

Article

A Problem Landscape Visualisation Method for Multi-Objective Optimisation

Zhiji Cui , Zimin Liang  and Miqing Li * 

School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK;
zxc990@student.bham.ac.uk (Z.C.); zxl525@student.bham.ac.uk (Z.L.)

* Correspondence: m.li.8@bham.ac.uk; Tel.: +44-(0)-121-41-43712

Abstract

Understanding the structure of multi-objective optimisation problems (MOPs) is essential for analysing search difficulty and supporting informed decision-making. In single-objective optimisation, fitness landscapes offer a spatial view of a problem, but extending such visualisations to MOPs is challenging due to the vector-valued nature of objectives. In this work, we introduce Pareto landscape, a fitness landscape visualisation technique for multi-objective optimisation on the basis of the Pareto dominance relation. We illustrate the main characteristics of a Pareto landscape, relate it to the classical fitness landscape, and examine its behaviour across benchmark suites, constrained problems, multimodal problems and real-world cases. We also show how it captures problem landscape structures relevant to optimisation difficulty. A comparison with gradient field heatmaps, PLOT, cost landscape, and constrained cost landscape further demonstrates that Pareto landscape offers complementary insight by highlighting structural patterns not visible with existing visualisation methods. Overall, the results indicate that the Pareto landscape provides a consistent way to observe problem structure across different classes of multi-objective optimisation problems.

Keywords: multi-objective optimisation; visualisation; fitness landscape; Pareto dominance

1. Introduction

Visualisation has long been central to understanding optimisation problems, enabling decision-makers to form more confident and informed judgements [1]. This becomes particularly acute in multi-objective optimisation (MOO), where outcomes are rarely a single best solution. Instead, decision-makers must select from a potentially large number of Pareto-optimal solutions (i.e., the Pareto front), each representing different trade-offs among conflicting objectives. Without appropriate visualisation support, such decisions may rely heavily on subjective interpretation and individual preferences, making it difficult to fully comprehend the implications of each choice [2].

In multi-objective optimisation, visualisation methods can be broadly categorised into two complementary domains, *objective space visualisation* and *problem landscape visualisation*, as surveyed in [3] and illustrated in Figure 1. Objective space methods plot solutions by their objective vectors, and a wide range of methods has been developed, including scatter- and matrix-based plots, heatmaps and parallel-coordinate-based plots, hierarchical layouts, and projection-based methods; see, for example, refs. [4–8] for overviews. These methods are effective at revealing trade-offs and dominance relations, but they provide very limited



Academic Editor: Oliver Schütze

Received: 27 December 2025

Revised: 14 April 2026

Accepted: 23 April 2026

Published: 27 April 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

insight into structural properties of the underlying search space, such as multimodality, attraction basins, or the connectivity of (locally) optimal regions [9,10].




Figure 1. A high-level categorisation of visualisation methods in multi-objective optimisation. Constrained cost landscape refers to the constrained extension of cost landscape proposed in [11].

Problem landscape visualisation methods address this limitation by revealing how solutions are distributed in the search space (also known as the decision space) with respect to their fitness. These methods can be broadly categorised into *abstract* and *native* methods. Abstract methods compress problem landscape structure into network abstractions, such as local optima networks and their multi-objective variants [12–15]. These methods offer valuable insights into connectivity and search dynamics, particularly in combinatorial settings, but they do not preserve the spatial structure of the decision space.

In contrast, native decision space methods display solutions according to their decision variables and assign solution quality using a height or colour scale. Among these, search-based methods such as gradient field heatmaps [16], PLOT [17], and local-dominance landscapes [18] can reveal local structures but depend on the behaviour of the search algorithm considered. Non-search-based methods avoid such artefacts by assigning quality independently of a search algorithm. Examples include the multi-contour plot [19], which overlays contour information of individual objectives to infer structure but may be difficult to interpret when objective interactions are complex; the cost landscape [20], which measures the number of dominating solutions but may not accurately reflect the underlying structure of the optimisation problem; and its constrained extension [11], which additionally represents infeasible regions through constraint violation.

In this paper, we introduce the *Pareto landscape*, a non-search-based landscape visualisation method that assigns solution quality using Pareto non-domination levels. We show that the Pareto landscape preserves dominance-consistent ordering and captures global structural characteristics and certain local structural patterns implied by the dominance-level topology, such as height differences between neighbouring solutions, isolated locally dominant regions, and neutral regions, although it does not explicitly identify locally efficient solutions or basin structures.

A preliminary version of this work was presented at PPSN 2024 [21]. The present journal version substantially extends that study in the following ways:

- It provides an in-depth analysis and discussion of the proposed visualisation, including its position within the field and its features and characteristics compared with existing methods.

- It considers more problem classes, including constrained, multimodal, and real-world problems, and shows how the visualisation reveals their structural characteristics and associated optimisation challenges.

The remainder of this paper is organised as follows. Section 2 provides the necessary background, including key terminology for multi-objective optimisation and landscape visualisation, followed by a review of existing visualisation methods for MOPs. Section 3 introduces Pareto landscape and discusses its fundamental properties and characteristics. Section 4 applies Pareto landscape to well-established benchmark suites (ZDT [22], DTLZ [23], WFG [24], and BBOB [25]) to reveal their structural features. Section 5 explores specific problem types, including constrained problems, multimodal problems, and real-world problems. Section 6 compares the Pareto landscape with existing landscape visualisation methods, specifically gradient field heatmap [16], PLOT [17], cost landscape [20], and its constrained extension [11]. Section 7 discusses the limitations of the proposed method. Finally, Section 8 concludes the paper with a summary of the characteristics revealed for different problem types and discusses future research directions.

2. Preliminaries

In this section, we first give the terminology used in multi-objective optimisation and in our visualisation method, followed by a brief review of representative visualisation methods of MOPs.

2.1. Terminology

Understanding the structural properties of MOPs, such as the distribution of globally and locally efficient regions, is a prerequisite for designing meaningful visualisation methods for the problem landscape. In this subsection, we introduce the key concepts that underpin the landscape representations considered in this work.

A multi-objective optimisation problem (MOP) involves the simultaneous optimisation of multiple conflicting objectives. Without loss of generality, all objectives are assumed to be minimised:

$$\min_{\mathbf{x} \in X} \mathbf{f}(\mathbf{x}) = \min_{\mathbf{x} \in X} (f_1(\mathbf{x}), \dots, f_m(\mathbf{x}))^T, \quad (1)$$

where $X \subseteq \mathbb{R}^n$ denotes the feasible decision space and $\mathbf{x} \in X$ is a decision vector. Each feasible solution maps to an objective vector $\mathbf{y} = \mathbf{f}(\mathbf{x}) \in Y_X$, where $Y_X = \mathbf{f}(X)$ is the feasible objective space.

As objectives typically conflict, solutions cannot generally be totally ordered. Comparisons are instead based on Pareto dominance. Given two objective vectors $\mathbf{y}, \mathbf{y}' \in Y_X$, we write $\mathbf{y} \prec \mathbf{y}'$ if

$$y_i \leq y'_i \forall i \quad \text{and} \quad \exists j : y_j < y'_j. \quad (2)$$

This relation extends to the decision space as $\mathbf{x} \prec \mathbf{x}'$ iff $\mathbf{f}(\mathbf{x}) \prec \mathbf{f}(\mathbf{x}')$.

A feasible solution $\mathbf{x}^* \in X$ is Pareto-optimal (also referred to as *globally efficient*) if no other feasible solution dominates it:

$$\nexists \mathbf{x}' \in X \text{ s.t. } \mathbf{f}(\mathbf{x}') \prec \mathbf{f}(\mathbf{x}^*). \quad (3)$$

All such solutions form the Pareto-optimal set:

$$PS = \{\mathbf{x} \in X : \mathbf{x} \text{ is globally efficient}\}, \quad (4)$$

and their corresponding objective vectors form the Pareto front:

$$PF = \{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in PS\}. \quad (5)$$

Beyond global efficiency, the decision space may contain multiple regions where no local improvement is possible. A solution $x \in X$ is said to be *locally efficient* if there exists a neighbourhood radius $\sigma > 0$ such that no solution within the open ball $B(x, \sigma) = \{x' \in X : \|x' - x\| < \sigma\}$ dominates it:

$$\nexists x' \in B(x, \sigma) \text{ s.t. } f(x') \prec f(x). \quad (6)$$

Given this, the set of all locally efficient solutions is

$$\text{Local PS} = \{x \in X : x \text{ is locally efficient}\}, \quad (7)$$

and the corresponding objective vector form is

$$\text{Local PF} = \{f(x) : x \in \text{local PS}\}. \quad (8)$$

A multi-objective optimisation problem is termed multimodal (MMOP) if it contains at least one locally efficient region in addition to the globally efficient region, or if two or more distinct globally efficient solutions map to the same point on the Pareto front [19,26–29].

In continuous MOO, the multi-objective gradient (MOG) [30] is closely linked to local efficiency. It indicates the steepest simultaneous descent direction of all objectives and satisfies $\nabla f(x) = 0$ if and only if x is locally efficient:

$$\nabla f(x) = \sum_{i=1}^m \alpha_i^* \nabla f_i(x), \quad (9)$$

where the coefficients α^* solve

$$\alpha^* = \arg \min_{\alpha} \left\| \sum_{i=1}^m \alpha_i \nabla f_i(x) \right\| \quad \text{s.t. } \alpha_i \geq 0, \sum_{i=1}^m \alpha_i = 1. \quad (10)$$

The MOG therefore characterises basins of attraction in the decision space.

Another widely used concept in MOO is the non-domination level [31]. Given a finite set $S \subset X$, the first non-dominated set is

$$F_1 = \{x \in S : \nexists x' \in S \text{ such that } x' \prec x\}, \quad (11)$$

and subsequent layers are obtained recursively:

$$F_k = \{x \in S \setminus (F_1 \cup \dots \cup F_{k-1}) : \nexists x' \in S \setminus (F_1 \cup \dots \cup F_{k-1}) \text{ such that } x' \prec x\}. \quad (12)$$

Each solution is assigned the level

$$\text{level}(x) = k \quad \text{if } x \in F_k. \quad (13)$$

In this work, the set S is not an intermediate population of an optimiser but samples drawn from the decision space or all feasible solutions. Under this interpretation, the first non-dominated set F_1 corresponds to the Pareto optimal solutions of the problem (or their sampled representatives), and higher levels reflect increasingly dominated regions of the decision space. These problem-based non-domination levels form the fitness used in Pareto landscape.

