



Characterization of constrained continuous multiobjective optimization problems: A feature space perspective

Aljoša Vodopija ^{a,b,*}, Tea Tušar ^{a,b}, Bogdan Filipič ^{a,b}

^a Department of Intelligent Systems, Jožef Stefan Institute, Ljubljana, Slovenia

^b Jožef Stefan International Postgraduate School, Ljubljana, Slovenia



ARTICLE INFO

Article history:

Received 22 December 2021

Received in revised form 20 April 2022

Accepted 27 May 2022

Available online 1 June 2022

Keywords:

Constrained multiobjective optimization

Problem landscape

Exploratory landscape analysis

Test problem

Benchmarking

ABSTRACT

Despite the increasing interest in constrained multiobjective optimization in recent years, constrained multiobjective optimization problems (CMOPs) are still insufficiently understood and characterized. For this reason, the selection of appropriate CMOPs for benchmarking is difficult and lacks a formal background. We address this issue by extending landscape analysis to constrained multiobjective optimization. By employing four exploratory landscape analysis techniques, we propose 29 landscape features (of which 19 are novel) to characterize CMOPs. These landscape features are then used to compare eight frequently used artificial test suites against a recently proposed suite consisting of real-world problems based on physical models. The experimental results reveal that the artificial test problems fail to adequately represent some realistic characteristics, such as strong negative correlation between the objectives and the overall constraint violation. Moreover, our findings show that all the studied artificial test suites have advantages and limitations, and that no “perfect” suite exists. Additionally, the effectiveness of the proposed features at predicting algorithm performance is demonstrated for two multiobjective optimization algorithms. Benchmark designers can use the obtained results to select or generate appropriate CMOP instances based on the characteristics they want to explore.

© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Many real-world continuous optimization problems involve multiple conflicting objectives that need to be optimized simultaneously and constraints that need to be respected [1]. Such optimization problems are called *constrained multiobjective optimization problems* (CMOPs) and have recently gained much interest in the optimization community. Indeed, many novel approaches for constraint handling and new benchmark suites of CMOPs have been proposed in the last years (e.g., [1–6]).

For this reason, never before has there been such a strong need for problem analysis that would inform the selection of CMOPs for benchmarking multiobjective optimization algorithms (MOA). Currently, the benchmark suites are still insufficiently understood, and their characteristics remain unrevealed. In addition to characteristics describing the number of variables, objectives and constraints, type of objective and constraint functions, and geometric properties of Pareto front shapes, there are only few and limited techniques proposed to explore the search space of CMOPs, especially from the field of exploratory landscape analysis (ELA) [7].

* Corresponding author.

E-mail address: aljosa.vodopija@ijs.si (A. Vodopija).

There are only limited studies on ELA in the multiobjective combinatorial context [8–10], an initial study in continuous multiobjective optimization [11], and some first attempts to visualize bi-objective continuous problem landscapes with up to three variables [12–15]. As for constrained optimization, two works characterize linear and quadratic constraints [16,17]. In addition to features for determining the relationship between the constraints, the total number of constraints and the feasibility ratio in the vicinity of the optimal solution were considered to measure the difficulty of the constrained single-objective optimization problems (CSOPs). In [18], the authors introduce five features to characterize CSOPs with arbitrary constraints. They aimed at describing feasibility ratio, the disjointedness of feasible regions, and relationships between the objectives and constraints. This work was later extended to multiobjective optimization [19].

Although in [19] the authors consider ELA techniques for multiple objectives and constraints simultaneously, many other crucial aspects still need to be considered. For example, recently proposed constraint handling techniques assume that CMOP violation landscapes consist of several “suboptimal” subregions—what we will call *violation multimodality*—or have a complex local structure that both severely aggravate algorithm performance [4]. To the best of our knowledge, none of these aspects has been studied from the ELA perspective yet [20].

Due to the lack of appropriate techniques to characterize CMOPs, preparing a sound and well-designed experimental setup for constrained multiobjective optimization is a challenging task. A poorly designed benchmark might lead to inadequate conclusions about CMOP landscapes and MOA performance. Two steps are required to resolve this issue: 1) providing new ELA techniques specialized for CMOPs and 2) using them to reveal the relationship between CMOP characteristics and MOA performance [21]. In this paper, we address step 1) by extending various ELA techniques to constrained multiobjective optimization and step 2), by building a regression model to predict the MOA performance based on the proposed ELA features.

The contributions of this work are as follows:

- We generalize the concept of a fitness landscape to CMOPs and introduce two mathematical definitions to formally describe violation multimodality in constrained multiobjective optimization.
- We use various ELA approaches including space-filling design, information content, random and adaptive walks to derive 29 ELA features, of which 19 are novel. These features are then used to characterize multimodality and local structure of the violation landscapes and measure the relationship between the objectives and constraints. In contrast to the related work, we also discuss the sensitivity and scalability of the proposed techniques.
- We compare the most frequently used artificial test suites against selected real-world problems from the RCM suite—a novel test suite consisting of real-world CMOPs [22]. Specifically, we assess whether the studied artificial test problems comprehensively represent the characteristics observed in the RCM suite. Note that only continuous and low-dimensional RCM problems are used in this study. The suites are also assessed with respect to the representatives of various CMOP characteristics by contrasting them against the set of all the considered CMOPs.
- We analyze the impact of the proposed ELA features on the performance of two well-known MOAs and discuss the importance of features for building regression models. The effectiveness of the novel features is also demonstrated.

Because of space limitations, we present only selected results in this paper. The interested reader can find the complete results online.¹

The rest of this paper is organized as follows. In Section 2, we provide the theoretical background for landscape analysis of CMOPs, while the theoretical advances are introduced in Section 3. Then, in Section 4, we introduce the ELA features for CMOPs and discuss the methodology used to derive them. Section 5 provides details on the experimental setup, presents the results, evaluates the existing test suites of CMOPs, investigates the relationship between the proposed ELA features and MOA performance, and discusses the sensitivity and scalability of the applied methodology. Finally, Section 6 concludes the work with a summary of findings and ideas for future work.

2. Theoretical background

This section provides the theoretical background for this study. After the definitions of the constrained multiobjective optimization problem and constraint violation, we present the concept of the fitness and violation landscape. Finally, we discuss the topological property of connectedness.

2.1. Constrained multiobjective optimization problems

A CMOP is, without loss of generality, formulated as:

$$\begin{aligned}
 &\text{minimize} && f_m(x), \quad m = 1, \dots, M \\
 &\text{subject to} && g_i(x) \leq 0, \quad i = 1, \dots, I \\
 &&& h_i(x) = 0, \quad i = I + 1, \dots, I + J
 \end{aligned} \tag{1}$$

¹ <https://vodopijaaljosa.github.io/cmop-web/>.

where $x = (x_1, \dots, x_D)$ is a search vector, $f_m : S \rightarrow \mathbb{R}$ are objective functions, $g_i : S \rightarrow \mathbb{R}$ inequality constraint functions, $h_i : S \rightarrow \mathbb{R}$ equality constraint functions, $S \subseteq \mathbb{R}^D$ is a search space of dimension D , and M, I and J are numbers of objectives, inequality and equality constraints, respectively.

In continuous optimization, the equality constraints are usually reformulated into inequality constraints as follows:

$$g_i(x) = |h_i(x)| - \eta \leq 0, \quad i = I + 1, \dots, I + J \tag{2}$$

where $\eta > 0$ is a user-defined tolerance value to relax equality constraints (in our study it is set to 10^{-4}). Using this definition, all constraints can be treated as inequality constraints. A solution x is said to be *feasible* if it satisfies all the constraints $g_i(x) \leq 0$ for $i = 1, \dots, I + J$.

One of the most important concepts in constrained optimization is the notion of the *constraint violation*. For a single constraint g_i it is defined as $v_i(x) = \max(0, g_i(x))$. For all constraints together it is combined as

$$v(x) = \sum_i^{I+J} v_i(x) \tag{3}$$

into the *overall constraint violation*. A solution x is feasible iff its overall constraint violation equals zero ($v(x) = 0$). Note that other definitions for overall constraint violation exist, and their use would impact the analysis performed in this study. However, the proposed definition for the overall constraint violation is by far the most commonly adopted in constrained optimization [20], and as such, it represents the most appropriate choice.

A feasible solution $x \in S$ *dominates* another solution $y \in S$ iff $f_m(x) \leq f_m(y)$ for all $1 \leq m \leq M$ and $f_m(x) < f_m(y)$ for at least one $1 \leq m \leq M$. Additionally, a solution $x^* \in S$ is *Pareto optimal* if there is no solution $x \in S$ such that it dominates x^* .

The set of all feasible solutions is called the *feasible region* and is denoted by $F = \{x \in S | v(x) = 0\}$. All nondominated feasible solutions represent a *Pareto-optimal set*, S_o . The image of the Pareto-optimal set in the objective space is the *Pareto front* and is denoted here by $P_o = \{f(x) | x \in S_o\}$.

2.2. Fitness and violation landscapes

The concept of a *fitness landscape* was formally defined in [23] by identifying the following three key elements:

- $S \subseteq \mathbb{R}^D$ is the search space,
- $f : S \rightarrow \mathbb{R}$ is the objective function,
- $d : S \times S \rightarrow \mathbb{R}$ is the distance metric.

The distance metric d is used to quantify the similarity between solutions in the search space and to define *neighborhoods*. In continuous optimization, the Euclidean distance is often used. Close solutions are those whose Euclidean distance is within a predefined threshold, δ . Formally, the solution $y \in S$ is part of the neighborhood of the solution $x \in S$, denoted by $\mathcal{N}(x, \delta)$, iff $d(x, y) \leq \delta$.

In [18], an adaptation of the fitness landscape was proposed for CSOPs by replacing the objective function with the overall constraint violation function. The resulting idea is known as the *violation landscape*.

2.3. Connectedness

A set $X \in S$ of a topological space is said to be *disconnected* if it can be represented as the union of two or more disjoint non-empty open subsets from X . Otherwise, X is said to be *connected*.

The maximal connected subsets (with respect to the inclusion order) of a non-empty space are called *connected components*. The components of any topological space X form a partition of X —they are disjoint, non-empty, and their union is the entire space.

3. Theoretical advances

In this section, we overview theoretical advances in landscape analysis for CMOPs. After defining constrained multiobjective problem landscapes, we introduce feasible components and basins of attraction in the violation landscape. Finally, we conclude with three examples depicting these novel concepts.

3.1. Constrained multiobjective problem landscapes

Following the definitions of fitness and violation landscapes from Section 2.2, we generalize these concepts to CMOPs by including multiple objective functions and the overall constraint violation function simultaneously. A *constrained multiobjective problem landscape* is defined with the following four elements:

- $S \subseteq \mathbb{R}^D$ is the search space,
- $f : S \rightarrow \mathbb{R}^M$ is the objective vector function,
- $v : S \rightarrow \mathbb{R}$ is the overall constraint violation function,
- $d : S \times S \rightarrow \mathbb{R}$ is the distance metric.

