min_j w_j < 1$ and $v(t) = 3$ function evaluations. There exists at least one objective vector $\in \mathcal{Y}_*^t$ that lies along the top or right side of $\mathcal{Y}$'s unshaded area (e.g., $\mathbf{y}^1$ and $\mathbf{y}^2$).

We have proved that $g_{\mathbf{w}}$ is Lipschitz-continuous for Lipschitz multi-objective problems. Note that this smoothness is unknown in general for black-box problem. For this, optimistic methods come to rescue. Inspired from [14], the pseudo-code of the proposed Weighted Optimistic Optimization (WOO) is shown in Algorithm 1. WOO grows a tree $\mathcal{T}$ over $\mathcal{X}$ by expanding at most one leaf node per depth in an iterative sweep across $\mathcal{T}$'s depths/levels. At depth $h \geq 0$, a leaf node (h, i) with lowest function value is expanded by partitioning its subspace along one dimension of $\mathcal{X}$. The algorithm inputs/parameters are listed as follows: i) weighted Tchebycheff function; ii) evaluation budget v ; iii) partition factor P ; and iv) partitioning tree depth $h_{max}(t)$. The last two parameters contribute to exploration-vs.-exploitation trade-off.

For better understanding of WOO, a bi-objective optimization problem is selected and Figure 3 shows working of WOO at different stages and approximation of Pareto-front.

Fig. 2 Hierarchical bandits partitioning the decision space $\mathcal{X}$ with a partition factor of $P = 3$ in a form of a P -ary tree, where each node is indexed by its depth h and its order i along the same: (h, i) .

Algorithm 1 The Weighted Optimistic Optimization (WOO) Algorithm

---

Input:
 weighted Tchebycheff function $g_{\mathbf{w}}$,
 search space $\mathcal{X}$,
 partition factor P ,
 evaluation budget v ,
 partitioning tree depth $h_{\max}(t)$.

Initialization:
 $t \leftarrow 1$, $\mathcal{T}_1 = \{(0, 0)\}$, Evaluate $g_{\mathbf{w}}(\mathbf{x}_{0,0})$.

- 1: **while** evaluation budget is not exhausted **do**
- 2: $\nu_{\min} \leftarrow \infty$
- 3:   **for** $l = 0$ **to** $\min\{\text{depth}(\mathcal{T}_t), h_{\max}(t)\}$ **do**
- 4:   Select $(l, o) = \arg \min_{(h,i) \in \mathcal{L}_{t,l}} g_{h,i}^*$
- 5:   **if** $g_{l,o}^* < \nu_{\min}$ **then**
- 6: $\nu_{\min} \leftarrow g_{l,o}^*$
- 7:   Expand (l, o) into its P child nodes
- 8:   Evaluate (l, o) 's P child nodes by $g_{\mathbf{w}}$
- 9:   Add (l, o) 's child nodes to $\mathcal{T}_t$
- 10:   **end if**
- 11: $\mathcal{T}_{t+1} \leftarrow \mathcal{T}_t$
- 12: $t \leftarrow t + 1$
- 13:   **end for**
- 14: **end while**
- 15: **return** $\mathcal{Y}_*^t = \{\mathbf{f}(\mathbf{x}^i) \mid i = 1, \dots, v(t), \mathbf{f}(\mathbf{x}^k) \not\prec \mathbf{f}(\mathbf{x}^i), \forall k = 1, \dots, v, k \neq i\}$

---

4 Convergence Analysis

In general, the performance of multi-armed bandit algorithms (strategies) is assessed through the notion of *regret/loss*: the difference between the present strategy's outcome and that of the optimal strategy. As our problem of interest is multi-objective, we define a Pareto-compliant regret in terms of the additive ϵ -quality indicator in terms of the obtained approximation set (5) as a function of the number of iterations as follows.

$$(7) \quad r(t) = I_{\epsilon+}^1(\mathcal{Y}_*^t) - I_{\epsilon+}^1(\{\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x}))\})$$

In the light of (7), the performance of W00 is analyzed two folds. First, a theoretical finite-time upper bound on the loss (7) is proved. Second, numerical experiments are setup to validate the proven performance on a set of synthetic bi-objective problems.