2.2. Related Work

Visualisation in multi-objective optimisation has been approached from several perspectives, with methods broadly grouped into *objective space* and *problem landscape* visualisation [3,32], as summarised in Figure 1. In this section, we review the major families of existing methods, with a particular emphasis on landscape visualisation methods that are most closely related to this work.

2.2.1. Objective Space Visualisation

Objective space methods plot solutions according to their objective vectors, and a wide range of methods has been developed [4–8]. Common methods that directly display objective values include scatter plots and scatter-plot matrices [33,34], value paths [35,36], heatmaps [37], and parallel coordinate plots (PCP) [38]. Building on these general methods, several extensions have been proposed to improve interpretability in many-objective settings. For example, heatmaps have similarly been enhanced through dominance-based grouping [39] and ordering [40] to reveal structure more clearly in high-dimensional problems, and PCP has been refined through interpretative guidance and objectives re-ordering [41,42].

Other families include circle- and polygon-based methods such as star diagrams [43,44] and radar charts [45]; icon-based methods such as Chernoff faces [46]; and hierarchical methods such as hierarchical trees [44,47]. Within this broader family of hierarchical structures, tree-based visualisation methods include space-filling layouts such as treemaps [48] and sunbursts [49], which have been applied to many-objective optimisation to visualise dominance hierarchies [50,51]. Projection-based approaches are also widely used, including PCA [52], MDS [53], self-organising maps [54], and RadVis [55]. Several projection-based methods tailored to many-objective optimisation have also been proposed, such as objective-reduction projections [56] and slice-based views, including projections [57,58]. Such projection-based tools are also commonly used to visualise algorithm behaviour by embedding population trajectories throughout an evolutionary run [59–62]. Finally, several visualisation-oriented test problems have been proposed to facilitate the systematic examination of objective space views, multimodality, and structural patterns in many-objective optimisation [63–65].

Despite the abundance of objective space methods, problem landscape visualisation has received comparatively little attention, and for a long time, cost landscape [20] was the only widely used method. Yet understanding structural properties is invaluable for algorithm design and effective decision-making, as these properties strongly influence optimisation difficulty and the interpretation of alternative solutions.

2.2.2. Problem Landscape Visualisation

Problem landscape visualisation methods aim to reveal the structural properties of an optimisation problem with respect to solution quality in the decision space. *Abstract* methods compress landscape information into graph-based representations, where nodes typically represent local optima or locally efficient sets, and edges denote possible transitions. Representative approaches include local optima networks and several multi-objective variants such as PLOSNet [13], CPLOSNet [14], PLON, DNON [66], and nearest-basin networks (NBN) [15]. These methods provide insight into connectivity and search dynamics, particularly in combinatorial settings, but they do not preserve geometric adjacency and therefore do not convey the spatial structure of the decision space.

Native decision space methods directly plot solutions according to their decision variables and encode solution quality through a height or colour scale. Such methods can be broadly divided into *search-based* approaches, such as local dominance landscape [18], gra-

dient field heatmap (GFH) [16], and plot of landscapes with optimal trade-offs (PLOT) [17], which derive solution quality by local search or multi-objective gradients, and *non-search-based* approaches, including the multi-contour plot [19], line cuts [25,67], cost landscape [20], and, for constrained problems, a constrained variant of cost landscape proposed by Tušar et al. [11]. For brevity, we hereafter refer to this constrained variant as the constrained cost landscape.

Among these methods, GFH, PLOT, cost landscape, and constrained cost landscape are most relevant to the present work, as they assign a scalar fitness value to each solution and thereby induce a height- or colour-based representation of the decision space that can be directly compared with the non-domination level used in Pareto landscape. By contrast, methods such as the multi-contour plot [19], which visualises objective contours, and the local dominance landscape [18], which highlights basins of attraction rather than defining a global fitness over all solutions, address different aspects of the landscape structure and are therefore not the primary focus of our comparison. Here we only give a brief description of GFH, PLOT, cost landscape, and constrained cost landscape; their characteristics and potential limitations will be analysed in detail in Section 6.

- **Cost landscape** [20] was proposed in Fonseca's PhD thesis. This method provides a global view of the structure of a multi-objective problem. In a cost landscape, the fitness value (represented through colour or height) of each solution corresponds to the number of other solutions that dominate it, which is commonly referred to as its *Pareto rank* [3,17,20,26]. In some studies and visualisation platforms [11,26,68], a normalised version is also used, referred to as the *dominance rank ratio*.
- **Constrained cost landscape** [11] extends the cost landscape to constrained multi-objective optimisation problems. For feasible solutions, it retains the dominance rank ratio encoding of the cost landscape; for infeasible solutions, it instead encodes the degree of constraint violation. This dual colouring scheme allows the spatial relationship between feasible quality and infeasible proximity to be simultaneously conveyed within a single visualisation.
- **Gradient field heatmap (GFH)** [16] is the first method to use gradient information explicitly for visualising the landscape of multi-objective problems. For each solution, GFH defines a fitness value (shown by colour) as the cumulative length of multi-objective gradients, i.e., the sum of normalised single-objective gradients, along a path towards a locally efficient solution. Starting from an initial solution with an accumulator initialised to zero, the method repeatedly moves to one of its eight neighbouring solutions in the direction indicated by the MOG, adding the normalised MOG length at each step. Once a locally efficient solution is reached, the accumulated length is assigned as the fitness of the starting solution. In this way, GFH provides a gradient-based view of local basin structures and has been recognised as a useful tool for analysing local landscape properties [3].
- **Plot of landscapes with optimal trade-offs (PLOT)** [17] combines the ideas of GFH and cost landscape to provide a joint view of basin structure and efficiency. As in GFH, the fitness of each solution is given by the cumulative MOG length, but PLOT uses a grey-scale colour scheme in which global and locally efficient solutions are emphasised according to their Pareto rank, following the cost landscape. This integration makes it possible to inspect both the organisation of basins of attraction and the global/local optimality structure within a single plot.

3. Pareto Landscape

The classical notion of a fitness landscape, first proposed by Wright [69], describes optimisation as navigating a three-dimensional "mountain range", where the horizontal

axes represent the decision variables and the vertical dimension denotes the fitness of solutions. This visual metaphor has been widely adopted in evolutionary computation and landscape analysis [70,71] due to its ability to intuitively depict key structural aspects of single-objective problems, including peaks and valleys, ridges, plateaus, funnels, and characteristics such as multimodality, ruggedness, and the spatial arrangement of local optima.

Extending this concept to multi-objective optimisation (MOO) is non-trivial, since solution quality is evaluated using multiple conflicting objective functions rather than a single scalar. A natural way to preserve essential trade-off information is to exploit the ordering induced by Pareto dominance. Based on this relation, solutions can be grouped into layers according to their level of non-domination, which can then serve as the vertical coordinate of a landscape representation. In this perspective, the “height” of a point reflects its non-domination level rather than a raw objective value (see the definition of non-domination levels in Section 2.1).

An illustrative example is provided in Figure 2 based on a four-objective multi-point distance minimisation problem (MP-DMP) [72,73]. (Unless otherwise stated, Pareto landscapes shown in this paper are generated using a uniform 500×500 sampling grid.) In this problem, each objective is defined as the Euclidean distance to one vertex of a regular polygon. As a result, the Pareto-optimal set corresponds exactly to the polygon boundary (Figure 2a). The corresponding Pareto landscape (Figure 2b) highlights this set as the highest plateau. Furthermore, as one moves away from the polygon, the dominance ordering results in progressively lower levels, with the boundary of the domain attaining the lowest positions. This demonstrates how Pareto landscapes directly capture the global structure of the problem landscape.




Figure 2. Illustration of a four-objective MP-DMP [72,73] instance. The decision space (a) contains a square region that maps exactly to the Pareto front. The corresponding Pareto landscape (b) places all Pareto-optimal points at the highest level, revealing a symmetric dominance-driven layering structure. From [21].

3.1. Properties and Characteristics

Figure 3 presents a simple example consisting of ten solutions a–j, each characterised by one decision variable and two minimisation objectives. Their objective vectors induce four non-domination levels: {b, g, h} form level 1, {a, e, j} form level 2, {c, i} form level 3, and {d, f} appear at level 4. As shown in Figure 3a, this ranking directly determines the vertical coordinate of each solution in Pareto landscape.




Figure 3. An illustrative example of ten solutions characterised by one decision variable and two minimisation objectives. Heights in the left figure reflect non-domination levels. From [21].

From this illustrative example, several parallels between Pareto landscape and the classical fitness landscape can be identified:

- Pareto-optimal solutions always appear at the highest peaks of the surface, such as **b**, **g** and **h** in Figure 3.
- Local structural patterns, such as isolated dominant regions and spatial separation between efficient regions, are visible in the landscape topology (e.g., **e** forms a locally dominant region, while **b** is spatially separated from **g** and **h**). Note that Pareto landscape does not explicitly identify locally efficient solutions; rather, such structures are inferred from the shape and connectivity of the dominance-level surface.
- Dominance is strictly reflected by height: whenever $x \prec x'$, x is placed at a higher level. For example, **b** dominates **a**, and **g** dominates **e**.
- Solutions with identical objective vectors always lie at the same height (e.g., **c** and **i**).
- The slope between neighbouring points indicates the magnitude of dominance change. For instance, the drop from **b** to **c** is visually steeper than that from **c** to **d**, consistent with a larger difference in non-domination level (2 vs. 1).