The notion of a neighborhood can be applied to CMOP landscapes without any modification. For example, a solution $x^* \in S$ is a *local Pareto-optimal solution* if it is feasible and there exists a $\delta > 0$ such that no feasible solution from the neighborhood $\mathcal{N}(x^*, \delta)$ dominates x^* . The set of all local Pareto-optimal solutions will be denoted by S_l . It is obvious that $S_0 \subseteq S_l$.

In [4], two topological structures concerning violation multimodality were identified that might cause premature algorithm stagnation: feasible subregions and areas that locally violate the constraints the least. The feasible region can consist of multiple disconnected subregions. The more such subregions there are, the harder it is for the algorithm to find a good approximation for the whole Pareto-optimal set. Moreover, the search space of a CMOP can have many areas that locally violate the constraints the least and are infeasible. Again, many such areas might aggravate the performance of an algorithm since it can get stuck in the infeasible region. Although these concepts are essential, they are not formally defined, which can lead to ambiguous interpretations. To overcome this situation, we provide a rigorous mathematical formulation and, in this way, facilitate their further consideration.

3.2. Feasible components

A connected component $\mathcal{F} \subseteq F$ of the feasible region is called a *feasible component*. Note that all feasible components form a partition of the feasible region. With feasible components, we provide a formal definition of feasible subregions. A feasible component can be treated as “suboptimal” when it contains no Pareto-optimal solutions. Obviously, if a problem landscape contains several “suboptimal” feasible components, it might be harder for an optimizer to find good approximations for all Pareto-optimal solutions.

3.3. Basins of attraction in the violation landscape

Basins of attraction in the fitness landscape are well-known objects widely used to characterize unconstrained single-objective optimization problems. In this paper, we reuse this concept in the violation landscape. However, we need to introduce some crucial notions first.

Similarly to the local Pareto-optimal solution, we can define a local “optimal” solution in the violation landscape—a solution that locally violates the constraints the least. Formally, x^* is a *local minimum-violation solution* if there exist a $\delta > 0$ such that $v(x^*) \leq v(x)$ for all $x \in \mathcal{N}(x^*, \delta)$. Additionally, if there exists no solution $x \in S$ such that $v(x^*) > v(x)$, then x^* is a (*global*) *minimum-violation solution*. The set of all local minimum-violation solutions will be denoted by F_l and of all minimum-violation solutions by F_m . It is obvious that $F_m \subseteq F_l$ and if there exist feasible solutions in S , then $F_m = F$.

A connected component $\mathcal{M} \subseteq F_l$ is called a *local minimum-violation component* and it provides a formal definition for areas that locally violate the constraints the least. All the local minimum-violation components form a partition of F_l .

Now, we can finally define a basin of attraction. Let us consider an abstract local search procedure as a mapping from the search space to the set of local minimum-violation solutions, $\mu : S \rightarrow F_l$, such that $\mu(x) = x$ for all $x \in F_l$. Then, a *basin of attraction* of a local minimum-violation component \mathcal{M} and local search μ is a subset of S in which μ converges towards a solution from \mathcal{M} , i.e., $\mathcal{B}(\mathcal{M}) = \{x \in S \mid \mu(x) \in \mathcal{M}\}$. In this case, the local minimum-violation component \mathcal{M} is said to be an *attractor* of the basin $\mathcal{B}(\mathcal{M})$. If there is only one basin in S , then the corresponding violation landscape is said to be *unimodal*. Otherwise, it is *multimodal*.

Each feasible component is a subset of one of the local minimum-violation components. The reverse is, in general, not true. However, if a local minimum-violation component contains feasible solutions, then it contains at least one feasible component as well. Consequently, one can discover feasible components by first identifying local minimum-violation components with feasible solutions.

3.4. Examples

Fig. 1 illustrates problem landscapes and violation landscapes for problems C2-DTLZ2 [24], MW6 [1], DAS-CMOP1 [2] and MW7 [1] with two variables and two objectives. All figures were obtained by evaluating a grid of 501×501 solutions from the search space. The first column shows the dominance ratio for each grid solution expressed as a proportion of grid solutions that dominate it [12]. Those solutions that are not dominated by any grid solution represent approximations of the Pareto-optimal solutions and are shown in black in the plots. The second column depicts the violation landscape in terms of constraint violation values with feasible components presented in white. The third column shows the problem landscape by visualizing the dominance ratio of the feasible regions in blue hues and the infeasible regions in pink. Again, black denotes approximations of Pareto-optimal solutions.

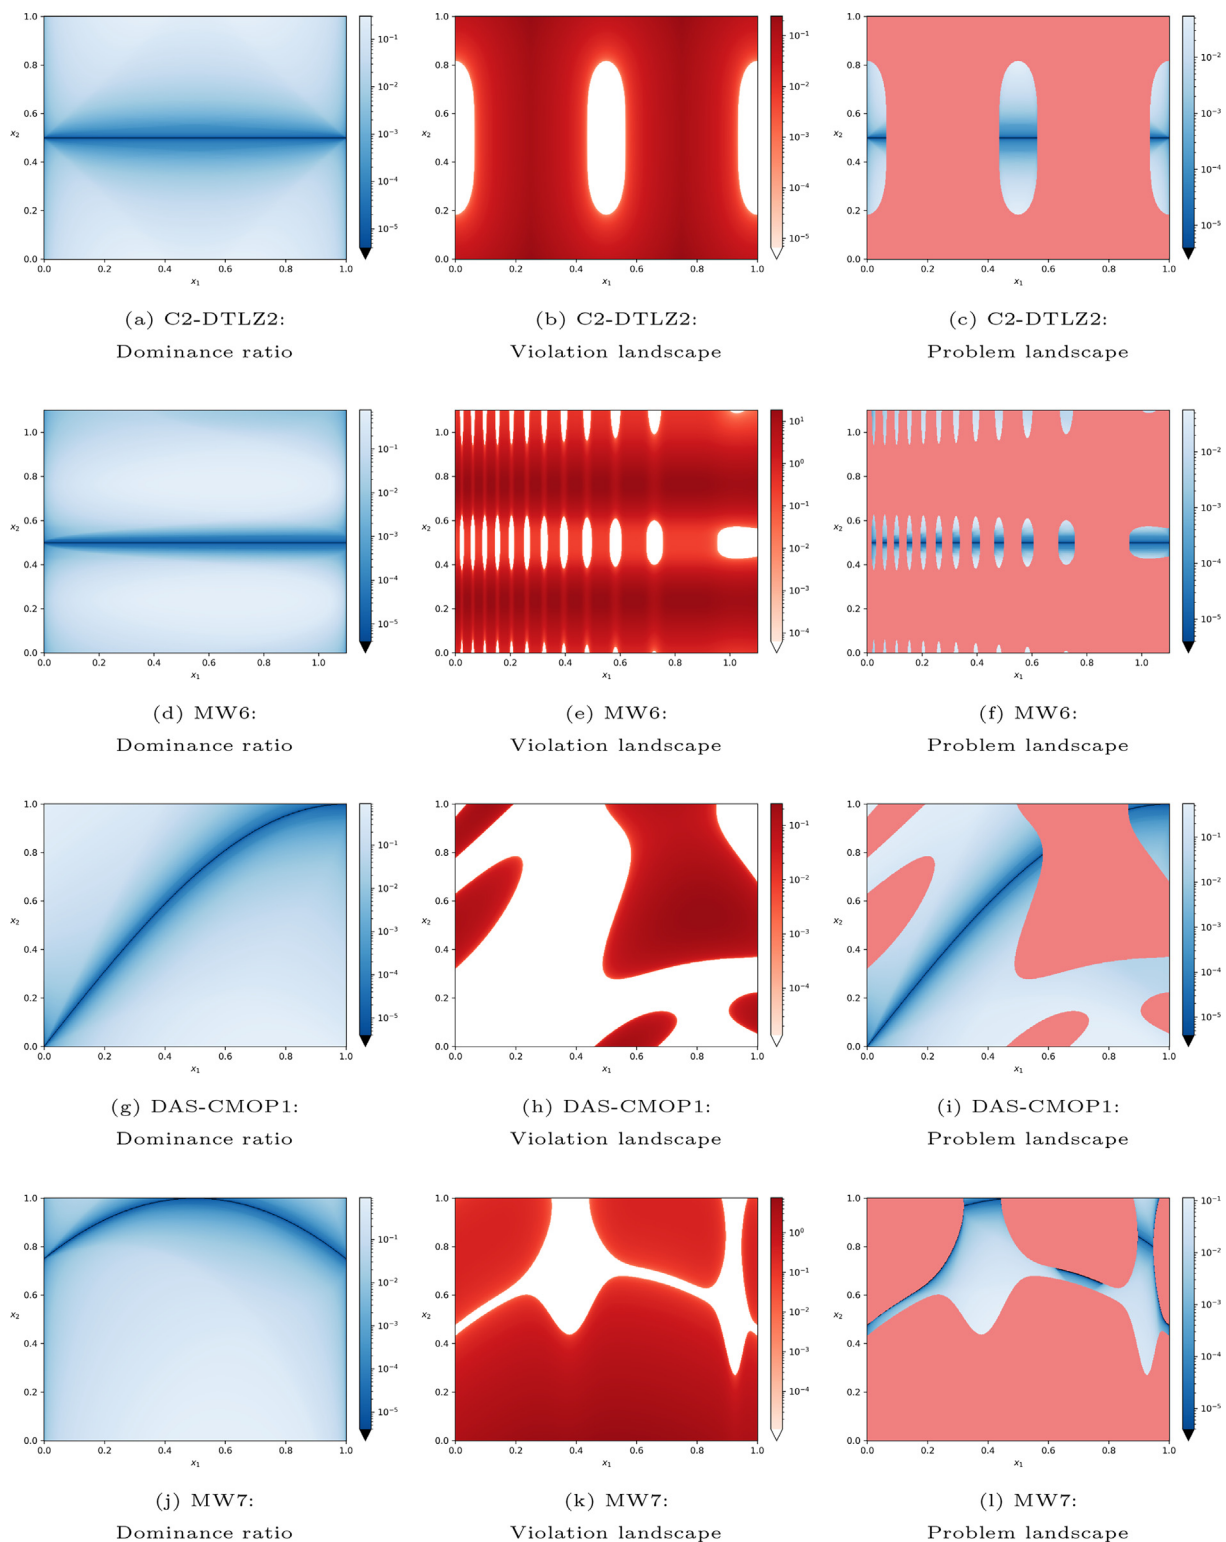


Fig. 1. Plots of the dominance ratio (the first column), the violation landscape (the second column), and the problem landscape (the third column) of three benchmark problems. The approximation of the Pareto-optimal set is depicted in black in dominance ratio and problem landscape plots.