4.1 Bounding the Indicator Loss

Building on Theorem 2, we employ the smoothness assumptions used in [4, 14]. Subsequently, W00's exploration of the decision space can be quantified in terms of the near-optimality dimension d [14, Definition 1]. From that, the following theorem is deduced.

Theorem 4 *Given t iterations, let us write $h(t)$ the smallest integer h such that*

$$(8) \quad Ch_{\max}(t) \sum_{l=0}^h \delta(l)^{-d} \geq t ,$$

where d is the near-optimality dimension and $\delta(l)$ is a decreasing sequence in $l > 0$ that captures $g_{\mathbf{w}}$ smoothness over the hierarchical partitioning $\{\mathcal{X}_{h,i}\}_{h \geq 0, 0 \leq i < P^h}$. Then the loss is bounded as

$$(9) \quad r(t) \leq \max_j \frac{1}{w_j} \delta(\min(h(t), h_{\max}(t) + 1)) .$$

Proof. Based on [14, Theorem 2], the loss in the quality of $g_{\mathbf{w}}(\mathbf{x}(t))$ is bounded by $\delta(h(t), h_{\max}(t) + 1)$, i.e.,

$$g_{\mathbf{w}}(\mathbf{x}(t)) - \min_{\mathcal{X}} g_{\mathbf{w}}(\mathbf{x}) \leq \delta(h(t), h_{\max}(t) + 1) .$$

Multiplying both sides by $\max_j \frac{1}{w_j}$ yields

$$\begin{aligned} & \underbrace{\max_j \frac{1}{w_j} g_{\mathbf{w}}(\mathbf{x}(t))}_{\geq I_{\epsilon+}^1(\mathcal{Y}_*^t)} - \underbrace{\max_j \frac{1}{w_j} \min_{\mathcal{X}} g_{\mathbf{w}}(\mathbf{x})}_{\geq I_{\epsilon+}^1(\{\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x}))\})} \\ & \leq \max_j \frac{1}{w_j} \delta(h(t), h_{\max}(t) + 1) , \end{aligned}$$

Also, based on Theorem 3, the two terms on the left side are lower bounded by terms under the braces, respectively. Therefore, the loss $r(t)$ of Eq. (7) can be bounded as follows.

$$\begin{aligned} r(t) &= I_{\epsilon+}^1(\mathcal{Y}_*^t) - I_{\epsilon+}^1(\{\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x}))\}) \\ &\leq \max_j \frac{1}{w_j} \delta(h(t), h_{\max}(t) + 1) . \end{aligned}$$

■

4.2 Numerical Validation

In this section, the presented finite-time loss bound in Theorem 4 is validated numerically. To this end, we used the setup described in [2]: a set of synthetic bi-objective ($m = 2$) problems of the form

$$(10) \quad \begin{aligned} f_1(\mathbf{x}) &= \|\mathbf{x} - \mathbf{x}_1^*\|_\infty^{\alpha_1}, \\ f_2(\mathbf{x}) &= \|\mathbf{x} - \mathbf{x}_2^*\|_\infty^{\alpha_2}, \end{aligned}$$

where $\mathcal{X} = [-1, 1]^n$, $n \in \{1, 2\}$, and $(\mathbf{x}_1^*, \mathbf{x}_2^*) \in \{(\mathbf{0}, \mathbf{1}), (\mathbf{0.21}, \mathbf{0.81}), (\mathbf{0.47}, \mathbf{0.61}), (\mathbf{0.57}, \mathbf{0.57})\}$ to capture a range of conflicting objectives. The theoretical bounds (8) and (9) for the problem instances at hand were coded in Python using the `SymPy` package [12] and computed as described in [2].

Moreover, the numerical indicator values $I_{\epsilon+}^1(\mathcal{Y}_*^t)$ (first term of (7)) at each iteration t are computed using a Python implementation of the W00 algorithm using a budget of $v = 10^3$ function evaluations. On the other hand, the second term of (7), $I_{\epsilon+}^1(\{\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x}))\})$, is computed for each of the eight problem instances using a budget of $v = 5 \times 10^5$ function evaluations.