There are, however, also distinctive characteristics that differentiate Pareto landscapes from traditional fitness landscapes:

- Higher altitude universally indicates better solution quality, independent of whether the objectives are defined for minimisation or maximisation.
- Mutually non-dominated solutions may fall into the same level or into different levels, depending on the dominance relations introduced by the remaining solutions in the set. For example, **e** and **j** share the same level, as neither dominates the other. Although **j** and **i** are also mutually non-dominated, **i** is lower because it is dominated by **e**.
- Solutions with substantially different objective values may be mapped to the same level if none dominates the other (e.g., **d** and **f**).

These observations stem from the fact that the Pareto landscape is determined by dominance relations rather than by raw objective values. Since dominance induces only a partial order within the evaluated set, mutually non-dominated solutions may occupy the same level even when their objective values differ greatly. Conversely, solutions that are both non-dominated globally may be assigned to different levels when additional solutions dominate one but not the other.

It is worth noting that among the existing visualisation methods reviewed in Section 2.2, cost landscape [20] is most closely related to Pareto landscape. Both methods share several similarities: they assign dominance-based scalar fitness to each solution, place globally optimal solutions at the highest level, and reveal global structural features of the decision space. However, several characteristics described above do not carry over to

the cost landscape. First, the Pareto landscape is based on non-domination levels, and therefore the slope between neighbouring points reflects discrete level differences, whereas in the cost landscape the fitness corresponds to the number of dominating solutions rather than a non-domination level. Second, mutually non-dominated solutions in the cost landscape are unlikely to share the same height, as their dominance counts are sensitive to the composition of the sampled set. Similarly, solutions with substantially different objective values are unlikely to map to the same height, since even small objective differences can alter the number of dominating solutions considerably. These distinctions are analysed further in Section 6.

3.2. On Different MP-DMP Instances

The MP-DMP problem family has been extensively used in the study of evolutionary multi-objective algorithms [74–77], as it allows controlled inclusion of a variety of structural features [73], including multimodality [72] and neutrality [18]. We demonstrate below that such features can be seen in the Pareto landscape.

Multimodality. In a multimodal MP-DMP instance, several congruent polygons are used instead of one, each defining a separate region of globally efficient solutions. Every region maps to the entire Pareto front. As illustrated in Figure 4, these regions appear as distinct plateaus in Pareto landscape, making the multiplicity of attraction basins visually apparent.




Figure 4. Multimodal four-objective MP-DMP with four Pareto-optimal regions (shaded). Each region maps to the full Pareto front and appears as a distinct plateau in Pareto landscape.

Neutrality. Neutral regions are created by assigning identical objective vectors to all solutions within a given subdomain [18]. An example is shown in Figure 5, where a square subregion ($x \in [0.7, 0.9]^2$ with $f_1 = \dots = f_4 = 2$) is neutral. In Pareto landscape, this entire region collapses to a single (worst) level, thereby exposing neutrality directly in the decision space.

Scalability. The number of objectives in MP-DMP can be increased by raising the number of vertices in the underlying polygon. Figure 6 compares three-, four-, and six-objective variants. Although the absolute number of dominance levels changes with dimensionality, the overall geometric form and topology of the landscape remain stable, indicating consistent structural complexity under different objective counts.

The examples above demonstrate that Pareto landscape captures key structural characteristics of MP-DMP problems, including the existence of multiple globally efficient regions, the presence of neutrality, and robustness to changes in the number of objectives.




Figure 5. Illustration of a four-objective MP-DMP with one Pareto-optimal region (shaded) and a neutral region (dashed). In the Pareto landscape, all points in the neutral region share the same worst level. From [21].




Figure 6. Pareto landscapes of MP-DMP instances with three, four, and six objectives. Although the number of non-domination levels differs, the overall landscape topology remains visually consistent. From [21].

4. Pareto Landscape on Representative Test Suites

In this section, we use Pareto landscape to visualise several widely used benchmark suites, to demonstrate that the proposed Pareto landscape can reproduce well-established structural properties. We explicitly indicate for which problems the results are consistent with existing visualisations reported in the literature, and for which they are clearly not (e.g., certain ZDT [22] and DTLZ [23] problems). The underlying reasons why different visualisation methods produce different representations are analysed in detail in Section 6. The benchmark suites considered include the following.

- **ZDT** [22]: a set of two-objective problems with varying search difficulty, including local optima and discontinuity.
- **DTLZ** [23]: designed for scalable MOO, supporting arbitrary numbers of objectives while retaining controlled features such as multimodality and disconnected efficiency regions.
- **WFG** [24]: targeting real-world characteristics such as separability, modality, parameter shift and bias, while enabling fine-grained control over problem difficulty.
- **BBOB (bbob-biobj)** [25]: combining well-studied single-objective functions, resulting in irregular and highly varied landscapes with limited a priori information.

For each BBOB problem, multiple instances exist; here, we consider the first instance. Note that we focus on representative problems from the above problem suites rather than exhaustively visualising every instance. For visualisation, all problems are set in a two-dimensional decision space with two objectives.

4.1. ZDT and DTLZ

Figure 7a,b show the Pareto landscapes of ZDT3 and ZDT4. A characteristic of the ZDT suite is that Pareto-optimal solutions lie in regions where all decision variables except x_1 are zero [22], as can be seen in both panels. ZDT3 exhibits five disconnected Pareto-optimal regions that correspond to separate segments of the Pareto front in objective space, consistent with visualisations reported in [3,78]. This feature can make it somewhat challenging for algorithms to recover all optimal regions. In addition, a locally efficient region is observable at the boundary of the feasible domain, arising from the domain truncation, which is also consistent with existing visualisations. ZDT4 is more challenging: from Figure 7b, many local optima are visible and can trap algorithms if the search step is not appropriate (e.g., too small), a structure also observable via the moPLOT platform [79,80]. It is worth noting, however, that for both ZDT3 and ZDT4, existing visualisations (particularly those based on native decision-space representations) [3,79,80] assign high solution quality to the region where $x_1 = 0$, which differs from the representation produced by Pareto landscape.




Figure 7. Pareto landscapes of ZDT3, ZDT4, DTLZ1 and DTLZ7.

The DTLZ suite [23] generally shares similar features with ZDT. Figure 7c,d show DTLZ1 and DTLZ7. DTLZ1, akin to ZDT4, presents a number of equally spaced local optimal regions (albeit with deeper valleys), as shown in Figure 7c, consistent with visualisations reported in [3,17,26] and also observable via the moPLOT platform [79,80]. DTLZ7 resembles ZDT3 in that it has disconnected Pareto-optimal regions when $x_2 = 0$ (see Figure 7d); however, the gaps between optimal regions are larger, making it more challenging for algorithms to transition from one region to another once an optimal region has been found. The structure of DTLZ7 is similarly observable via the moPLOT platform [79,80]. It is worth noting, similar to the ZDT cases, existing results [3,17,26,79,80] based on native decision-space representations assign high solution quality to the regions where $x_1 = 0$ and $x_1 = 1$ for DTLZ1, and where $x_1 = 0$ for DTLZ7, which differs from the representation produced by the Pareto landscape.

4.2. WFG

Figure 8 presents WFG1, WFG4, WFG7 and WFG8. Compared with the previous benchmark suites, the WFG problems exhibit more deliberately constructed decision-space patterns, reflecting the fine-grained control over separability, modality, and parameter dependencies built into the WFG toolkit [24].




Figure 8. Pareto landscapes of several WFG problems.

In Figure 8a, WFG1 shows a relatively smooth and regular landscape, with a single straight Pareto-optimal ridge. This is consistent with the fact that WFG1 is designed to be separable and unimodal [24]. By contrast, WFG4 in Figure 8b exhibits a clearly rugged multimodal structure, with repeated oscillatory variations across the surface. This agrees with the intended multimodality of WFG4 [24]. Moreover, Huband et al. [24] noted that its hills are larger than those of some other multimodal WFG problems, which makes the problem more likely to trap search algorithms. Figure 8c,d show that WFG7 and WFG8 differ markedly in the shape of their Pareto-optimal regions. WFG7 presents a single straight Pareto-optimal ridge and an otherwise regular landscape structure. WFG8, in contrast, exhibits a curved Pareto-optimal ridge together with a more asymmetric surrounding landscape. This is consistent with the additional parameter dependencies introduced in WFG8, where distance-related parameters depend on position-related parameters, making the problem non-separable [24].

4.3. BBOB

A common feature across ZDT, DTLZ and WFG is the presence of “man-made patterns” (e.g., controllable Pareto-optimal regions), including optimal solutions clustered where certain variables are fixed, or equally spaced local optima. Such regularities can be exploited by search algorithms—for example, by decoupling variables or tuning step sizes. This is not the case in the BBOB (bbob-biobj) suite, where each objective is composed of well-studied single-objective functions; consequently, the Pareto-optimal region is neither known nor directly controllable [3,25].

Figure 9 shows eight bbob-biobj problems. Their landscapes exhibit substantial variety: there are cases of uni-modality (BBOB1, BBOB29), high ruggedness (BBOB10, BBOB47, BBOB55), global structure (BBOB47), large local optima area (BBOB10, BBOB55), and small spike-like local optima (BBOB18, BBOB39, BBOB53). The Pareto-optimal regions are also irregular, including a single straight line (BBOB1), a broken line (BBOB29), several separate lines (BBOB10, BBOB55), disconnected points (BBOB47), and mixtures of lines and points (BBOB18, BBOB39, BBOB53). These structural observations are in line with those reported in [3,17,25,26].




Figure 9. Pareto landscapes of several BBOB (bbob-biobj) problems.

The disconnectedness of Pareto-optimal regions in many BBOB problems poses a significant challenge [26]. To approximate the entire Pareto-optimal region, algorithms must preserve diversity during the search; they cannot simply follow a single direction to one optimum and then explore locally. In contrast, most problems in ZDT, DTLZ and

WFG lack this level of disconnectedness, making it comparatively easier for algorithms to recover the whole Pareto-optimal region.

5. Pareto Landscape on Certain Types of Problems

5.1. Constrained Problems

In multi-objective optimisation, constraints partition the decision space into feasible and infeasible regions. The feasible region may be small in measure, disconnected, or irregularly shaped, and the resulting constrained Pareto set and Pareto front may differ substantially from their unconstrained counterparts. These structural effects make constrained MOPs particularly challenging, since an algorithm must balance convergence towards the Pareto front with effective exploration and the maintenance of feasibility. To study such effects systematically, a number of constrained benchmark suites have been proposed [81].

