

Pareto Landscape: Visualising the Landscape of Multi-Objective Optimisation Problems

Zimin Liang[†], Zhiji Cui[†], and Miqing Li^{*}

University of Birmingham, Birmingham B15 2TT, UK

Abstract. Fitness landscape is a valuable framework to understand optimisation problems. In single-objective optimisation, by displaying fitness landscape in a 3D space with the “height” representing the fitness (objective function value) of solutions, one can easily comprehend a variety of problem characteristics (optimality, multi-modality, level of ruggedness, etc) and spatial features of the search space (basin, ridge, funnel, etc). However, such straightforward visualisation cannot be directly extended to the multi-objective optimisation case in which a solution corresponds to a vector of values on multiple objective functions. In this paper, we make an attempt to address this issue. Instead of objective function values, we use the Pareto dominance relation to stratify solutions, introducing a method we term *Pareto landscape* for visualising multi-objective problem landscape. We compare Pareto landscape with well-established fitness landscape visualisation methods, including *cost landscape*, *gradient field heatmap* and *PLOT*, and show that Pareto landscape can capture problem characteristics that the other methods cannot do. Lastly, we present the Pareto landscapes of commonly used benchmark problems (ZDT, DTLZ, WFG and BBOB) in the domain, and discuss their features and characteristics.

1 Introduction

Fitness landscape plots serve as a valuable framework for understanding optimisation problems. They facilitate an intuitive grasp of complex search spaces (also known as decision spaces) and challenges confronted by search heuristics, inspiring potential improvements and development of new optimisation algorithms.

In the context of multi-objective optimisation, there is a growing interest in developing visualisation tools for presenting fitness landscapes [28,3]. This ranges from simply integrating landscape information (e.g., local Pareto optimal sets) into the objective space [28] to designing networks that graphically represent optimal solutions [21,22] and sets [7], as well as the behaviour of search algorithms [24]. While these visualisation tools effectively capture the structure of (local) optimal solutions/sets and their connectedness, the compression and abstract representation involved make it difficult for people to envision the spatial location of solutions/sets and their associated adjacent structures.

^{*}Corresponding author.

[†]These authors contributed equally to this work.

Another approach, which can overcome the above issue, is to directly display the decision space of an optimisation problem with an additional element (e.g., colour) representing solution quality [28]. This encompasses early attempts like cost landscape (aka dominance ratio) [11] and recent efforts such as multi-contour plot [18], gradient field heatmaps [17], line cuts [30,2], local dominance landscape [8], and PLOT [26]. These methods are similar to the fitness landscape plots in single-objective optimisation, where the native decision space is considered, allowing them to present important problem characteristics such as location, shape and size of attraction basins. However, they may not be very accurate in some cases due to the effect of the behaviour of the search algorithm used to locate local optima (e.g., in gradient field heatmaps [17] and PLOT [26]), or the criterion considered to measure the “fitness” of solutions (e.g., in cost landscape [11]), which we will show later on.

In this paper, we attempt to develop a fitness landscape visualisation method for multi-objective optimisation, termed Pareto landscape, that can accurately represent problem characteristics and spatial features. The idea of Pareto landscape is simple. It considers the level of the Pareto non-domination to which a solution belongs when determining its fitness. To evaluate Pareto landscape, we first compare it with its counterpart in single-objective optimisation to identify their similarities and differences. Then, by utilising a class of test functions flexible for embedding rich features [19,15,9], we demonstrate that Pareto landscape can capture various problem characteristics (e.g., multi-modality and neutrality), as well as spatial features of the decision space (e.g., valleys, plateaus, ridges and funnels). Moreover, we compare Pareto landscape with three relevant landscape visualisation methods and discuss their differences. Lastly, we use Pareto landscape to visualise commonly-used multi-objective benchmark functions in the domain, including ZDT [32], DTLZ [6], WFG [13] and BBOB (bbob-biobj) [2], to delineate their features and characteristics, particularly the challenges they may pose for search heuristics.