Figure 4 presents the theoretical and numerical results obtained on the eight instances of problem (10). It can be noted that the theoretical measures bound the numerical indicator throughout the algorithm iterations. Furthermore, the bound gets tighter (closer) on problems of less conflicting objectives ($i = 3$ and $i = 4$). That is, the optimal solutions of the objectives are close from each other in the decision space — $(\mathbf{0.47}, \mathbf{0.61}), (\mathbf{0.57}, \mathbf{0.57})$, respectively. This is in line with the observation amount of exploration (the near-optimality dimension's constant) grows linearly with the number of optimal solutions [15]. The code and data of the numerical validation will be made available at the project website.

5 Conclusion

This paper has established the Lipschitz continuity of the weighted Tchebycheff function of Lipschitz-continuous multi-objective problems. The derived, yet unknown smoothness motivated formulating the weighted Tchebycheff problem in a multi-armed bandits setting following the optimism in the face of uncertainty. As a result, we presented and described the Weighted Optimistic Optimization algorithm (W00). W00 looks for the optimal solution by building deterministic hierarchical bandits (in the form of a space-partitioning tree) over the decision space $\mathcal{X}$.

The presented finite-time analysis has established an upper bound on the Pareto-compliant additive ϵ -indicator value of the approximation set obtained from W00's sampled function values as a function of the number of iterations. Numerical experiments on a set of synthetic problems of varying difficulty and dimensionality confirmed the theoretical bounds. To the best of our knowledge, this is the first time, a Pareto-compliant quality indicator is investigated with respect to the weighted Tchebycheff problem in finite time. Potential future research directions include studying the effect of randomizing the weighting vector $\mathbf{w}$ in a stochastic multi-armed bandits setting.

Fig. 3 A visual illustration of W00's sampling and partitioning of the decision space on the *Fonseca and Fleming* benchmark problem [7], where $\mathcal{X} = [-4, 4]^2$, $f_1(\mathbf{x}) = 1 - e^{-\sum_{i=1}^2 (x_i - \frac{1}{\sqrt{2}})^2}$, and $f_2(\mathbf{x}) = 1 - e^{-\sum_{i=1}^2 (x_i + \frac{1}{\sqrt{2}})^2}$. The hierarchical partitioning is performed coordinate-wise in a round-robin fashion across iterations.

Fig. 4 Empirical Validation of Theorem 4 using eight instances of problem (10) having the following configurations: (a) $\mathbf{x}_1^* = \mathbf{0}$, $\mathbf{x}_2^* = \mathbf{1}$, and $\mathcal{X} = [-1, 1]^1$; (b) $\mathbf{x}_1^* = \mathbf{0.21}$, $\mathbf{x}_2^* = \mathbf{0.81}$, and $\mathcal{X} = [-1, 1]^1$; (c) $\mathbf{x}_1^* = \mathbf{0.47}$, $\mathbf{x}_2^* = \mathbf{0.61}$, and $\mathcal{X} = [-1, 1]^1$; (d) $\mathbf{x}_1^* = \mathbf{0.57}$, $\mathbf{x}_2^* = \mathbf{0.57}$, and $\mathcal{X} = [-1, 1]^1$; (e) $\mathbf{x}_1^* = \mathbf{0}$, $\mathbf{x}_2^* = \mathbf{1}$, and $\mathcal{X} = [-1, 1]^2$; (f) $\mathbf{x}_1^* = \mathbf{0.21}$, $\mathbf{x}_2^* = \mathbf{0.81}$, and $\mathcal{X} = [-1, 1]^2$; (g) $\mathbf{x}_1^* = \mathbf{0.47}$, $\mathbf{x}_2^* = \mathbf{0.61}$, and $\mathcal{X} = [-1, 1]^2$; (h) $\mathbf{x}_1^* = \mathbf{0.57}$, $\mathbf{x}_2^* = \mathbf{0.57}$, and $\mathcal{X} = [-1, 1]^2$. As a function of W00's iterations, the figure shows three curves: the orange, solid curve denotes the additive ϵ -indicator, while the dashed, dotted curve represents the quantity $\bar{I}_{\epsilon+}^1 = \max_j \frac{1}{w_j} (\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x})) + 1) + I_{\epsilon+}^1(\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x})))$. In addition, the top region of each subplot shows the maximum tree depth reached $h_{\max}(t)$ at iteration t . The dashed curve upper bounds the orange, solid curve substantiating Theorem 3's result: $I_{\epsilon+}^1(\mathcal{J}^t) - I_{\epsilon+}^1(\{\mathbf{f}(\arg \min_{\mathbf{x} \in \mathcal{X}} g_{\mathbf{w}}(\mathbf{x}))\}) \leq \max_j \frac{1}{w_j} \delta(h(t), h_{\max}(t) + 1)$.