It is helpful to distinguish constrained problems from the box-constrained setting considered by most existing landscape visualisation methods, where the decision space is restricted only by variable bounds and no additional feasibility constraints are imposed, as in the ZDT and DTLZ problems discussed in the previous section. The native decision-space visualisation methods most closely related to Pareto landscape, namely GFH [16], PLOT [17], and cost landscape [20], were all developed for this box-constrained setting. For constrained problems specifically, Tušar et al. [11] proposed a constrained variant of cost landscape, which extends cost landscape by colouring feasible points according to their dominance rank ratio and infeasible points according to their constraint violation values. Comparisons between Pareto landscape and these existing methods are provided in Section 6.

In this section, we adopt six representative constrained benchmarks—SRN [82], TNK [83], CTP [84], CF [85], C-DTLZ [86], and CMOP [87]—each targeting distinct structural properties:

- **SRN [82]**: an early bi-objective problem with non-linear objectives and non-linear inequality constraints. It highlights how strongly curved feasibility boundaries shape the constrained PF.
- **TNK [83]**: a two-objective problem whose trigonometric and circular constraints produce a feasible region with disconnected components, often used to assess diversity preservation when feasible solutions occur in isolated subsets.
- **CTP [84]**: a family obtained by imposing trigonometric constraints on unconstrained functions to produce controlled feasible regions. CTP1–CTP7 vary the feasible ratio, disconnectedness, and PF shape.
- **CF [85]**: introduced for CEC 2009, combining DTLZ/WFG-type objectives with a wide range of non-linear constraints. CF1–CF10 feature irregular feasible regions and deceptive feasibility boundaries.
- **C-DTLZ [86]**: constrained variants of DTLZ designed to scale to many objectives; constraints can remove large portions of the search space, yielding very low feasible volume and reshaping the PF.
- **CMOP [87]**: a toolkit that constructs CMOP1–CMOP8 with controllable feasible ratios and dimensionality; constraints are categorised by whether they primarily affect PF geometry, feasible ratio, or both.

Because many constrained problems exhibit similar structural patterns in two dimensions, we show only a subset of representative cases. Unless otherwise stated, the same visualisation settings are used throughout the subsequent analyses.

5.1.1. SRN and TNK

Figure 10a,b show SRN and TNK. In SRN, two quadratic constraints leave only a narrow band of Pareto-optimal solutions (shown in green), while the infeasible domain forms a flat plateau. This illustrates how a smooth unconstrained structure can be truncated by feasibility, with only a small subset of solutions remaining valid.

TNK is more challenging: a trigonometric constraint combined with a circular inequality yields several disconnected PS. Pareto-optimal solutions appear in multiple clusters separated by infeasible space, which complicates diversity maintenance because solutions in one feasible component provide no direct guidance toward another.




Figure 10. Pareto landscapes of several constrained test problems: (a) SRN, (b) TNK, (c,d) CTP4 and CTP6, (e,f) CF5 and CF6, (g) C-DTLZ1, and (h,i) CMOP2 and CMOP7. Note that for C-DTLZ1, its infeasible region is highlighted in cyan due to its extremely small size, making it visible in the visualisation.

5.1.2. CTP

Figure 10c,d depict CTP4 and CTP6. CTP4 does not produce a single large infeasible basin; instead, many small infeasible pockets surround the global Pareto region, creating scattered local optima that appear as pits in the landscape. Progress near the true PF is delicate: small steps can be trapped by nearby local PS, whereas larger steps may overshoot the narrow globally optimal area. This calls for careful step-size control to traverse a fragmented feasible neighbourhood without losing global optimality.

CTP6 is qualitatively different: constraints truncate the domain into two distant feasible segments, yielding one global and one local PS far apart. Solutions in one feasible region cannot be used to easily discover the other, risking entrapment in the local PS.

5.1.3. CF

Figure 10e,f show CF5 and CF6. CF5 features two curved corridors in the decision space—a superior left corridor and an inferior right one. The infeasible region blocks portions of the left corridor, making some right corridor solutions Pareto optimal. This splits the PS into two disconnected components separated by infeasible regions, hindering algorithms from maintaining diversity across both parts. Unlike CF5, the infeasible regions in CF6 simply reshape the PS rather than fragmenting it. While the PS remains connected, the steep feasible boundary demands high precision—small perturbations can easily lead to constraint violations.

5.1.4. C-DTLZ

Figure 10g presents C-DTLZ1. The unconstrained DTLZ1 pattern (akin to ZDT4; Section 4.1, Figure 7) exhibits regularly spaced local optima. The added constraints reshape feasibility around the PS, fragmenting it into several disconnected PS. This raises difficulty: once converged to one PS, escaping across infeasible regions to discover others becomes hard, increasing the risk of partial-front convergence.

5.1.5. CMOP

Figure 10h,i show CMOP2 and CMOP7. In CMOP2, periodic constraints slice feasibility into several strips, producing multiple regularly spaced PS; algorithms must traverse infeasible gaps to recover all of them. CMOP7 features large infeasible blocks that separate distant PS and also introduce several local PS, increasing the chance of premature convergence.

5.2. Multimodal Problems

In MOO, *multimodality* refers to the presence of multiple PS that map to the same (or nearly identical) PF, and/or the existence of local PF [19,27,28]. Such situations occur frequently in applications, e.g., [88]. Prior studies [27,29,89] show that maintaining diversity solely in objective space is insufficient: failure to recognise multiple PS restricts the decision maker to a narrow set of options and can degrade decision quality.

We consider six representative multimodal problems suites: Omni-test, SYM, Two-On-One, SS-UF, the MMF suite, and IDMP/IDMP_e, covering symmetric/asymmetric structures and the presence of local PF.

- **Omni-test** [90]: objectives built from sine/cosine yield a highly repetitive PS structure; the true PS/PF is analytically characterisable.
- **SYM-PART** [91]: geometric symmetry (e.g., “simple” and “rotated” variants) generates clusters of equivalent PS, useful to test balanced coverage of symmetric sets.
- **Two-On-One** [92]: bi-objective families with one or two optima; by shifting/lifting minima, one can examine linkage between PS and the PF.
- **SS-UF** [27]: problems with multiple PS mapping to the same PF through symmetric/shifted transformations; functions introduce symmetric/asymmetric multimodality and decision space redundancy (generalised later by MMF).
- **MMF** [89]: a systematic suite distinguishing overlap of PS, PF convexity, and the number/geometry of PS.
- **IDMP/IDMP_e** [28]: introduces local PF to emulate cases where objectives are nearly equivalent but decision variables differ substantially; IDMP_e emphasises hierarchical ranking across global/local PF.

To visualise these multimodal problems, we adopt a modified legend that highlights the relationship between decision-space structure and the Pareto front. Specifically, solutions sharing the same colour correspond to the same point on the Pareto front. As shown in Figure 11, this encoding is illustrated for the SS-UF1 problem by pairing the Pareto

landscape with its corresponding objective-space view. The colour spectrum is consistent across both representations, allowing multimodal solutions that map to the same PF point to be directly identified in the decision space.




Figure 11. Objective space (a) and Pareto landscape (b) of SS-UF1. Colours are consistent across both views: solutions sharing the same colour in the decision space map to the same point on the Pareto front.

For problems with a single global optimum, locally efficient solutions are highlighted in cyan, while the globally efficient solution is highlighted in green. Unlike gradient-based visualisation methods such as GFH or PLOT, Pareto landscape does not explicitly identify local optima through gradient descent. Therefore, to highlight locally efficient regions in the visualisations, we approximate local optima directly from the sampled grid. Specifically, a solution is treated as locally efficient if it is non-dominated within its immediate grid neighbourhood, excluding globally efficient solutions.

Due to finite sampling resolution, we treat solutions as equivalent when their objective vectors are sufficiently close in a relative sense. Specifically, two solutions are considered equivalent if, for every objective k ,

$$|f_i^k - f_j^k| \leq \tau \cdot (f_{\max}^k - f_{\min}^k),$$

where $f_{\max}^k$ and $f_{\min}^k$ are the observed maximum and minimum for objective k , and τ is a relative tolerance. We set $\tau = 10^{-3}$ in all experiments.

5.2.1. Omni-Test and SYM

Figure 12a,b show the Omni-test and SYM. The Omni-test exhibits multiple global optima on a regular grid [90]. Each PS maps to the same PF, yielding many equivalent PS regions; identifying a single PS region is insufficient, and the periodic structure truncated at domain boundaries can introduce additional local PS regions. The SYM problem is related but simpler to configure; it produces multiple PS regions mapping to the same PF with less complex geometry, making it a lightweight test for balanced maintenance of equivalent sets.

5.2.2. Two-On-One

Figure 12c depicts the Two-On-One problem in case 1. The construction combines a bimodal quartic function with a unimodal sphere. In this case, the two quartic optima are symmetric and project to identical objective values, while the optima of the sphere lie midway along the line connecting the two quartic optima. As a result, a single continuous Pareto set is formed, exhibiting symmetric representations in the landscape. This corre-

sponds to a basic form of multimodality, in which identical objective values are shared by multiple decision vectors due to symmetry in the decision space.




Figure 12. Pareto landscapes of several multimodal test problems: (a) Omni-test, (b) SYM-PART1, (c) Two-On-One cases 1, (d,e) SS-UF1 and SS-UF3, (f,g) MMF4 and MMF6, (h) IDMPM2T2, and (i) IDMPM2T2_e.

5.2.3. SSUF

Figure 12d,e show SS-UF1 and SS-UF3. In SS-UF1, the (nearly) continuous PS is symmetric, with mild periodic undulations that create local PS close to the global set. Meanwhile, SS-UF3 exhibits a clearly multimodal structure. The global Pareto-optimal set consists of two disconnected continuous branches, each forming a smooth ridge in the landscape. These branches are separated by regions exhibiting strong oscillations, leading to a highly rugged landscape. This results in multiple attraction basins associated with the different global/local optima.