The Pareto-optimal set of C2-DTLZ2 without considering constraints is the horizontal line $L : x_2 = 0.5$ (Fig. 1a). Similarly holds for MW6 (Fig. 1d), while the Pareto-optimal sets of DAS-CMOP1 and MW7 without constraints are curves in the search space (Figs. 1g and 1i).

As we can see from Fig. 1b, C2-DTLZ2 has three feasible components that are also the only three local minimum-violation components. The corresponding basins are separated with two vertical lines $L_1 : x_1 = 0.25$ and $L_2 : x_1 = 0.75$. All the feasible components and basins intersect with the Pareto-optimal set (Fig. 1c). In other words, each basin attracts the search towards nondominated feasible solutions, making this problem easy for constraint handling.

On the other hand, MW6 consists of 35 feasible components and 72 basins (Fig. 1e). Between two adjacent horizontal components, there is an additional basin that contains no feasible solutions. It is also interesting that not all the feasible components contain Pareto-optimal solutions (Fig. 1f). This makes the MW6 problem harder to handle, especially for separation-based constraint handling techniques that separately handle feasible and infeasible solutions and strictly favor feasible solutions over the infeasible ones. The Pareto-optimal set is a subset of the Pareto-optimal set corresponding to MW6 without constraints.

The DAS-CMOP1 problem has three feasible components: a small one in the upper left corner, a medium-sized one in the upper right corner, and a large one in the middle. In contrast to C2-DTLZ2 and MW6, the feasible components have irregular shapes. The feasible components are also the only local minimum-violation components. The Pareto-optimal set is a subset of the Pareto-optimal set corresponding to DAS-CMOP1 without considering constraints.

In contrast, the Pareto-optimal set of the MW7 problem consists of two disconnected parts of the Pareto-optimal set belonging to the MW7 problem without constraints and parts of the boundary between the feasible and infeasible regions (Fig. 1j). This problem has one feasible component and also only one basin of attraction (Fig. 1k).

To identify and count the basins of attraction, we use the truncated Newton method [25] as a local search procedure μ . This method approximates the Newton's direction by approximately solving the Newton equation using an iterative technique. In this study, the conjugate gradient technique is used as the iterative solver. Note that any other deterministic local search could be used instead and possibly result in a different number of basins. Nevertheless, the truncated Newton method proved to be efficient and reliable for identifying basins of attraction.

4. Proposed ELA features

Most of the novel constraint handling techniques proposed in multiobjective optimization are designed to tackle severe violation multimodality, complex local structure of the violation landscapes (disjoint feasible components and rugged constraint functions), and small feasibility ratios [4,26,20,19]. This preposition was the primary motivation for selecting the ELA approaches used in this study. They are all well-known and frequently used methods to derive various features for characterizing problem landscapes. In particular, we include four ELA approaches that we find reliable in deriving features for describing the multimodality and local structure of the violation landscapes and providing some global information such as feasibility ratio, namely space-filling design, information content, random walks, and adaptive walks.

4.1. Space-filling design features

The space-filling design approach is used to derive various features for characterizing feasible components, e.g., the number of components, their sizes and optimality. This information can be in turn used to quantify the violation multimodality. Additionally, features expressing the feasibility ratio, and correlations between the objectives and constraints are also derived within this approach. It comprises the following three steps:

1. Generate an initial sample $X_S \subseteq S$ of solutions following a selected space-filling design and derive feasible solutions.
2. Cluster feasible solutions to obtain approximations for feasible components.
3. Find nondominated solutions among the obtained feasible solutions.

The first feature of this group is the number of feasible components, $N_{\mathcal{F}}$, and it is estimated as the total number of the obtained clusters. Next, the proportion of solutions from X_S that form a particular cluster is used to measure the size of the corresponding feasible component. Three features were derived with this respect: the smallest feasible component, \mathcal{F}_{\min} , the median feasible component, \mathcal{F}_{med} , and the largest feasible component, \mathcal{F}_{\max} .

The following two features considered within this group express the “optimality” of feasible components. The first feature reflects the “optimality” of the largest feasible component, $\mathcal{O}(\mathcal{F}_{\max})$, and it is expressed as the proportion of nondominated feasible solutions in the largest cluster. The second feature is the size of the “optimal” feasible component, \mathcal{F}_{opt} . It is approximated as the proportion of all solutions from the initial sample that form a cluster with the maximum number of nondominated feasible solutions.

The proportion of feasible solutions in the initial sample is used to estimate the feasibility ratio, $\rho_{\mathcal{F}}$, of the corresponding CMOP.

In addition, two features express the correlations between the objectives and constraints. They are the minimum and maximum correlations among objectives and the overall constraint violation and are denoted by ρ_{\min} and ρ_{\max} , respectively.

Finally, the proportion of Pareto-optimal solutions on the feasible region boundary is approximated as the proportion of nondominated solutions on the boundary of one of the obtained clusters and is denoted by $\rho_{\partial S_0}$.

4.2. Information content features

The information content approach was originally used to analyze a sequence of fitness values obtained by sorting a sample of solutions and characterizes the local structure of a fitness landscape. We extend this approach to the violation landscape. Instead of analyzing fitness sequences, we analyze sequences of overall constraint violation values. This way we can derive information about the smoothness, ruggedness and neutrality of the violation landscape.

The approach follows the description provided in [27] and it is here presented only briefly. The interested reader can find a detailed explanation in the original paper. First, a sample of solutions $X_1 = \{x^{(1)}, \dots, x^{(n)}\}$ is generated following a space-filling design. The generated solutions from X_1 are sorted in a sequence as follows: The first solution is selected at random, while each subsequent solution is selected as the nearest solution to the current one. Then, a new sequence, $\phi(\lambda) = \{\phi^{(1)}(\lambda), \dots, \phi^{(n-1)}(\lambda)\}$, is generated following the rule

$$\phi^{(i)}(\lambda) = \begin{cases} \searrow & \text{if } \frac{\Delta v^{(i)}}{\|\Delta x^{(i)}\|} < -\lambda \\ \rightarrow & \text{if } \frac{|\Delta v^{(i)}|}{\|\Delta x^{(i)}\|} \leq \lambda \\ \nearrow & \text{if } \frac{\Delta v^{(i)}}{\|\Delta x^{(i)}\|} > \lambda \end{cases} \tag{4}$$

where $\lambda > 0$, $\Delta x^{(i)} = x^{(i+1)} - x^{(i)}$, $\Delta v^{(i)} = v^{(i+1)} - v^{(i)}$ and $v^{(i)} = v(x^{(i)})$. In more detail, the symbol \searrow indicates a decrease of the overall constraint violation, \rightarrow a neutral area within the threshold λ , and \nearrow an increase of the overall constraint violation. In particular, two consecutive symbols compose a block, ab , where $a, b \in \{\searrow, \rightarrow, \nearrow\}$, which represents an object in the violation landscape, i.e., a slope, peak, etc. For example, $\nearrow \nearrow$ indicates a constant increase of the overall constraint violation, while $\searrow \nearrow$ indicates a bottom. In particular, two identical consecutive symbols represent a smooth area in the violation landscape, while two different consecutive symbols indicate a rugged area.

Finally, the information content of X_1 is defined as the Shannon entropy of the distribution for the blocks ab :

$$H(\lambda) = -\sum_{a \neq b} p_{ab} \log_6 p_{ab} \tag{5}$$

where p_{ab} is the probability of finding one of possible blocks ab , where $a \neq b$, within the sequence $\phi(\lambda)$. The logarithm base is six as this is the number of different blocks where $a \neq b$. The main intuition behind information content is that a large value of $H(\lambda)$ is an indication of perpetual changes in the violation landscape and therefore its ruggedness. Another important measure is the partial information content defined as $M(\lambda) = |\phi(\lambda)|/(n-1)$, where $\phi(\lambda)$ is a sequence derived from $\phi(\lambda)$ by removing all \rightarrow and repeated symbols.

From $H(\lambda)$ and $M(\lambda)$, various features can be derived representing smoothness of the violation landscape. In this work, we study three features often used in the literature [21]: the maximum information content $H_{\max} = \max_i \{H(\lambda)\}$, the settling sensitivity $\varepsilon_S = \log_{10}(\min_{\lambda} \{\lambda | H(\lambda) < 0.05\})$, and the initial partial information $M_0 = M(0)$. Features H_{\max} and M_0 express the violation landscape smoothness (or ruggedness). Smooth landscapes have small values of H_{\max} and M_0 , and vice versa. The feature ε_S , on the other hand, represents the maximum change of the overall constraint violation within the original sequence.

4.3. Random walk features

A random walk through the problem search space is a well-known and often used ELA technique in single-objective optimization. A walk is an ordered sequence of solutions $X_R = \{x^{(1)}, \dots, x^{(n)}\}$ such that $x_i \in \mathcal{N}(x^{(i-1)}, \delta)$ for all $i \in \{2, \dots, n\}$. During a random walk, there is no particular criterion to select the neighboring solution at each step, i.e., a random neighbor is selected [10]. This approach is used to describe the disjointedness of feasible regions, thus providing an insight into the local structure of feasible components.

For constrained single-objective optimization, a ratio of feasible boundary crossings, $\rho_{\partial F}$, was introduced in [18]. This feature quantifies the number of boundary crossings from feasible to infeasible space and vice versa, as encountered by a random walk through the search space. Given a sequence X_R generated from a random walk, a binary sequence is derived $B = \{b^{(1)}, \dots, b^{(n)}\}$ such that $b_i = 0$ iff x_i is a feasible solution. The feature $\rho_{\partial F}$ is defined as the proportion of steps in the random walk that cross the boundary of the feasible region:

$$\rho_{\partial F} = \frac{1}{n-1} \sum_{i=2}^n |b^{(i)} - b^{(i-1)}|. \tag{6}$$

Since the ratio of feasible boundary crossings does not consider the objective values, it is a feature characterizing violation landscapes. Therefore, it can be used in multiobjective optimization without any modifications [19].

To provide a more robust measure, [18] defines $\rho_{\partial F}$ as the average value over N independent random walks. However, this aggregate value provides no local information at all. To avoid this shortcoming, we also consider the maximum and minimum values of $\rho_{\partial F}$ observed in N random walks. The resulting features are: the minimum ratio of feasible boundary crossings $(\rho_{\partial F})_{\min}$, the median ratio of feasible boundary crossings $(\rho_{\partial F})_{\text{med}}$, and the maximum ratio of feasible boundary crossings $(\rho_{\partial F})_{\max}$. A value of $(\rho_{\partial F})_{\min}$ close to zero indicates that there exists at least one larger feasible (or infeasible) region. Similarly, a large value of $(\rho_{\partial F})_{\max}$ indicates that there is an area in the violation landscape with a heavily dissected feasible component.

4.4. Adaptive walk features

In contrast to random walks, an improving neighbor is selected at each step of an adaptive walk, as a local search would do. In unconstrained single-objective optimization, adaptive walks are often used to derive various features to characterize basins of attraction of the fitness landscape. In this paper, we extend these approaches to the violation landscape to quantitatively describe its basins of attraction. Similarly to the space-filling design, this approach is mainly used to quantify the violation multimodality, however, this time by also studying local minimum-violation components without feasible solutions.