References

1. Al-Dujaili, A., Suresh, S.: Dividing rectangles attack multi-objective optimization. In: IEEE Congress on Evolutionary Computation (CEC), 2016, pp. 1–8. IEEE, Vancouver, Canada (2016)
2. Al-Dujaili, A., Suresh, S.: Multi-objective simultaneous optimistic optimization. arXiv preprint arXiv:1612.08412 (2016)
3. Al-Dujaili, A., Suresh, S.: A naive multi-scale search algorithm for global optimization problems. *Information Sciences* **372**, 294–312 (2016)
4. Al-Dujaili, A., Suresh, S., Sundararajan, N.: MSO: a framework for bound-constrained black-box global optimization algorithms. *Journal of Global Optimization* pp. 1–35 (2016). DOI 10.1007/s10898-016-0441-5. URL <http://dx.doi.org/10.1007/s10898-016-0441-5>
5. Custódio, A.L., Madeira, J.A., Vaz, A.I.F., Vicente, L.N.: Direct multisearch for multiobjective optimization. *SIAM Journal on Optimization* **21**(3), 1109–1140 (2011)
6. Das, I., Dennis, J.E.: A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. *Structural and multidisciplinary optimization* **14**(1), 63–69 (1997)
7. Fonseca, C.M., Fleming, P.J.: An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary computation* **3**(1), 1–16 (1995)
8. Kim, I.Y., de Weck, O.L.: Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Structural and multidisciplinary optimization* **29**(2), 149–158 (2005)
9. Knowles, J., Thiele, L., Zitzler, E.: A tutorial on the performance assessment of stochastic multi-objective optimizers. TIK-Report 214, Computer Engineering and Networks Laboratory, ETH Zurich, Gloriastrasse 35, ETH-Zentrum, 8092 Zurich, Switzerland (2006)
10. Loshchilov, I.: Surrogate-Assisted Evolutionary Algorithms. Theses, Université Paris Sud - Paris XI ; Institut national de recherche en informatique et en automatique - INRIA (2013). URL <https://tel.archives-ouvertes.fr/tel-00823882>
11. Messac, A., Mattson, C.A.: Normal constraint method with guarantee of even representation of complete pareto frontier. *AIAA journal* **42**(10), 2101–2111 (2004)
12. Meurer, A., Smith, C.P., Paprocki, M., Čertík, O., Kirpichev, S.B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J.K., Singh, S., et al.: Sympy: symbolic computing in python. *PeerJ Computer Science* **3**, e103 (2017)
13. Miettinen, K.: Nonlinear multiobjective optimization, volume 12 of international series in operations research and management science (1999)
14. Munos, R.: Optimistic optimization of a deterministic function without the knowledge of its smoothness. In: *Advances in Neural Information Processing Systems* 24, pp. 783–791. Curran Associates, Inc. (2011)
15. Munos, R.: From bandits to Monte-Carlo Tree Search: The optimistic principle applied to optimization and planning. *Foundations and Trends in Machine Learning* **7**(1), 1–130 (2014). URL <http://hal.archives-ouvertes.fr/hal-00747575>
16. Pareto, V.: *Manual of political economy*. Augustus M. Kelley Publishers, New York (1971)
17. Stadler, W.: A survey of multicriteria optimization or the vector maximum problem, part i: 1776–1960. *Journal of Optimization Theory and Applications* **29**(1), 1–52 (1979)
18. Van Moffaert, K., Nowé, A.: Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research* **15**(1), 3483–3512 (2014)
19. Zadeh, L.: Optimality and non-scalar-valued performance criteria. *IEEE transactions on Automatic Control* **8**(1), 59–60 (1963)
20. Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., Da Fonseca, V.G.: Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary*

Computation **7**(2), 117–132 (2003)