5.2.4. MMF

Figure 12f,g show MMF4 and MMF6. MMF4 yields two symmetric but fully separated global PS; converging to only one leaves half the global structure undiscovered. MMF6 introduces larger periodic perturbations, breaking the PS into many smaller fragments and adding spurious local PS, making traversal more difficult.

5.2.5. IDMP

Figure 12h,i illustrate IDMPM2T2 and IDMPM2T2_e, respectively. In IDMPM2T2, multiple solutions correspond to the same Pareto-optimal objective vector, giving rise to multimodality with several equivalent global PSs. By contrast, IDMPM2T2_e introduces an additional parameter ε to control the separation between the global and local PFs,

thereby breaking this symmetry. Beyond creating a global–local multimodal structure, the introduction of ε also alters the local landscape characteristics, leading to a flattening of regions around local optima. This indicates that the extended formulation affects not only the PF/PS structure but also the local geometry of the problem, whose impact may become increasingly unpredictable for more complex problem instances.

5.3. Real-World Problems

Beyond synthetic benchmarks, several real-world problem suites [93–97] evaluate algorithm behaviour under practical conditions. These problems integrate domain-inspired structures, irregular feasible regions, and non-linear constraints not fully captured by standard tests.

Because our visualisation approach is limited to two decision variables, only a subset of such problems can be analysed. We therefore focus on representative cases from RWC-MOP [95] and RE [93] (e.g., RWC-MOP10, RWC-MOP16, and RE2-2-4), and additionally include a map-based problem from [98].

- **Map-based problems** [98]: facility placement on geographic maps with distance-based objectives, capturing spatial trade-offs (e.g., accessibility vs. coverage) and producing highly irregular landscape structures.
- **RWC-MOP suite** [95]: difficulty-controllable constrained MOP mimicking real applications; constraints tune feasible ratios and disconnectedness, while objectives reflect typical conflicts (e.g., cost vs. performance).
- **RE problems** [93]: derived from practical engineering formulations (e.g., welded beam design, hatch-cover optimisation, and water resource management) with non-linear objectives/constraints reflecting physical requirements.

We now visualise some of these real-world problems. Visualisation settings remain as before.

5.3.1. Map-Based Problem

Figure 13a shows a map-based facility-placement problem with distance-based objectives. The landscape is highly irregular, exhibiting isolated global solutions together with multiple locally efficient regions shaped by the geographical features.

5.3.2. RWC-MOP

In contrast, RWC-MOP10 and RWC-MOP16 (each defined with only two decision variables) yield comparatively simpler structures. Low-dimensional models are often formulated with reduced non-linear coupling, making their landscapes visually less complex. RWC-MOP10 (Figure 13b) features a narrow ridge-like global PS; RWC-MOP16 (Figure 13c) introduces large infeasible regions while maintaining a continuous global PS.

5.3.3. RE

Figure 13d shows the Pareto landscape for the RE2-2-4 problem. Although the overall landscape is smooth, the PS exhibits a clear turning point. This turning point corresponds to a change in the shape of the PF, consistent with the observations reported in [93].




Figure 13. Pareto landscapes of representative real-world problems: (a) map-based problem, (b,c) two instances from the RWCMOP suite, and (d) RE2-2-4.

6. Comparison with Existing Methods

In this section, we compare Pareto landscape with four representative problem landscape visualisation methods: gradient field heatmap (GFH) [16], PLOT [17], cost landscape [20], and constrained cost landscape of Tušar et al. [11]. As reviewed in Section 2.2, GFH and PLOT are search-based methods that rely on multi-objective gradient (MOG) descent, whereas cost landscape and constrained cost landscape are non-search-based. The first three methods were developed for box-constrained problems, where the decision space is restricted only by lower and upper bounds on the decision variables, whereas the last explicitly addresses constrained problems.

For the visualisations used in this section, GFH and PLOT are reproduced from the moPLOT platform [79,80]. Since these two methods are shown in 2D form, Pareto landscape is also presented in 2D form, with the height dimension omitted, for comparison in Section 6.1. For the remaining comparisons, cost landscape adopts the same colour scale as Pareto landscape. Constrained cost landscape, discussed in Section 6.3, follows the colour settings of the original work [11], uses Pareto rank rather than dominance rank ratio, and is presented in 3D.

Table 1 summarises the main characteristics of all five methods considered in this section. We next discuss each comparison in turn, beginning with the search-based methods.

Table 1. Summary comparison of the methods considered in Section 6. ✓: yes; △: partial; ×: no.

Property	GFH	PLOT	Cost Landscape	Constrained Cost Landscape	Pareto Landscape
Globally efficient set clearly distinguishable	△	✓	✓	✓	✓
Locally efficient set explicitly identified	✓	✓	×	×	×
Attraction basins explicitly shown	✓	✓	×	×	×
Dominance-consistent ordering	×	△	△	△	✓
Search-free	×	×	✓	✓	✓
Applicable to constrained problems	×	×	×	✓	✓
Constraint violation magnitude encoded	×	×	×	✓	×

6.1. Comparison with Search-Based Visualisation Methods

GFH [16] and PLOT [17] start from a regular grid (e.g., 500×500 points) and iteratively follow the steepest MOG descent direction until a point with $\|\text{MOG}\| \leq \delta$ is reached, where the threshold δ is set to 10^{-4} by default. In practice, however, points with very small MOG values may not necessarily be locally efficient, leading to possible misidentification of the Pareto-optimal region.

Figure 14 illustrates this issue for the BBOB11 test problem [25]. The true Pareto-optimal region is a rectangular band in the centre of the decision space. GFH and PLOT fail to mark this region correctly, whereas cost landscape and Pareto landscape accurately highlight it.

Furthermore, search in GFH and PLOT is restricted to the eight grid-adjacent directions. This may deviate from the true steepest descent of the MOG. Additionally, they use the cumulative path length as the fitness proxy, meaning points of identical objective quality may have different fitness values purely due to differing traversal paths.



**Figure 14.** Comparison of GFH, PLOT, cost landscape and Pareto landscape on BBOB11. From [21].

Figure 15 provides such an example on ZDT1 [22]. A diagonal region appears visually inferior in GFH and PLOT, not because those solutions are worse but because the descent paths are longer. Consequently, gradient-based search visualisations may misrepresent dominance relations.

**Figure 15.** Comparison of GFH, PLOT, cost landscape and Pareto landscape on ZDT1. From [21].

It should be noted that GFH and PLOT have a capability that is not provided by the Pareto landscape or cost landscape. That is, locally efficient solutions can be identified

through gradient descent. This allows the basin structure and local optima to be explicitly revealed. Consequently, the choice of method depends on the analysis objective. GFH or PLOT are suitable when the identification of local optima is required, whereas the Pareto landscape provides a dominance-consistent view of the global structure.

Beyond the comparison with search-based methods, differences between cost landscape and Pareto landscape are also visible in Figure 15. Unlike Figure 14, cost landscape and Pareto landscape no longer coincide perfectly. For instance, cost landscape marks the top-left region as worst, while Pareto landscape marks the top region instead. We analyse this distinction in the next subsection.

6.2. Comparison with Cost Landscape

The cost landscape [20] assigns to each point the number of sampled points that dominate it. Like the Pareto landscape, it yields a *relative* fitness over the sampled set; however, it is typically more sensitive to sampling changes. Whenever a new point that dominates a given solution is added, that solution's Pareto rank increases. By contrast, Pareto landscape uses the nondomination *level* (front index) from nondominated sorting; levels change only when front boundaries change, making the representation more stable under insertions that do not alter the front structure.

For example, in Figure 3, inserting additional *non-dominated* points between $\mathbf{g} = (2, 2)$ and $\mathbf{h} = (3, 1)$ does not affect any solution's nondomination level. In the cost landscape, however, such insertions increase the dominance counts of solutions like $\mathbf{d}$, $\mathbf{c}$, $\mathbf{i}$ and $\mathbf{f}$ (which are dominated by the new points), thereby changing their displayed fitness. This illustrates the greater robustness of the Pareto landscape to sample perturbations that preserve the front structure.

It is also worth noting that while the Pareto landscape is more stable under insertions that preserve the front structure, insertions that do alter the front may affect its representation more substantially. In such cases, the dominance-count information provided by the cost landscape may better reflect asymmetries in dominance relations.

In a separate aspect of the difference between the two methods, Figure 16 provides a concrete comparison on ZDT1. With two variables X_1 and X_2 , one can verify the following: (i) fixing X_1 and increasing X_2 always produces a solution dominated by the original; (ii) fixing X_2 and varying X_1 yields solutions that are mutually non-dominated. In the cost landscape (Figure 16a), points A and B share nearly the same height despite A dominating B; conversely, C and D (non-dominated to each other) appear at markedly different heights.




Figure 16. 3D cost landscape (a) and Pareto landscape (b) of ZDT1. Points A–D: A (0,0.2), B (0,0.4), C (0.2,1), D (0.4,1).

A similar inconsistency is observed for DTLZ2. In Figure 17a, the cost landscape assigns almost identical heights to solutions A and B despite A dominating B, while solutions C and D, which are mutually non-dominated, appear at substantially different heights. In contrast, the Pareto landscape (Figure 17b) is strictly consistent with these dominance properties: A lies above B, while C and D are on the same level.




Figure 17. 3D cost landscape (a) and Pareto landscape (b) of DTLZ2. Points A–D: A (0,0.6), B (0,0.8), C (0.2,1), D (0.4,1).

6.3. Comparison with Constrained Cost Landscape

Figure 18 provides a comparison between constrained cost landscape [11] and Pareto landscape on CMOP8 [87]. In the feasible region, constrained cost landscape is based on Pareto rank (or dominance rank ratio), whereas Pareto landscape is based on non-domination level. As a result, the two visualisations exhibit different structural patterns even over the same feasible region. In the Pareto landscape, within each disconnected feasible region, the surface deteriorates in a more nearly layer-by-layer manner as one moves away from the Pareto-optimal region, producing an approximately linear change in height. By contrast, in the constrained cost landscape, the deterioration is initially milder and then becomes steeper.