The approach considered in this study is similar to the technique described in [28]. However, various novel features specialized for CMOPs are introduced here for the first time. The approach consists of four steps:

1. Generate an initial sample $X_A \subseteq S$ of solutions following a selected space-filling design.
2. Perform a local search for each solution $x \in X_A$ to obtain local minimum-violation solutions.
3. Cluster local minimum-violation solutions to obtain approximations for local minimum-violation components and basins of attraction.
4. Find nondominated feasible solutions among the obtained local minimum-violation solutions.

Using the above approach, we can derive various features to characterize basins of attraction of the violation landscape. The resulting total number of the found clusters provides an approximation for the number of basins, $N_{\mathcal{B}}$, the first feature of this group. The proportion of all solutions from the initial sample that converge to a specific cluster is used to measure the size of the corresponding basin. The resulting features are: the smallest basin, \mathcal{B}_{\min} , the median basin, \mathcal{B}_{med} , and the largest basin, \mathcal{B}_{\max} . Similarly, we derive three equivalent features for feasible basins (basins containing feasible solutions): the smallest feasible basin, $(\mathcal{B}_F)_{\min}$, the median feasible basin, $(\mathcal{B}_F)_{\text{med}}$, and the largest feasible basin, $(\mathcal{B}_F)_{\max}$. In addition, the proportion of all solutions that converge to one of the feasible basins reflects the size of the union of all feasible basins, $\cup \mathcal{B}_F$.

Next, the minimum constraint violation value of all local minimum-violation solutions in a cluster approximates the constraint violation of the corresponding basin defined as $v(B) = \inf\{v(x) | x \in B\}$. Three features are considered: the median constraint violation over all basins, $v(\mathcal{B})_{\text{med}}$, the maximum constraint violation of all basins, $v(\mathcal{B})_{\max}$, and the constraint violation of the largest basin, $v(\mathcal{B}_{\max})$.

Finally, the last two features considered in this group indicate the “optimality” of basins of attraction. The first feature is the “optimality” of the largest basin, $\mathcal{O}(\mathcal{B}_{\max})$, expressed as the number of nondominated feasible solutions in the largest cluster. The second feature is the size of the “optimal” basin, \mathcal{B}_{opt} , expressed as the proportion of solutions that converge to the cluster with the maximum number of nondominated feasible solutions.

5. Experimental analysis

In this section, we first introduce the test suites chosen for the experiments and discuss the experimental setup, including the specific techniques and parameter settings used to derive the ELA features. These features are then used to evaluate the existing test suites. Additionally, the effectiveness of the features in predicting MOA performance is demonstrated on two MOAs. Finally, we present the sensitivity analysis of the parameters used for feature extraction and comment on the scalability of the applied ELA approaches.

5.1. Test suites

The most notable artificial test suites of CMOPs used to assess the performance of constrained multiobjective optimization algorithms are CTP [29], CF [30], C-DTLZ [24], NCTP [31], DC-DTLZ [32], LIR-CMOP [33], DAS-CMOP [2], and MW [1].

All the mentioned test suites consist of artificially designed CMOPs and are not derived from real-world applications. To overcome this weakness, a novel suite named RCM was proposed in [22]. The RCM suite collects 50 real-world optimization problems based on physical models, including problems from mechanical design, chemical engineering, power electronics, etc. The problem instances come in various dimensions and numbers of objectives. After excluding all the problems containing discrete variables and those with more than five variables, 11 RCM problems remained for our experimental analysis and are presented in Table 2. As we can see, ten of the selected RCM problems have two objectives while the remaining one has five objectives. The number of constraints varies from one to eight, and all the constraints are inequalities.

The basic characteristics of the test suites are summarized in Table 3.

Table 1

The proposed ELA features to characterize CMOPs categorized into four groups: space-filling design, information content, random walk, and adaptive walk. “New” indicates that the corresponding feature is proposed in this paper.

Space-filling design features		
$N_{\mathcal{F}}$	Number of feasible components	New
\mathcal{F}_{\min}	Smallest feasible component	New
\mathcal{F}_{med}	Median feasible component	New
\mathcal{F}_{\max}	Largest feasible component	New
$\mathcal{O}(\mathcal{F}_{\max})$	Proportion of Pareto-optimal solutions in \mathcal{F}_{\max}	New
\mathcal{F}_{opt}	Size of the “optimal” feasible component	New
$\rho_{\mathcal{F}}$	Feasibility ratio	[19]
ρ_{\min}	Minimum correlation	[18] ^a
ρ_{\max}	Maximum correlation	[18] ^a
$\rho_{\partial\mathcal{S}_o}$	Proportion of boundary Pareto-optimal solutions	New
Information content features		
H_{\max}	Maximum information content	[27] ^b
ϵ_s	Settling sensitivity	[27] ^b
M_0	Initial partial information	[27] ^b
Random walk features		
$(\rho_{\partial\mathcal{F}})_{\min}$	Minimum ratio of feasible boundary crossings	[18,19]
$(\rho_{\partial\mathcal{F}})_{\text{med}}$	Median ratio of feasible boundary crossings	[18,19]
$(\rho_{\partial\mathcal{F}})_{\max}$	Maximum ratio of feasible boundary crossings	[18,19]
Adaptive walk features		
$N_{\mathcal{B}}$	Number of basins	[28] ^b
\mathcal{B}_{\min}	Smallest basin	New
\mathcal{B}_{med}	Median basin	New
\mathcal{B}_{\max}	Largest basin	New
$(\mathcal{B}_{\mathcal{F}})_{\min}$	Smallest feasible basin	New
$(\mathcal{B}_{\mathcal{F}})_{\text{med}}$	Median feasible basin	New
$(\mathcal{B}_{\mathcal{F}})_{\max}$	Largest feasible basin	New
$\cup\mathcal{B}_{\mathcal{F}}$	Proportion of feasible basins	New
$v(\mathcal{B})_{\text{med}}$	Median constraint violation over all basins	New
$v(\mathcal{B})_{\max}$	Maximum constraint violation of all basins	New
$v(\mathcal{B}_{\max})$	Constraint violation of \mathcal{B}_{\max}	New
$\mathcal{O}(\mathcal{B}_{\max})$	Proportion of Pareto-optimal solutions in \mathcal{B}_{\max}	New
\mathcal{B}_{opt}	Size of the “optimal” basin	New

^aFor the first time applied in multiobjective optimization in our work.

^bFor the first time applied on the violation landscape in our work.

Table 2

Characteristics of the selected RCM problems: dimension of the search space D , number of objectives M , and number of constraints I (all constraints are inequalities).

Test problem	Name	D	M	I
RCM2	Vibrating platform design	5	2	5
RCM3	Two bar truss design	3	2	3
RCM4	Welded beam design	4	2	4
RCM5	Disc brake design	4	2	4
RCM10	Two bar plane truss design	2	2	2
RCM11	Water resources management	3	5	7
RCM12	Simply supported beam design	4	2	1
RCM14	Multiple disk clutch brake design	5	2	8
RCM16	Cantilever beam design	2	2	2
RCM18	Front rail design	3	2	3
RCM20	Hydro-static thrust bearing design	4	2	7

5.2. Experimental setup

The proposed features are demonstrated on the test suites listed in Section 5.1. In particular, bi-objective C-DTLZ and DC-DTLZ problems were considered with the default number of constraints. Additionally, a difficulty triplet of (0.5, 0.5, 0.5) was used for the DAS-CMOP suite as this is by far the most frequently used difficulty triplet in the literature.

Three dimensions of the search space (2, 3, 5) were considered to measure the scalability of various techniques used to derive the ELA features. For each dimension, initial sample sizes were decided based on some initial experiments and were

Table 3

Characteristics of test suites: number of problems, dimension of the search space D , number of objectives M , and number of constraints I (all constraints are inequalities). The characteristics of the selected RCM problems (see Table 2) are shown in parentheses.

Test suite	#problems	D	M	I
CTP [29]	8	*	2	2, 3
CF [30]	10	*	2, 3	1, 2
C-DTLZ [24]	6	*	*	1, *
NCTP [31]	18	*	2	1, 2
DC-DTLZ [32]	6	*	*	1, *
DAS-CMOP [2]	9	*	2, 3	7, 11
LIR-CMOP [33]	14	*	2, 3	2, 3
MW [1]	14	*	2, *	1–4
RCM [22]	50 (11)	2–34 (2–5)	2, 5	1–29 (1–8)

*Scalable parameter.

selected as the minimum number of solutions needed for the feature values to converge. The details are presented in Table 4. Note that because the real-world problems from the RCM test suite are not scalable, the parameter values used for artificial problems with five variables were used also for the four-dimensional RCM problems.

To produce the initial samples used by all the proposed techniques to derive ELA features, we used the space-filling design based on Latin hypercube sampling.

We used the density-based spatial clustering of applications with noise (DBSCAN) [34] as the clustering algorithm for the identification of feasible components and basins, as well as to discover nondominated solutions on the boundary of the feasible region. The following DBSCAN configuration was used in the experiments: distance metric was set to the Euclidean distance, the number of samples in a neighborhood for a point to be considered a core point was set to five. At the same time, the maximum distance between two solutions for one to be considered as in the neighborhood of the other, ϵ , was defined based on initial experiments for each dimension separately (see Table 4).

To express the relationship between the objectives and constraints (ρ_{\min} and ρ_{\max}), the Spearman’s rank correlation coefficient was used.

Following the recommendation from [27], the initial sample size X_1 for the information content features was set to $1000D$.

Simple random walks [35] were employed for calculating features representing ratios of feasible boundary crossings. The number of steps X_R was set to 1000, each with maximum step size, δ , of 1% of the range of the domain for each test problem [18]. Thirty independent runs were conducted to obtain the values for $(\rho_{\partial F})_{\min}$, $(\rho_{\partial F})_{\text{med}}$ and $(\rho_{\partial F})_{\max}$.

Finally, we used the truncated Newton method [25] as a local search procedure used in adaptive walks.

5.3. Implementation details

All CMOPs and techniques to derive the ELA features were implemented in the Python programming language [36]. We used the `pymoo` [37] implementation for CTP, DAS-CMOP and MW, while the rest of the suites was reimplemented from scratch. Next, `pymoo` was used for NSGA-III and C-TAEA, `pyDOE` [38] for Latin hypercube sampling, `scikit-learn` [39] for DBSCAN and Random forest, `pflacco` [40] for the calculation of information content features, and `SciPy` [41] for the truncated Newton method and to calculate Spearman’s rank correlation coefficients. The rest of the functionalities, including the simple random walk, was implemented from scratch.

5.4. Results

Figs. 2 and 3 show violin plots of distributions for seven selected features \mathcal{F}_{opt} , ρ_{\min} , H_{\max} , $N_{\mathcal{B}}$, $\mathcal{V}(\mathcal{F}_{\max})$, $(\rho_{\partial F})_{\text{med}}$ and $\mathcal{V}(\mathcal{B}_{\max})$. The first column represents the distribution for the set of all the considered CMOPs, while the rest of the columns correspond to each suite separately. Each black dot represents one problem instance. Its y -axis depicts the feature value, while the x -axis has no specific meaning and is used for better visualization only. The violin plot (colored area) approximates the probability density function for the feature distribution. For example, Fig. 2b shows there are more problem instances in the CF suite (third column, violin plot in orange) with $\rho_{\min} \approx 0$ than those with $\rho_{\min} \approx -0.75$. In addition, there are no problem instances with $\rho_{\min} \approx -1$. The violin plot in light blue behind each suite corresponds to the distribution of RCM in Fig. 2,

Table 4

Parameters used in the experimental analysis: Dimension of the search space D , initial sample size for space-filling design $|X_S|$ and adaptive walk $|X_A|$, and maximum distance between two solutions ϵ .

D	2	3	4, 5
$ X_S $	25 000	100 000	250 000
$ X_A $	10 000	25 000	50 000
ϵ	0.02	0.04	0.12

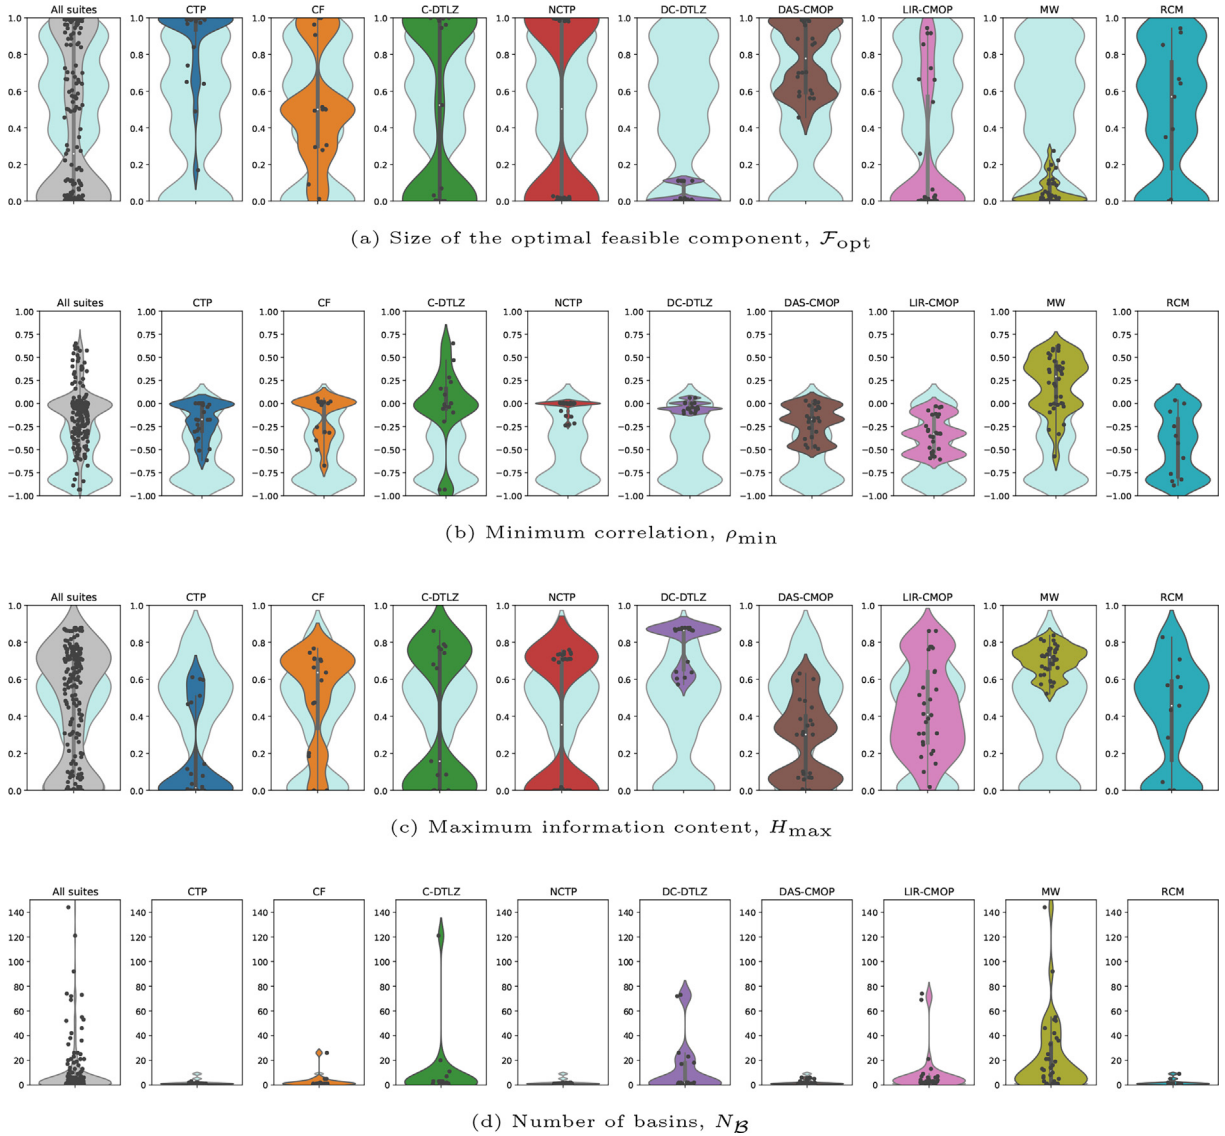


Fig. 2. Violin plots depicting selected feature distributions for the set of all the considered CMOPs (the first column) and for each suite separately (the remaining columns). The violin plot in light blue behind each suite corresponds to the RCM suite.

and in light gray to the distribution of the set of all CMOPs in Fig. 3. The idea is to show the coverage of RCM characteristics by the artificial test suites (Fig. 2) and to expose the differences between a single test suite and the rest of the suites (Fig. 3).

In addition, Fig. 4 shows two heat maps depicting the coverage of RCM characteristics by the artificial test suites, and the coverage of characteristics present in the set of all the considered CMOPs. The coverage is expressed as one minus the mean value of all distances between the feature values from the RCM suite (or the set of all CMOPs) and the nearest feature value from a selected single suite. Note that feature values need to be normalized to [0, 1] before calculating the coverage values. A coverage value near one indicates excellent coverage. This means that for each feature value from RCM (or the set of all CMOPs), there exists at least one problem from the selected single suite whose feature value is close to that one from RCM (or the set of all CMOPs). In Fig. 4, a lighter color indicates that a feature value is well covered by a test suite, and vice versa. For example, in Fig. 4a, $N_{\mathcal{F}}$ (first row) as represented by RCM is well covered by all the artificial suites, while \mathcal{F}_{\max} (fourth row) is poorly covered, especially by DC-DTLZ and MW.

5.5. Test suites evaluation

According to [42], a good benchmark suite should “include the difficulties that are typical of real world instances of the problem class under investigation”. To evaluate whether considered artificial test suites contain the characteristics observed

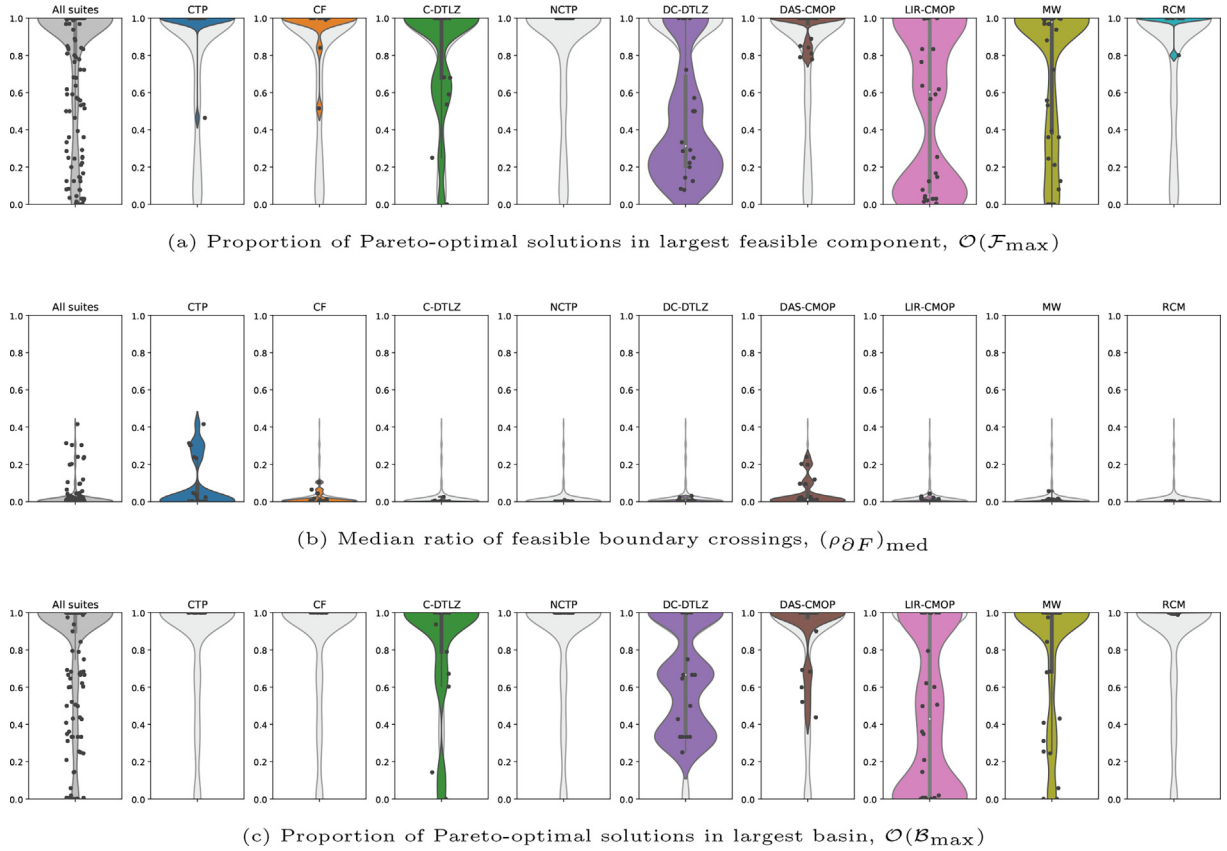


Fig. 3. Violin plots depicting selected feature distributions for the set of all the considered CMOPs (the first column) and for each suite separately (the remaining columns). The violin plot in light gray behind each suite corresponds to the set of all CMOPs.

in real-world CMOPs, we use 11 continuous and low-dimensional problems from the RCM suite and the coverage metric proposed in Section 5.4. Note that the low number of considered real-world problems limits the generality of the conclusions that can be drawn from these results. For example, we cannot claim that some problem characteristics that are missing from these 11 problems are in fact unrealistic. Nevertheless, we can identify which characteristics of these real-world problems are not present in the artificial problems.