Figure 18. Comparison between constrained cost landscape (a) and Pareto landscape (b) on CMOP8. In both cases, infeasible regions are rendered as a flat surface at the lowest level. In (a), the feasible region is coloured according to Pareto rank, while the infeasible region is coloured in red according to the degree of constraint violation. In (b), feasible solutions are coloured according to non-domination level, while the infeasible region is uniformly shown in grey.

This is because the number of solutions dominating a point accumulates in most directions away from the Pareto-optimal region and often increases more rapidly as the point moves further into the inferior part of the decision space. This interpretation is also consis-

tent with the observation in Section 3 that even small objective differences may substantially alter the number of dominating solutions, and is also reflected in Figures 16 and 17, although special directions such as $x_1 = 0$ in ZDT1 and $x_1 = 0$ or $x_1 = 1$ in DTLZ2 form exceptions, as along these directions one objective does not deteriorate further.

For the infeasible region in Figure 18, the two visualisations share the same geometric shape. However, the constrained cost landscape provides additional information by colouring this surface according to the degree of constraint violation. More specifically, over the whole decision space, the violation becomes more severe as (x_1, x_2) approaches $(0, 0)$ or $(1, 1)$. In addition, for each infeasible band between two parallel feasible regions, the violation is highest near the middle of the band. In this regard, constrained cost landscape may offer useful insight into algorithm behaviour near feasibility boundaries [11] and may help indicate which parts of the infeasible region are closer to feasibility and which may constitute stronger constraint barriers.

7. Limitations

Despite the insights provided by the Pareto landscape, several limitations of the proposed method should be acknowledged.

First, Pareto landscape is a native decision-space visualisation that relies on the geometric structure of a continuous Euclidean decision space. Consequently, its applicability is primarily restricted to continuous optimisation problems with a small number of decision variables. For discrete or combinatorial problems, where neighbourhood relations are defined through combinatorial operators rather than geometric adjacency, the spatial interpretation underlying the Pareto landscape becomes less meaningful, as the decision space lacks a natural geometric embedding. In such settings, abstract landscape representations, such as local optima networks and their multi-objective extensions [12–15], may provide a more appropriate way to analyse structural properties. A limited exception exists for low-dimensional problems with ordered discrete variables, such as integer-valued decision spaces, where regular grid sampling remains feasible and the spatial interpretation of the Pareto landscape is preserved.

Second, the visualisation relies on a sampled approximation of the decision space. In this work, Pareto landscape is constructed on a regular grid, and its interpretability therefore depends on the sampling resolution. In particular, when the grid is sufficiently fine, major structural patterns can usually be observed, whereas under very coarse sampling, these patterns may become blurred or may fail to appear.

Third, Pareto landscape does not explicitly identify locally efficient solutions. Instead, local structural features must be inferred from the topology of the dominance-level surface. By contrast, search-based visualisation methods such as gradient field heatmaps and PLOT can explicitly reveal basins of attraction and locally efficient solutions through gradient descent mechanisms [16,17]. Consequently, the Pareto landscape is more suited to providing a global structural overview rather than a precise characterisation of local optimality structures.

Finally, although dominance-based measures such as non-domination levels tend to be more stable than dominance-count measures when the overall front structure is preserved, a systematic sensitivity analysis with respect to sampling density and alternative sampling strategies remains an important direction for future work.

As discussed in Section 6, the Pareto landscape should not be interpreted as replacing existing landscape visualisation methods. Rather, it provides a complementary perspective. Methods such as cost landscape, constrained cost landscape, gradient field heatmaps, and PLOT capture different aspects of dominance relations, infeasibility, and local optimality structure. In particular, for constrained problems, the Pareto landscape does not

reveal variation in infeasibility across the decision space as directly as a constrained cost landscape. The choice of method therefore depends on the analysis objective.

8. Conclusions

In this work, we introduced the Pareto landscape, a non-search-based problem landscape visualisation that assigns solution quality through non-domination levels. This provides a natural extension of the classical fitness landscape to multi-objective optimisation. We discussed the similarities and differences between them. The Pareto landscape differs from the latter in that it provides a “relative” measure of solution quality within the set considered. However, like a fitness landscape, it can effectively capture a range of problem characteristics (e.g., optimality, multimodality, neutrality, level of ruggedness, and location of attraction basins) and spatial features of the decision space (e.g., peaks, valleys, plateaus, ridges and funnels).

We used the Pareto landscape to visualise various multi-objective optimisation problems, including well-established benchmarks, constrained and multimodal functions, and real-world problems. Our findings were as follows.

- Many established benchmark suites (e.g., ZDT [22], DTLZ [23] and WFG [24]) exhibit human-made features that may be exploited by search configurations, for example, by separately searching for optimum on different variables or employing specific search step-sizes during the search process. The BBOB suite may better represent challenging real-world optimisation scenarios with diverse, irregular landscapes.
- For constrained MOPs, the scale and pattern of infeasibility vary dramatically—from small scattered pockets in CTP4 to large blocking regions in CTP6 [84], and from regular periodic strips in CMOP2 [87] to irregular corridors in CF5 [85]. These varied constraints fragment the decision space into disconnected feasible regions, leading to distinct structural consequences: multiple isolated Pareto optimal regions (TNK [83], CF5 [85], C-DTLZ1 [86]) that appear as separate peaks in the landscape, additional local Pareto optimal regions at feasibility boundaries (CTP6 [84], CMOP7 [87]), and narrow feasible bands near the Pareto optimal regions (SRN [82], CF6 [85]).
- Most multimodal benchmarks are concerned with global optimality (e.g., Omnitest [90], SYM-PART [91], and SS-UF [27]), used to test search algorithms to find multiple identical global optimal regions. Very few focus on local optimality, i.e., test search algorithms not being trapped in local optimal regions. However, the latter can be much more common in the real world.
- Many low-dimensional real-world problems (e.g., RWCMOP [95] and RE [93]) often exhibit comparatively simpler structures due to reduced non-linear coupling in their formulations. However, map-based problems [98] can have strong irregularities, featuring isolated global solutions and multiple locally efficient regions.

A comparison with gradient field heatmaps (GFH) [16], PLOT [17], cost landscape [20], and constrained cost landscape [11] further demonstrates that Pareto landscape offers complementary insight: it provides a dominance-consistent, search-free global view across different classes of multi-objective optimisation problems, whereas the other methods are more suitable when information about infeasibility variation, locally efficient solutions, or local basin structure is required.

Lastly, it is worth pointing out that a key limitation, shared with all native decision space visualisation methods, is its restriction to problems with two or three decision variables. Extending the approach to higher-dimensional decision spaces, or coupling Pareto dominance with suitable dimensionality-reduction techniques, is an important direction for future research.

Author Contributions: Conceptualisation, M.L.; Methodology, Z.C., Z.L. and M.L.; Formal analysis, Z.C.; Writing—original draft, Z.C.; Writing—review & editing, Z.L. and M.L.; Visualisation, Z.C.; Supervision, M.L.; Project administration, M.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data and source code used to generate the results and visualisations in this study are publicly available at: <https://github.com/zxc990/pareto-landscape-visualisation> (accessed on 27 December 2025).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Eberhard, K. The effects of visualization on judgment and decision-making: A systematic literature review. *Manag. Rev. Q.* **2023**, *73*, 167–214. [CrossRef]
2. Dy, B.; Ibrahim, N.; Poorthuis, A.; Joyce, S. Improving visualization design for effective multi-objective decision making. *IEEE Trans. Vis. Comput. Graph.* **2021**, *28*, 3405–3416. [CrossRef] [PubMed]
3. Schäpermeier, L.; Grimme, C.; Kerschke, P. Plotting impossible? Surveying visualization methods for continuous multi-objective benchmark problems. *IEEE Trans. Evol. Comput.* **2022**, *26*, 1306–1320. [CrossRef]
4. Miettinen, K. Survey of methods to visualize alternatives in multiple criteria decision making problems. *OR Spectr.* **2014**, *36*, 3–37. [CrossRef]
5. Ibrahim, A.; Rahnamayan, S.; Martin, M.V.; Deb, K. 3D-RadVis: Visualization of Pareto front in many-objective optimization. In *Proceedings of the 2016 IEEE Congress on Evolutionary Computation (CEC)*; IEEE: Piscataway, NJ, USA, 2016; pp. 736–745.
6. Filipič, B.; Tušar, T. A taxonomy of methods for visualizing pareto front approximations. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Kyoto, Japan, 15–19 July 2018; pp. 649–656.
7. Hakanen, J.; Gold, D.; Miettinen, K.; Reed, P.M. Visualisation for decision support in many-objective optimisation: State-of-the-art, guidance and future directions. In *Many-Criteria Optimization and Decision Analysis: State-of-the-Art, Present Challenges, and Future Perspectives*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 181–212.
8. Walker, D.; Fieldsend, J.; Everson, R. Visualising many-objective populations. In *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*, Philadelphia, PA, USA, 7–11 July 2012; pp. 451–458.
9. Abdelaal, M.; Galuschka, M.; Zorn, M.; Kannenberg, F.; Menges, A.; Wortmann, T.; Weiskopf, D.; Kurzhals, K. Visual analysis of fitness landscapes in architectural design optimization. *Vis. Comput.* **2024**, *40*, 4927–4940. [CrossRef]
10. Zou, F.; Chen, D.; Liu, H.; Cao, S.; Ji, X.; Zhang, Y. A survey of fitness landscape analysis for optimization. *Neurocomputing* **2022**, *503*, 129–139. [CrossRef]
11. Tušar, T.; Vodopija, A.; Filipič, B. Visual Exploration of the Effect of Constraint Handling in Multiobjective Optimization. In *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2023; pp. 3–16.
12. Ochoa, G.; Tomassini, M.; Vérel, S.; Darabos, C. A study of NK landscapes’ basins and local optima networks. In *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation*, Atlanta, GA, USA, 12–16 July 2008; pp. 555–562.
13. Liefoghe, A.; Derbel, B.; Verel, S.; López-Ibáñez, M.; Aguirre, H.; Tanaka, K. On Pareto local optimal solutions networks. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 232–244.
14. Liefoghe, A.; Ochoa, G.; Verel, S.; Derbel, B. Pareto local optimal solutions networks with compression, enhanced visualization and expressiveness. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Lisbon, Portugal, 15–19 July 2023; pp. 713–721.
15. Diao, Y.; Li, C.; Zeng, S.; Yang, S.; Coello, C.A.C. Nearest-Better Network for Fitness Landscape Analysis of Continuous Optimization Problems. *IEEE Trans. Evol. Comput.* **2024**, *29*, 2089–2103. [CrossRef]
16. Kerschke, P.; Grimme, C. An expedition to multimodal multi-objective optimization landscapes. In *Proceedings of the 9th International Conference Evolutionary Multi-Criterion Optimization (EMO)*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 329–343.
17. Schäpermeier, L.; Grimme, C.; Kerschke, P. One PLOT to show them all: Visualization of efficient sets in multi-objective landscapes. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 154–167.
18. Fieldsend, J.E.; Chugh, T.; Allmendinger, R.; Miettinen, K. A feature rich distance-based many-objective visualisable test problem generator. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Prague, Czech Republic, 13–17 July 2019; pp. 541–549.

19. Kerschke, P.; Wang, H.; Preuss, M.; Grimme, C.; Deutz, A.; Trautmann, H.; Emmerich, M. Towards analyzing multimodality of continuous multiobjective landscapes. In *Proceedings of the Parallel Problem Solving from Nature—PPSN XIV: 14th International Conference, Edinburgh, UK, 17–21 September 2016*; Proceedings 14; Springer: Berlin/Heidelberg, Germany, 2016; pp. 962–972.
20. Fonseca, C.M.M.d. Multiobjective Genetic Algorithms with Application to Control Engineering Problems. Ph.D. Thesis, University of Sheffield, Sheffield, UK, 1995.
21. Liang, Z.; Cui, Z.; Li, M. Pareto landscape: Visualising the landscape of multi-objective optimisation problems. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*; Springer: Berlin/Heidelberg, Germany, 2024; pp. 299–315.
22. Zitzler, E.; Deb, K.; Thiele, L. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evol. Comput.* **2000**, *8*, 173–195. [[CrossRef](#)]
23. Deb, K.; Thiele, L.; Laumanns, M.; Zitzler, E. Scalable test problems for evolutionary multiobjective optimization. In *Evolutionary Multiobjective Optimization: Theoretical Advances and Applications*; Abraham, A., Jain, L., Goldberg, R., Eds.; Springer: London, UK, 2005; pp. 105–145.
24. Huband, S.; Hingston, P.; Barone, L.; While, L. A review of multiobjective test problems and a scalable test problem toolkit. *IEEE Trans. Evol. Comput.* **2006**, *10*, 477–506. [[CrossRef](#)]
25. Brockhoff, D.; Auger, A.; Hansen, N.; Tušar, T. Using well-understood single-objective functions in multiobjective black-box optimization test suites. *Evol. Comput.* **2022**, *30*, 165–193. [[CrossRef](#)]
26. Grimme, C.; Kerschke, P.; Aspar, P.; Trautmann, H.; Preuss, M.; Deutz, A.H.; Wang, H.; Emmerich, M. Peeking beyond peaks: Challenges and research potentials of continuous multimodal multi-objective optimization. *Comput. Oper. Res.* **2021**, *136*, 105489. [[CrossRef](#)]
27. Liang, J.J.; Yue, C.; Qu, B.Y. Multimodal multi-objective optimization: A preliminary study. In *Proceedings of the 2016 IEEE Congress on Evolutionary Computation (CEC)*; IEEE: Piscataway, NJ, USA, 2016; pp. 2454–2461.
28. Li, W.; Yao, X.; Zhang, T.; Wang, R.; Wang, L. Hierarchy ranking method for multimodal multiobjective optimization with local Pareto fronts. *IEEE Trans. Evol. Comput.* **2022**, *27*, 98–110. [[CrossRef](#)]
29. Ren, S.; Qiu, Z.; Bian, C.; Li, M.; Qian, C. Maintaining diversity provably helps in evolutionary multimodal optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, Jeju, Republic of Korea, 3–9 August 2024; pp. 7012–7020.
30. Fliege, J.; Svaiter, B.F. Steepest descent methods for multicriteria optimization. *Math. Methods Oper. Res.* **2000**, *51*, 479–494. [[CrossRef](#)]
31. Goldberg, D.E. *Genetic Algorithms in Search, Optimization and Machine Learning*; Addison-Wesley: Boston, MA, USA, 1989.
32. Brockhoff, D.; Tušar, T. GECCO 2023 tutorial on benchmarking multiobjective optimizers 2.0. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, Lisbon, Portugal, 15–19 July 2023; pp. 1183–1212.
33. Chambers, J.M.; Kleiner, B. 10 graphical techniques for multivariate data and for clustering. *Handb. Stat.* **1982**, *2*, 209–244.
34. Carr, D.; Nicholson, W. *Evaluation of Graphical Techniques for Data in Dimensions 3 to 5: Scatter Plot Matrix, Glyph and Stereo Examples*; Technical Report; Pacific Northwest Labs.: Richland, WA, USA, 1985.
35. Geoffrion, A.M.; Dyer, J.S.; Feinberg, A. An interactive approach for multi-criterion optimization, with an application to the operation of an academic department. *Manag. Sci.* **1972**, *19*, 357–368. [[CrossRef](#)]
36. Schilling, D.A.; Reville, C.; Cohon, J. An approach to the display and analysis of multiobjective problems. *Socio Econ. Plan. Sci.* **1983**, *17*, 57–63. [[CrossRef](#)]
37. Pryke, A.; Mostaghim, S.; Nazemi, A. Heatmap visualization of population based multi objective algorithms. In *Proceedings of the Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2007; pp. 361–375.
38. Inselberg, A. The plane with parallel coordinates. *Vis. Comput.* **1985**, *1*, 69–91. [[CrossRef](#)]
39. Walker, D.J.; Everson, R.; Fieldsend, J.E. Visualizing mutually nondominating solution sets in many-objective optimization. *IEEE Trans. Evol. Comput.* **2012**, *17*, 165–184. [[CrossRef](#)]
40. Walker, D.J.; Everson, R.M.; Fieldsend, J.E. Visualisation and ordering of many-objective populations. In *Proceedings of the IEEE Congress on Evolutionary Computation*; IEEE: Piscataway, NJ, USA, 2010; pp. 1–8.
41. Li, M.; Zhen, L.; Yao, X. How to read many-objective solution sets in parallel coordinates [educational forum]. *IEEE Comput. Intell. Mag.* **2017**, *12*, 88–100. [[CrossRef](#)]
42. Zhen, L.; Li, M.; Cheng, R.; Peng, D.; Yao, X. Adjusting parallel coordinates for investigating multi-objective search. In *Proceedings of the Asia-Pacific Conference on Simulated Evolution and Learning*; Springer: Piscataway, NJ, USA, 2017; pp. 224–235.
43. Manas, M. Graphical methods of multicriterial optimization. *ZAMM* **1982**, *62*, T375–T377.
44. Chambers, J.M. *Graphical Methods for Data Analysis*; Chapman and Hall/CRC: Boca Raton, FL, USA, 2018.
45. Kasanen, E.; Östermark, R.; Zeleny, M. Gestalt system of holistic graphics: New management support view of MCDM. *Comput. Oper. Res.* **1991**, *18*, 233–239. [[CrossRef](#)]
46. Chernoff, H. Graphical Representations as a Discipline. In *Information Linkage Between Applied Mathematics and Industry*; Academic Press: Cambridge, MA, USA, 1979; pp. 49–59.