The results show that several features characterizing feasible components are not adequately covered by the studied artificial test suites (Fig. 4a, features from $N_{\mathcal{F}}$ to \mathcal{F}_{opt}). Especially the recently proposed artificial test suites NCTP, DC-DTLZ, LIR-CMOP, and MW insufficiently represent problem characteristics of the studied RCM problems. For example, many of the considered RCM problems have only one feasible component, while recent artificial test problems have many feasible components. As a consequence, the RCM problems result in feasible components of any size, while, for example, the MW suite contains problems with rather small feasible components only. This shortcoming of the artificial test suites can be observed nicely for \mathcal{F}_{opt} (Fig. 2a).

Similarly, the existing artificial test suites fail to satisfactorily represent the conflicts between the objectives and the overall constraint violation. As we can see in Fig. 2b, strong negative correlations, $\rho_{min} \in [-1.00, -0.75]$, observed in the studied RCM problems, are found in the C-DTLZ suite only. We can argue that a general multiobjective optimization algorithm would most likely struggle to maintain good convergence and diversity when the objectives and the overall constraint violation are strongly negatively correlated. However, our findings show that the studied artificial test suites do not cover this particular characteristic well.

The results also indicate that the feasibility ratio and information content features are poorly covered by the artificial test suites except for CFs and LIR-CMOPs (Fig. 4a, features ρ_F , H_{max} , ϵ_s and M_0). While the RCM problems take a full range of feature values, the artificial test problems have often extremely small or large feasible ratios and extremely smooth or very rugged violation landscapes (Fig. 2c).

On the other hand, the features related to basins of attraction are mostly covered well by the studied artificial test suites (Fig. 4a, features from $N_{\mathcal{B}}$ to \mathcal{B}_{opt}). The suites that are not diverse enough considering the basin-related features are CF, NCTP, and DAS-CMOP. Additionally, the ratios of feasible boundary crossings are well covered by the considered artificial test suites, while ρ_{os_0} is very well covered by all the artificial test suites except for DAS-CMOP.



Fig. 4. Heat maps showing the coverage of characteristics present (a) in the RCM suite and (b) in the set of all the considered CMOPs.

Finally, the studied artificial test suites have many basin-related characteristics that cannot be observed in the selected RCM problems. While the considered RCM problems often contain only one basin where all Pareto-optimal solutions are located, many recently proposed artificial test problems have several basins, and the Pareto-optimal solutions are spread among multiple basins (Figs. 2d, 3a and 3c). This is in line with the recent development in constrained multiobjective optimization where the primary source of complexity is severe violation multimodality (e.g., Fig. 1f) [20]. However, our study suggests that other relevant aspects, such as the sizes of feasible components and basins of attraction, negatively correlated objectives and the overall constraint violation, diverse ruggedness and feasibility ratios, are frequently insufficiently addressed by the literature.

An additional analysis is devoted to measuring whether the considered test suites contain a wide variety of problems with different characteristics. This is done in the same way as the comparison with RCM, except that in this case the test suites are contrasted against the set of all CMOPs. The result are shown in Fig. 4b.

Interestingly, the results are not drastically different from the previous ones, indicating that RCM contains problems with various characteristics. For example, the coverage of most of the space-filling design and information content features is very similar (Fig. 4b, features from $N_{\mathcal{F}}$ to M_0). This is also true for seven features concerning basins of attraction \mathcal{B}_{med} , \mathcal{B}_{max} , $(\mathcal{B}_{\mathcal{F}})_{\text{min}}$, $(\mathcal{B}_{\mathcal{F}})_{\text{med}}$, $(\mathcal{B}_{\mathcal{F}})_{\text{max}}$, $\cup \mathcal{B}_{\mathcal{F}}$ and \mathcal{B}_{opt} , except for the NCTP suite which has significantly worse coverage when contrasted against the set of all CMOPs.

Within the space-filling design, the only true exception is $\mathcal{O}(\mathcal{F}_{\text{max}})$. The coverage for this feature is significantly worse when comparing the suites against the set of all CMOPs. This happens because RCM does not cover $\mathcal{O}(\mathcal{F}_{\text{max}})$ particularly well (Fig. 3a), which has already been explained before. As we can see, this aspect is even more drastically manifested for $\mathcal{O}(\mathcal{B}_{\text{max}})$ (Fig. 3c).

Finally, the coverage of random walk features is slightly worse when the suites are compared against the set of all CMOPs (Fig. 4b, features $(\rho_{\partial \mathcal{F}})_{\text{min}}$, $(\rho_{\partial \mathcal{F}})_{\text{med}}$ and $(\rho_{\partial \mathcal{F}})_{\text{max}}$). This is true because RCM contains only problems with small values of these particular features. However, in general, this set of features is well covered. Almost all the problems have only small ratios of

feasible boundary crossings except for few problems from CTP, CF and DAS-CMOP (Fig. 3b). The same observation holds for the space-filling design feature $N_{\mathcal{F}}$, and for the following adaptive walk features $N_{\mathcal{B}}$, \mathcal{B}_{\min} , $v(\mathcal{B})_{\text{med}}$, $v(\mathcal{B})_{\text{max}}$ and $v(\mathcal{B}_{\text{max}})$.

In summary, we can see that all the studied test suites have some advantages and some limitations. For this reason, we recommend to include CMOPs from various suites to construct a comprehensive and sound benchmark.

5.6. Feature-based MOA performance prediction

To analyze the relationship between the proposed ELA features and empirical problem hardness, we built a regression model to predict the performance of two well-known MOAs: the Nondominated sorting genetic algorithm III (NSGA-III) [43,24] and a more recently proposed Two-archive evolutionary algorithm for constrained multiobjective optimization (C-TAEA) [32], both equipped with their default constraint handling techniques. The MOA performance was measured as the difference between the hypervolume of the Pareto front and the hypervolume of the Pareto front approximation generated by that algorithm. In more detail, both MOAs were run with a population of 200 solutions for 500 generations. For every 50 generations (10% of the total budget), the difference in hypervolume between the obtained front and the Pareto front was recorded. These differences in hypervolume represented the multi-target variables for the regression analysis. We selected the multi-target approach as we wanted to study the relationship between the proposed features and MOA performance for different budgets. The features from Table 1 and three additional dimensionality features (D , M , and I) represented the predictor variables.

To build regression models, we used the well-known Random forest algorithm [44]. This approach was selected as it is known to generate reliable models without hyperparameter tuning or feature selection [7], and can be used to measure the feature importance. The evaluation of the models was performed using a leave-one-suite-out cross-validation, i.e., the models were repeatedly trained on all the suites except one and then tested on problems from the unseen suite. The results in terms of mean absolute error (MAE), mean square error (MSE), and coefficient of determination (R^2) are reported in Table 5. The scores are averaged over multiple outputs of the regression model and all the suites. Values for MAE and MSE close to 0 indicate good performance of the model. In contrast, R^2 reaches 0 for the baseline model, which always predicts the mean of the differences in hypervolume, and 1 when the model exactly matches the observed values. If the model performs worse than the baseline model, R^2 is less than 0.


The results show the impact of the considered groups of features. As we can see, the most impactful features for both MOAs are those derived from the space-filling design and adaptive walks. Less effective are features related to the local structure of the violation landscape, namely information content, and random walk features. The least valuable features are those associated with the dimensionality of the problems. Indeed, the models trained solely with dimensionality features perform worse than the baseline model. Finally, the best overall results are obtained using all the features.

Additionally, Fig. 5 shows the relative importance based on the mean decrease in impurity of the ELA features at predicting the MOA performance. As we can see, the most important features are similar for both MOAs. In both cases, the most informative feature is \mathcal{B}_{med} showing that basin sizes contain relevant information. Next in the ranking are the features expressing the correlations between the objectives and constraints (ρ_{\min} and ρ_{\max}), and feasibility ratio (ρ_F). This additionally confirms the findings already discussed in [18,19]. Then come feature related to the local structure of the violation landscape, M_0 and $(\rho_{\partial F})_{\text{max}}$ for NSGA-III, and M_0 and H_{max} for C-TAEA.

Table 5

Leave-one-suite-out cross-validation results: Mean absolute error (MAE), mean square error (MSE), and coefficient of determination (R^2). The mean (μ) and standard deviation (σ) calculated over all the suites are reported.

NSGA-III			
Feature group	MAE [$\mu \pm \sigma$]	MSE [$\mu \pm \sigma$]	R^2 [$\mu \pm \sigma$]
Space filling	0.081 ± 0.032	0.019 ± 0.017	0.619 ± 0.075
Info. content	0.103 ± 0.039	0.030 ± 0.025	0.353 ± 0.238
Random walk	0.101 ± 0.038	0.028 ± 0.024	0.457 ± 0.111
Adaptive walk	0.085 ± 0.022	0.018 ± 0.010	0.617 ± 0.026
Dimensionality	0.131 ± 0.028	0.049 ± 0.025	-0.108 ± 0.237
All	0.066 ± 0.020	0.010 ± 0.006	0.787 ± 0.013
C-TAEA			
Feature group	MAE [$\mu \pm \sigma$]	MSE [$\mu \pm \sigma$]	R^2 [$\mu \pm \sigma$]
Space filling	0.090 ± 0.039	0.024 ± 0.022	0.623 ± 0.059
Info. content	0.122 ± 0.061	0.048 ± 0.053	0.294 ± 0.283
Random walk	0.121 ± 0.072	0.048 ± 0.050	0.320 ± 0.199
Adaptive walk	0.105 ± 0.051	0.037 ± 0.035	0.457 ± 0.114
Dimensionality	0.134 ± 0.070	0.060 ± 0.072	-0.091 ± 0.495
All	0.070 ± 0.012	0.011 ± 0.004	0.851 ± 0.002



(a) NSGA-III

(b) C-TAEA

Fig. 5. Feature importance based on mean decrease in impurity for MOA performance prediction: (a) NSGA-III and (b) C-TAEA. Bars correspond to the mean importance value averaged over all budgets, the additional marks denote feature novelty (see Table 1).

Interestingly, the features characterizing feasible components ($N_{\mathcal{F}}$, \mathcal{F}_{\min} , \mathcal{F}_{med} , \mathcal{F}_{\max} , $\mathcal{O}(\mathcal{F}_{\max})$, and \mathcal{F}_{opt}) are important at predicting C-TAEA performance, while they appear irrelevant for predicting NSGA-III performance. A possible explanation is that NSGA-III has difficulties in escaping from local minimum-violation components without feasible solutions, thus the distribution of the basins has a higher impact on the performance of NSGA-III. In contrast, C-TAEA avoids local minimum-violation components without feasible solutions efficiently, therefore, the distribution of feasible components is more informative for its performance.

Finally, the results show that none of the dimensionality features is among the most important ones. The most relevant at predicting NSGA-III performance is D , while for C-TAEA is I .

5.7. Sensitivity analysis

The crucial step in deriving space-filling design and adaptive walk features is the accurate identification of feasible components and basins of attraction, respectively. Once the feasible components and basins are identified, all other features can be unequivocally derived. On the other hand, if the clustering step incorrectly determines the feasible components and basins of attraction, the rest of the features is also affected. The identification of both feasible components and basins of attraction strongly depends on two parameters: initial sample size ($|X_S|$ and $|X_A|$) and the maximum distance between two solutions ε used by the DBSCAN algorithm.

The goal is to find the smallest sample size and the smallest ε that would still yield reliable results. Small sample sizes are preferred since they require fewer function evaluations. Additionally, a small ε is necessary to not erroneously merge two or more disjoint feasible components (or basins of attraction). On the other hand, a too small ε can result in the separation of feasible components (or basins) into two or more components (or basins). For this reason, a careful selection of these parameters is necessary.

Space-filling design features are more sensitive to parameter setting than adaptive walk features. All infeasible solutions in the space-filling design approach are discarded and used only for enumeration. On the other hand, almost all initial solutions for adaptive walk features produce local-minimum violation solutions, thus facilitating the clustering procedure. For this reason, we investigate the sensitivity of space-filling design only. The presented analysis is very similar for adaptive walk features, except that smaller initial sample sizes are required in this case.

The goal of the presented analysis is to measure the parameter sensitivity for the identification of feasible components. More specifically, we study the extraction of the number of feasible components $N_{\mathcal{F}}$ with respect to 1) different parameter values and 2) repetition of the experiments, since a stochastic approach is used to obtain the initial sample. To address 1), we test the space-filling design approach using all combinations of the following parameter values: $|X_S| \in \{10000, 25000, 50000, 100000, 250000, 500000\}$ and $\epsilon \in \{0.01, 0.02, \dots, 0.14, 0.15\}$. For 2), we repeat all the experiments 30 times, each time with a different initial sample. The analysis is performed for $D = 2, 3$, as the correctness of the derived components can be visually checked, while this is not possible for $D = 5$.

The sensitivity analysis is performed on two CMOPs: C2-DTLZ2 and DAS-CMOP1. On the one hand, C2-DTLZ2 represents a simple problem concerning feature extraction. It has three well-shaped feasible components that are spaced apart (Fig. 1b). On the other hand, the feasible components of DAS-CMOP1 have irregular shapes, and the small one in the upper left corner is located close to the large one in the middle (Fig. 1h). This is precisely what makes the identification of components hard for DBSCAN. The 2-D DAS-CMOP1 problem with two variables has three components, while the DAS-CMOP1 with three variables has five components.

The results are shown in the form of heat maps in Fig. 6, where the median logarithmic difference between the exact number of feasible components and the derived number of clusters is depicted. The values are normalized with the number of feasible components and logarithmic to obtain better visualization. Smaller values in lighter colors indicate better performance, and vice versa. Specifically, white color indicates a perfect match.

As we can see from Fig. 6a, for C2-DTLZ2 ($D = 2$) the exact number of feasible components can be obtained for all $|X_S| \geq 25000$ and all $\epsilon \geq 0.02$, while for DAS-CMOP1 ($D = 2$) already the combination $|X_S| = 10000$ and $\epsilon = 0.02$ is sufficient to identify the feasible components (Fig. 6c). Interestingly, for DAS-CMOP1 ($D = 2$) we can see that for $\epsilon \geq 0.12$ one component is recognized incorrectly. What happens is that the small component in the top left corner is merged with the large one due to the oversized ϵ (Fig. 7). Nevertheless, a combination of $|X_S| = 25000$ and $\epsilon = 0.02$ proved to be the most reliable for $D = 2$, resulting in a perfect match for all 30 repetitions on C2-DTLZ2 and DAS-CMOP1. Moreover, Fig. 7 shows that DBSCAN is really effective also when dealing with CMOPs with irregularly-shaped feasible components.

For $D = 3$, Figs. 6b and 6d clearly show we need larger $|X_S|$ and ϵ to obtain the exact number of components. For both CMOPs a combination of $|X_S| = 100000$ and $\epsilon = 0.04$ proved to be the most effective, correctly matching all the components in all 30 repetitions. Additionally, we can see that for DAS-CMOP1 two or more components are incorrectly merged for $\epsilon \geq 0.06$. An interesting observation is that the technique seems to work better for smaller initial sample sizes when using




Fig. 6. Median log differences between the exact number of feasible components and the derived number of clusters for each combination of $|X_S|$ and ϵ for the 2-D and 3-D problems C2-DTLZ2 and DAS-CMOP1.




Fig. 7. Feasible components for DAS-CMOP1 as identified by space-filling design approach for $\varepsilon = 0.02$ (left), $\varepsilon = 0.12$ (right) and $|X_S| = 25000$. Different colors correspond to disjoint feasible components. For $\varepsilon = 0.12$, we can see that the small component in the upper left corner is merged with the large one in the middle.

$\varepsilon \leq 0.02$. The reason behind this behavior is that DBSCAN finds no cluster at all in such cases and the number of identified feasible components is thus zero.

In conclusion, for feasible components that are spaced apart, the space-filling design and similarly adaptive walk approaches are highly reliable and insensitive to small changes in parameter values. Furthermore, the performance is not affected by the shape of the feasible components. On the other hand, for problems with feasible components that are placed close together, the techniques are more sensitive to the selection of parameter values.

A discussion on the parameter sensitivity for information content and random walks can be found in [27,18]. For this reason, we just briefly summarize the key findings. In [27], a comprehensive sensitivity analysis was performed to measure the robustness of the information content approach. Among others, the authors addressed the effect of varying the size of X_1 . The experimental results showed that already a sample size of $100D$ proved to be robust for deriving information content features. Similarly, in [18], the authors demonstrated that the random walk approach was reliable for deriving ratio of feasible boundary crossings. This was shown on a wide variety of violation landscapes by obtaining small standard deviations while repeating the experiments 30 times.

Finally, the space-filling design and information content methods do not require additional function evaluations except those from the initial sample. For this reason, the number of required function evaluations equals $|X_S|$ and $|X_I|$, respectively. Similarly, the random walk approach requires the same number of function evaluations as there are steps in the walk, $|X_R|$. The exact number of function evaluations for space-filling design, information content, and random walks can be found in Section 5.2. In contrast, in the proposed adaptive walk approach, the Newton method requires additional function evaluations for each solution from the initial sample. The total number of function evaluations required by the adaptive walk approach is $N_f|X_A|$ in the median case, where N_f is 7.8, 12.4, and 16.1 for D values of 2, 3, and 5, respectively.

5.8. A note on scalability

For problems with up to five variables the derived ELA feature values closely coincided with actual values. In contrast, it was virtually impossible to investigate high-dimensional CMOPs since the size of the initial sample used by the space-filling design techniques grows exponentially. Indeed, a more detailed experimental analysis showed that to obtain meaningful results, we would need at least 2 000 000 solutions for problems with seven variables and 15 000 000 solutions for problems with ten variables. This is not unexpected since it is a well-known consequence of the curse of dimensionality.

In addition, the feasible components of the artificial test problems often have the shape of an ellipsoid (Figs. 1b and 1e). As the dimension of the search space increases, a feasible component (ellipsoid) becomes an insignificant volume relative to that of the decision space (rectangular area). For this reason, it is very hard to obtain relevant results using space-filling design techniques for such large-scale problems.

It should also be noted that although the proposed ELA features are appropriate to characterize problems with up to five variables, the required number of function evaluations for space-filling design and adaptive walk approaches is significant even for low-dimensional CMOPs. This implies that these features cannot be derived for expensive real-world applications such as simulation-based optimization.

6. Conclusions

In this work, we have presented a comprehensive investigation of CMOPs using landscape analysis. Firstly, we have extended the concepts of fitness and violation landscapes to constrained multiobjective optimization and proposed a math-

emational formulation to describe multimodality in the violation landscape. Secondly, four ELA techniques, namely space-filling design, information content, random walks, and adaptive walks, have been adapted for CMOPs. Next, we have used the resulting features to analyze eight artificial test suites and compare them against the RCM suite, which consists of real-world problems based on physical models. In particular, the artificial test suites were assessed with respect to the representativeness of various RCM characteristics. Additionally, by contrasting the test suites against a set of all the considered CMOPs, we have also assessed whether the test suites contain problems with various characteristics. The effectiveness of the proposed features at predicting algorithm performance was also assessed on two MOAs. Finally, we have discussed the advantages and limitations of the existing artificial test suites as well as the sensitivity and limitations of our methodology. The experimental results show that the proposed ELA techniques are adequate to characterize constrained multiobjective optimization problem landscapes with up to five decision variables.

Our findings indicate that the artificial test problems fail to satisfactorily represent some real-world problem characteristics. In particular, the sizes of feasible components, correlations between the objectives and the overall constraint violation, feasibility ratios, and ruggedness of the violation landscapes are not well represented by the existing artificial test suites. Moreover, the predominant source of complexity in novel artificial test problems is very often violation multimodality. This can lead to biased interpretations and conclusions since many other relevant characteristics are missed. All this clearly suggests that further advances in landscape analysis for CMOPs are required to overcome these drawbacks.

Additionally, the experimental analysis on feature-based MOA performance prediction shows a great potential of the proposed ELA features. The most valuable features are those derived from adaptive walks and space-filling design. However, the results indicate that the synergy gained from all the feature groups yields especially prospective results.

This study is an important step towards enhancing the understanding of test suites developed for constrained multiobjective optimization. It has revealed that all the studied artificial test suites have certain advantages and limitations and showed that no “perfect” test suite exists.

However, we are aware that the presented work has a shortcoming. The proposed approach to characterize violation multimodality is not sufficiently scalable, and thus only low-dimensional problems were addressed in this study. Similarly, these features cannot be used to characterize expensive real-world CMOPs as their derivation requires many function evaluations. Nevertheless, the results show that the proposed ELA features include relevant information for predicting MOA performance, thus their research is relevant and should be further considered. For this reason, in the future, we plan to extend the proposed methodology to large-scale and expensive optimization problems, as well as to further analyze the behavior of constraint handling techniques subject to various problem characteristics.

7. CRediT authorship contribution statement

Aljoša Vodopija: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing – original draft. Tea Tušar: Conceptualization, Methodology, Software, Formal analysis, Writing – review & editing, Visualization. Bogdan Filipič: Conceptualization, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors acknowledge the financial support from the Slovenian Research Agency (young researchers program and research core funding No. P2-0209). The authors also acknowledge the project “Constrained Multiobjective Optimization Based on Problem Landscape Analysis” (No. N2-0254) was financially supported by the Slovenian Research Agency. Finally, this work is part of a project that has received funding from the European Union’s Horizon 2020 research and innovation program under Grant Agreement No. 692286.