47. Kleiner, B.; Hartigan, J.A. Representing points in many dimensions by trees and castles. *J. Am. Stat. Assoc.* **1981**, *76*, 260–269. [\[CrossRef\]](#)
48. Johnson, B.; Shneiderman, B. Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *Proceedings of the 1991 Proceeding Visualization. IEEE Computer Society*; IEEE: Piscataway, NJ, USA, 1991; pp. 284–285.
49. Stasko, J.; Catrambone, R.; Guzdial, M.; McDonald, K. An evaluation of space-filling information visualizations for depicting hierarchical structures. *Int. J. Hum. Comput. Stud.* **2000**, *53*, 663–694. [\[CrossRef\]](#)
50. Walker, D.J. Visualising multi-objective populations with treemaps. In *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation*, Madrid, Spain, 11–15 July 2015; pp. 963–970.
51. Walker, D.J. Visualisation with treemaps and sunbursts in many-objective optimisation. *Genet. Program. Evolvable Mach.* **2018**, *19*, 421–452. [\[CrossRef\]](#)
52. Wold, S.; Esbensen, K.; Geladi, P. Principal component analysis. *Chemom. Intell. Lab. Syst.* **1987**, *2*, 37–52. [\[CrossRef\]](#)
53. Kruskal, J.B.; Wish, M. *Multidimensional Scaling*; SAGE Publications, Inc.: Newbury Park, CA, USA, 1978.
54. Kohonen, T. *Self-Organizing Maps*; Springer: Berlin/Heidelberg, Germany, 1995.
55. Hoffman, P.; Grinstein, G.; Marx, K.; Grosse, I.; Stanley, E. DNA visual and analytic data mining. In *Proceedings of the Visualization'97 (Cat. No. 97CB36155)*; IEEE: Piscataway, NJ, USA, 1997; pp. 437–441.
56. Zhen, L.; Li, M.; Peng, D.; Yao, X. Objective reduction for visualising many-objective solution sets. *Inf. Sci.* **2020**, *512*, 278–294. [\[CrossRef\]](#)
57. Tušar, T.; Filipič, B. Visualization of Pareto front approximations in evolutionary multiobjective optimization: A critical review and the prosecution method. *IEEE Trans. Evol. Comput.* **2014**, *19*, 225–245. [\[CrossRef\]](#)
58. Tušar, T.; Filipič, B. Visualizing 4D approximation sets of multiobjective optimizers with prosecutions. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, Dublin, Ireland, 12–16 July 2011; pp. 737–744.
59. De Lorenzo, A.; Medvet, E.; Tušar, T.; Bartoli, A. An analysis of dimensionality reduction techniques for visualizing evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Prague, Czech Republic, 13–17 July 2019; pp. 1864–1872.
60. Walter, M.J.; Walker, D.J.; Craven, M.J. Visualising evolution history in multi-and many-objective optimisation. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 299–312.
61. Walter, M.J.; Walker, D.J.; Craven, M.J. An explainable visualisation of the evolutionary search process. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Boston, MA, USA, 9–13 July 2022; pp. 1794–1802.
62. Walter, M.J.; Walker, D.J.; Craven, M.J. Visualizing population dynamics to examine algorithm performance. *IEEE Trans. Evol. Comput.* **2022**, *26*, 1501–1510. [\[CrossRef\]](#)
63. Li, M.; Grosan, C.; Yang, S.; Liu, X.; Yao, X. Multiline distance minimization: A visualized many-objective test problem suite. *IEEE Trans. Evol. Comput.* **2017**, *22*, 61–78. [\[CrossRef\]](#)
64. Li, M.; Yang, S.; Liu, X. A test problem for visual investigation of high-dimensional multi-objective search. In *Proceedings of the 2014 IEEE Congress on Evolutionary Computation (CEC)*; IEEE: Piscataway, NJ, USA, 2014; pp. 2140–2147.
65. Tušar, T.; Filipič, B. Scaling and visualizing multiobjective optimization test problems with knees. *Knee* **2012**, *40*, f2.
66. Fieldsend, J.E.; Alyahya, K. Visualising the landscape of multi-objective problems using local optima networks. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Prague, Czech Republic, 13–17 July 2019; pp. 1421–1429.
67. Volz, V.; Naujoks, B.; Kerschke, P.; Tušar, T. Single-and multi-objective game-benchmark for evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Prague, Czech Republic, 13–17 July 2019; pp. 647–655.
68. Brockhoff, D.; Auger, A.; Hansen, N.; Tušar, T. BBOB Biobj Visualizations. 2021. Available online: <https://numbbo.github.io/bbob-biobj/vis> (accessed on 27 December 2025).
69. Wright, S. The Roles of Mutation, Inbreeding, Crossbreeding, and Selection in Evolution. In *Proceedings of the Sixth International Congress of Genetics*, Ithaca, NY, USA, 24–31 August 1932; pp. 356–366.
70. Jones, T. *Evolutionary Algorithms, Fitness Landscapes and Search*. Ph.D. Thesis, The University of New Mexico, Albuquerque, NM, USA, 1995.
71. Ochoa, G.; Malan, K. Recent advances in fitness landscape analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Prague, Czech Republic, 13–17 July 2019; pp. 1077–1094.
72. Ishibuchi, H.; Hitotsuyanagi, Y.; Tsukamoto, N.; Nojima, Y. Many-objective test problems to visually examine the behavior of multiobjective evolution in a decision space. In *Proceedings of the International Conference on Parallel Problem Solving from Nature (PPSN)*, Kraków, Poland, 11–15 September 2010; pp. 91–100.
73. Fieldsend, J.E.; Chugh, T.; Allmendinger, R.; Miettinen, K. A visualizable test problem generator for many-objective optimization. *IEEE Trans. Evol. Comput.* **2021**, *26*, 1–11. [\[CrossRef\]](#)
74. Singh, H.K.; Isaacs, A.; Ray, T.; Smith, W. A study on the performance of substitute distance based approaches for evolutionary many-objective optimization. In *Simulated Evolution and Learning*; Springer: Berlin/Heidelberg, Germany, 2008; pp. 401–410.

75. Li, M.; Yang, S.; Liu, X.; Shen, R. A comparative study on evolutionary algorithms for many-objective optimization. In Proceedings of the 7th International Conference on Evolutionary Multi-Criterion Optimization (EMO), Sheffield, UK, 19–22 March 2013; pp. 261–275.
76. Ishibuchi, H.; Akedo, N.; Ohyanagi, H.; Nojima, Y. Behavior of EMO algorithms on many-objective optimization problems with correlated objectives. In *Proceedings of the 2011 IEEE Congress of Evolutionary Computation (CEC)*; IEEE: Piscataway, NJ, USA, 2011; pp. 1465–1472.
77. Liu, Y.; Ishibuchi, H.; Nojima, Y.; Masuyama, N.; Shang, K. A double-niched evolutionary algorithm and its behavior on polygon-based problems. In *Proceedings of the 15th International Conference on Parallel Problem Solving from Nature (PPSN)*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 262–273.
78. Suresh, A.; Deb, K. On a Better Understanding of Unique Identifiers of Pareto Solutions for Multi-criterion Optimization, Visualization, and Decision-making. In Proceedings of the 2024 IEEE Congress on Evolutionary Computation (CEC), Yokohama, Japan, 30 June–5 July 2024; pp. 1–8. [\[CrossRef\]](#)
79. Schäpermeier, L.; Grimme, C.; Kerschke, P. moPLOT Website. 2021. Available online: <https://schaepermeier.shinyapps.io/moPLOT/> (accessed on 27 December 2025).
80. Schäpermeier, L.; Grimme, C.; Kerschke, P. To boldly show what no one has seen before: A dashboard for visualizing multi-objective landscapes. In *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2021; pp. 632–644.
81. Liang, J.; Ban, X.; Yu, K.; Qu, B.; Qiao, K.; Yue, C.; Chen, K.; Tan, K.C. A survey on evolutionary constrained multiobjective optimization. *IEEE Trans. Evol. Comput.* **2022**, *27*, 201–221. [\[CrossRef\]](#)
82. Srinivas, N.; Deb, K. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evol. Comput.* **1994**, *2*, 221–248. [\[CrossRef\]](#)
83. Tanaka, M.; Watanabe, H.; Furukawa, Y.; Tanino, T. GA-based decision support system for multicriteria optimization. In *Proceedings of the 1995 IEEE International Conference on Systems, Man and Cybernetics. Intelligent Systems for the 21st Century*; IEEE: Piscataway, NJ, USA, 1995; Volume 2, pp. 1556–1561.
84. Deb, K.; Pratap, A.; Meyarivan, T. Constrained test problems for multi-objective evolutionary optimization. In *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2001; pp. 284–298.
85. Zhang, Q.; Zhou, A.; Zhao, S.; Suganthan, P.N.; Liu, W.; Tiwari, S. *Multiobjective Optimization Test Instances for the CEC 2009 Special Session and Competition*; Special Session on Performance Assessment of Multi-Objective Optimization Algorithms, Technical Report; University of Essex: Colchester, UK; Nanyang Technological University: Singapore, 2008; Volume 264, pp. 1–30.
86. Jain, H.; Deb, K. 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.* **2013**, *18*, 602–622. [\[CrossRef\]](#)
87. Fan, Z.; Li, W.; Cai, X.; Li, H.; Hu, K.; Yin, H. Difficulty controllable and scalable constrained multi-objective test problems. In *Proceedings of the 2015 International Conference on Industrial Informatics-Computing Technology, Intelligent Technology, Industrial Information Integration*; IEEE: Piscataway, NJ, USA, 2015; pp. 76–83.
88. Xu, Z.; Zhang, K.; Del Ser, J.; Li, M.; Xu, X.; He, J.; Wu, N. Multi-Objective Optimization for Multimodal Multi-Objective Multi-Point Shortest Path Problem Considering Unforeseeable Road Eventualities. *IEEE Trans. Intell. Transp. Syst.* **2025**, *26*, 8622–8640. [\[CrossRef\]](#)
89. Yue, C.; Qu, B.; Liang, J. A multiobjective particle swarm optimizer using ring topology for solving multimodal multiobjective problems. *IEEE Trans. Evol. Comput.* **2017**, *22*, 805–817. [\[CrossRef\]](#)
90. Deb, K.; Tiwari, S. Omni-optimizer: A procedure for single and multi-objective optimization. In *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2005; pp. 47–61.
91. Rudolph, G.; Naujoks, B.; Preuss, M. Capabilities of EMOA to detect and preserve equivalent Pareto subsets. In *Proceedings of the International Conference on Evolutionary Multi-Criterion Optimization*; Springer: Berlin/Heidelberg, Germany, 2007; pp. 36–50.
92. Preuss, M.; Naujoks, B.; Rudolph, G. Pareto set and EMOA behavior for simple multimodal multiobjective functions. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*; Springer: Berlin/Heidelberg, Germany, 2006; pp. 513–522.
93. Tanabe, R.; Ishibuchi, H. An easy-to-use real-world multi-objective optimization problem suite. *Appl. Soft Comput.* **2020**, *89*, 106078. [\[CrossRef\]](#)
94. Zapotecas-Martínez, S.; García-Nájera, A.; Menchaca-Méndez, A. Engineering applications of multi-objective evolutionary algorithms: A test suite of box-constrained real-world problems. *Eng. Appl. Artif. Intell.* **2023**, *123*, 106192. [\[CrossRef\]](#)
95. Kumar, A.; Wu, G.; Ali, M.Z.; Luo, Q.; Mallipeddi, R.; Suganthan, P.N.; Das, S. A benchmark-suite of real-world constrained multi-objective optimization problems and some baseline results. *Swarm Evol. Comput.* **2021**, *67*, 100961. [\[CrossRef\]](#)
96. He, C.; Tian, Y.; Wang, H.; Jin, Y. A repository of real-world datasets for data-driven evolutionary multiobjective optimization. *Complex Intell. Syst.* **2020**, *6*, 189–197.

97. Palakonda, V.; Kang, J.M.; Jung, H. Benchmarking Real-World Many-Objective Problems: A Problem Suite with Baseline Results. *IEEE Access* **2024**, *12*, 49275–49290. [[CrossRef](#)]
98. Ishibuchi, H.; Akedo, N.; Nojima, Y. A many-objective test problem for visually examining diversity maintenance behavior in a decision space. In Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation, Dublin, Ireland, 12–16 July 2011; pp. 649–656.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.