References

- [1] Z. Ma, Y. Wang, Evolutionary constrained multiobjective optimization: Test suite construction and performance comparisons, *IEEE Trans. Evol. Comput.* 23 (6) (2019) 972–986, <https://doi.org/10.1109/TEVC.2019.2896967>.
- [2] Z. Fan, W. Li, X. Cai, H. Li, C. Wei, Q. Zhang, K. Deb, E. Goodman, Difficulty adjustable and scalable constrained multiobjective test problem toolkit, *Evol. Comput.* 28 (3) (2019) 339–378, <https://doi.org/10.1162/evco-a-00259>.
- [3] Z.-Z. Liu, Y. Wang, Handling constrained multiobjective optimization problems with constraints in both the decision and objective spaces, *IEEE Trans. Evol. Comput.* 23 (5) (2019) 870–884, <https://doi.org/10.1109/TEVC.2019.2894743>.
- [4] Q. Zhu, Q. Zhang, Q. Lin, A constrained multiobjective evolutionary algorithm with detect-and-escape strategy, *IEEE Trans. Evol. Comput.* 24 (5) (2020) 938–947, <https://doi.org/10.1109/TEVC.2020.2981949>.
- [5] Z. Ma, Y. Wang, Shift-based penalty for evolutionary constrained multiobjective optimization and its application, *IEEE Trans. Cybern.* (2021) 1–13, <https://doi.org/10.1109/TCYB.2021.3069814>.

- [6] Z.-Z. Liu, Y. Wang, B.-C. Wang, Indicator-based constrained multiobjective evolutionary algorithms, *IEEE Trans. Syst. Man Cybern. Syst.* 51 (9) (2021) 5414–5426, <https://doi.org/10.1109/TSMC.2019.2954491>.
- [7] O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, G. Rudolph, Exploratory landscape analysis, in: *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation (GECCO '11)*, ACM, 2011, pp. 829–836, <https://doi.org/10.1145/2001576.2001690>.
- [8] S. Verel, A. Liefvooghe, L. Jourdan, C. Dhaenens, On the structure of multiobjective combinatorial search space: MNK-landscapes with correlated objectives, *Eur. J. Oper. Res.* 227 (2) (2013) 331–342, <https://doi.org/10.1016/j.ejor.2012.12.019>.
- [9] F. Daolio, A. Liefvooghe, S. Verel, H. Aguirre, K. Tanaka, Problem features versus algorithm performance on rugged multiobjective combinatorial fitness landscapes, *Evol. Comput.* 25 (4) (2017) 555–585, https://doi.org/10.1162/evco_a_00193.
- [10] A. Liefvooghe, F. Daolio, S. Verel, B. Derbel, H. Aguirre, K. Tanaka, Landscape-aware performance prediction for evolutionary multiobjective optimization, *IEEE Trans. Evol. Comput.* 24 (6) (2020) 1063–1077, <https://doi.org/10.1109/TEVC.2019.2940828>.
- [11] A. Liefvooghe, S. Verel, B. Lacroix, A.-C. Zăvoianu, J. McCall, Landscape features and automated algorithm selection for multi-objective interpolated continuous optimisation problems, in: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '21)*, 2021, pp. 421–429. doi:10.1145/3449639.3459353.
- [12] C.M. Fonseca, *Multiobjective genetic algorithms with application to control engineering problems (Ph.D. thesis)*, University of Sheffield, 1995.
- [13] P. Kerschke, H. Wang, M. Preuss, C. Grimme, A. Deutz, H. Trautmann, M. Emmerich, Towards analyzing multimodality of continuous multiobjective landscapes, in: *International Conference on Parallel Problem Solving from Nature (PPSN XIV)*, 2016, pp. 962–972, https://doi.org/10.1007/978-3-319-45823-6_90.
- [14] P. Kerschke, C. Grimme, An expedition to multimodal multi-objective optimization landscapes, in: *International Conference on Evolutionary Multi-Criterion Optimization (EMO 2017)*, 2017, pp. 329–343, https://doi.org/10.1007/978-3-319-54157-0_23.
- [15] L. Schäpermeier, C. Grimme, P. Kerschke, To boldly show what no one has seen before: A dashboard for visualizing multi-objective landscapes, in: *International Conference on Evolutionary Multi-Criterion Optimization (EMO 2021)*, 2021, pp. 632–644, https://doi.org/10.1007/978-3-030-72062-9_50.
- [16] S. Poursoltan, F. Neumann, A feature-based analysis on the impact of linear constraints for ϵ -constrained differential evolution, *2014 IEEE Congress on Evolutionary Computation (CEC)* (2014) 3088–3095, <https://doi.org/10.1109/CEC.2014.6900572>.
- [17] S. Poursoltan, F. Neumann, A feature-based comparison of evolutionary computing techniques for constrained continuous optimisation, in: S. Arik, T. Huang, W.K. Lai, Q. Liu (Eds.), *Neural Information Processing*, Springer International Publishing, Cham, 2015, pp. 332–343.
- [18] K.M. Malan, J.F. Oberholzer, A.P. Engelbrecht, Characterising constrained continuous optimisation problems, *2015 IEEE Congress on Evolutionary Computation (CEC)* (2015) 1351–1358, <https://doi.org/10.1109/CEC.2015.7257045>.
- [19] C. Picard, J. Schiffmann, Realistic constrained multiobjective optimization benchmark problems from design, *IEEE Trans. Evol. Comput.* 25 (2) (2021) 234–246, <https://doi.org/10.1109/TEVC.2020.3020046>.
- [20] B. Filipič, A. Vodopija, Multiobjective optimization in the presence of constraints, in: *2021 IEEE Congress on Evolutionary Computation (CEC)*, 2021.
- [21] M.A. Muñoz, Y. Sun, M. Kirley, S.K. Halgamuge, Algorithm selection for black-box continuous optimization problems: A survey on methods and challenges, *Inf. Sci.* 317 (2015) 224–245, <https://doi.org/10.1016/j.ins.2015.05.010>.
- [22] A. Kumar, G. Wu, M.Z. Ali, Q. Luo, R. Mallipeddi, P.N. Suganthan, S. Das, A benchmark-suite of real-world constrained multi-objective optimization problems and some baseline results, *Swarm Evol. Comput.* 67. doi:10.1016/j.swevo.2021.100961.
- [23] P.F. Stadler, *Fitness landscapes*, in: M. Lässig, A. Valleriani (Eds.), *Biological Evolution and Statistical Physics*, Springer, 2002, pp. 183–204.
- [24] H. Jain, K. Deb, An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part II: Handling constraints and extending to an adaptive approach, *IEEE Trans. Evol. Comput.* 18 (4) (2014) 602–622, <https://doi.org/10.1109/TEVC.2013.2281534>.
- [25] S.G. Nash, *User's guide for TN/TNBC: Fortran routines for nonlinear optimization*, Technical Report 397, Mathematical Sciences Department, The Johns Hopkins University, Baltimore, MD, 1984.
- [26] Z. Ma, Y. Wang, W. Song, A new fitness function with two rankings for evolutionary constrained multiobjective optimization, *IEEE Trans. Syst. Man Cybern. Syst.* 51 (8) (2021) 5005–5016, <https://doi.org/10.1109/TSMC.2019.2943973>.
- [27] M.A. Muñoz, M. Kirley, S.K. Halgamuge, Exploratory landscape analysis of continuous space optimization problems using information content, *IEEE Trans. Evol. Comput.* 19 (1) (2015) 74–87, <https://doi.org/10.1109/TEVC.2014.2302006>.
- [28] P. Kerschke, H. Trautmann, Comprehensive feature-based landscape analysis of continuous and constrained optimization problems using the R-package Flacco, in: *Applications in Statistical Computing: From Music Data Analysis to Industrial Quality Improvement*, Springer, 2019, pp. 93–123. doi:10.1007/978-3-030-25147-5_7.
- [29] K. Deb, A. Pratap, T. Meyarivan, Constrained test problems for multi-objective evolutionary optimization, in: *International Conference on Evolutionary Multi-Criterion Optimization (EMO 2001)*, 2001, pp. 284–298.
- [30] Q. Zhang, A. Zhou, S. Zhao, P.N. Suganthan, W. Liu, S. Tiwari, Multiobjective optimization test instances for the CEC 2009 special session and competition, Technical report CES-487, The School of Computer Science and Electronic Engineering, University of Essex, Colchester, UK, 2008.
- [31] J.P. Li, Y. Wang, S. Yang, Z. Cai, A comparative study of constraint-handling techniques in evolutionary constrained multiobjective optimization, *2016 IEEE Congress on Evolutionary Computation (CEC)* (2016) 4175–4182, <https://doi.org/10.1109/CEC.2016.7744320>.
- [32] K. Li, R. Chen, G. Fu, X. Yao, Two-archive evolutionary algorithm for constrained multiobjective optimization, *IEEE Trans. Evol. Comput.* 23 (2) (2019) 303–315, <https://doi.org/10.1109/TEVC.2018.2855411>.
- [33] Z. Fan, W. Li, X. Cai, H. Huang, Y. Fang, Y. You, J. Mo, C. Wei, E. Goodman, An improved epsilon constraint-handling method in MOEA/D for CMOPs with large infeasible regions, *Soft Comput.* 23 (23) (2019) 12491–12510, <https://doi.org/10.1007/s00500-019-03794-x>.
- [34] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in: *International Conference on Knowledge Discovery and Data Mining (KDD-96)*, 1996, pp. 226–231.
- [35] K.M. Malan, A.P. Engelbrecht, A progressive random walk algorithm for sampling continuous fitness landscapes, *2014 IEEE Congress on Evolutionary Computation (CEC)* (2014) 2507–2514, <https://doi.org/10.1109/CEC.2014.6900576>.
- [36] G. Van Rossum, F.L. Drake, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA, 2009.
- [37] J. Blank, K. Deb, *pymoo: Multi-objective optimization in Python*, *IEEE Access* 8 (2020) 89497–89509.
- [38] A. Lee, *pyDOE: Design of experiments for Python*, v. 0.3.8 (2015). <https://pypi.org/project/pyDOE/>.
- [39] F. Pedregosa et al., *Scikit-learn: Machine learning in Python*, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [40] R.P. Prager, *pflacco: A Python interface of the R package flacco*, v. 0.4 (2019). <https://pypi.org/project/pflacco/>.
- [41] P. Virtanen et al., *SciPy 1.0: Fundamental algorithms for scientific computing in Python*, *Nat. Methods* 17 (2020) 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.
- [42] T. Bartz-Beielstein, et al., Benchmarking in optimization: Best practice and open issues, arXiv:2007.03488v2 (2020). arXiv:2007.03488v2.
- [43] K. Deb, H. Jain, An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: Solving problems with box constraints, *IEEE Trans. Evol. Comput.* 18 (4) (2014) 577–601, <https://doi.org/10.1109/TEVC.2013.2281535>.
- [44] L. Breiman, Random forests, *Mach. Learn.* 45 (2001) 5–32, <https://doi.org/10.1023/A:1010933404324>.