2 Preliminaries

2.1 Terminology in Multi-Objective Optimisation

For multi-objective optimisation problems (MOPs), without loss of generality, we seek to simultaneously minimise m objective functions $f(x) = (f_1(x), \dots, f_m(x))$, where $f(x)$ ($x \in X$) is a mapping of an n -dimensional variable vector to an m -dimensional objective vector; namely, $f : X \rightarrow Z$ where $X \subseteq \mathbb{R}^n$ and $Z \subseteq \mathbb{R}^m$.

Unlike single-objective optimisation, where there are only three relations between solutions with respect to their quality, *better*, *worse* and *equal*, in multi-objective optimisation there is a situation that two solutions are not comparable; namely, one solution is better on some objective(s) and the other is better on some other objective(s). This situation is termed as Pareto non-domination. Formally, a decision vector x is said to (Pareto) dominate another x' , denoted by $x \prec x'$, iff

$$f_i(x) \leq f_i(x') \quad \forall i \in \{1, \dots, m\} \wedge f(x) \neq f(x') \quad (1)$$

For a solution set S , a solution x ($x \in S$) is called a (*Pareto*) *nondominated solution* if there does not exist any $x' \in S$ such that $x' \prec x$. The set of such nondominated solutions is called the *nondominated set* of S , and its mapping in the objective space is called the *nondominated front*. For an MOP, in the case that S represents all possible feasible solutions (i.e., $S = X$), the nondominated set is called the *Pareto optimal set* (or *global efficient set*). The mapping of the Pareto optimal set in the objective space is called the *Pareto optimal front*.

Besides the global efficient set, there usually exist some local optima in the search space, namely the *local efficient sets*, each consisting of a set of *local efficient points*. In continuous multi-objective optimisation, a solution x is said to be local efficient if there does not exist $y \in B_\epsilon$ that $y \prec x$ [4], where $B_\epsilon \subset X$ represents an ϵ ball for a sufficiently small $\epsilon > 0$, meaning that all neighbours of x within a sufficiently small distance do not dominate x .

An important property in continuous multi-objective optimisation is the *multi-objective gradient* (MOG) [10]. MOG is a vector that is oriented towards a descent direction by the sum of its normalised single-objective components. It has a desired characteristic that $\text{MOG} = 0$ if a solution is a local or global efficient solution. Formally, MOG is defined as

$$\nabla f(x) = \sum_{i=1}^m \alpha_i^* \nabla f_i(x) \quad (2)$$

where α^* is a weight vector determined based on

$$\alpha^* = \arg \min_{\alpha} \left\{ \left\| \sum_{i=1}^m \alpha_i \nabla f_i(x) \right\| \mid \alpha_i \geq 0, \sum_{i=1}^m \alpha_i = 1 \right\} \quad (3)$$

2.2 Related Work

Amongst various fitness landscape visualisation methods in multi-objective optimisation [28,3], we consider the ones that work on the native decision space (rather than a compressed, abstract space). Specifically, we consider three visualisation methods, cost landscape [11], gradient field heatmap [17] and PLOT [26], which are relevant to the proposed method. We now briefly introduce them (in Section 4 we will compare our method against them).

Cost Landscape [11]. Proposed in Carlos M. Fonseca’s PhD thesis, cost landscape is a useful visualisation tool for observing the global structure of a multi-objective problem. In some visualisation platforms (e.g., BBOB visualisation website [1]), it is also called dominance ratio. In cost landscape, the fitness (which can be reflected by colour or height in the plot) of a point is the number of points that dominate it in the point set considered (i.e., *Pareto rank* [11]).

Gradient Field Heatmap (GFH) [17]. GFH, proposed by Kerschke and Grimme, is the first attempt to employ gradient to visualise the landscape of multi-objective problems. In GFH, the fitness (represented by colour) of a point is obtained by the cumulative lengths of MOGs (Eq. (2), the sum of normalised

single objective gradients) on the path to a locally efficient point. More specifically, starting from a point and a variable to store the cumulative MOGs, GFH performs a search that moves to one of its 8 neighbouring points according to its MOG, and at the same time, accumulates the MOG of each step. The move is iterated until it hits a local efficient point. Then the fitness of the starting point is the cumulative length of the MOGs. GFH is considered a good tool to capture the local structures of the problem [28].

Plot of Landscapes with Optimal Trade-offs (PLOT) [26]. More recently, Schäpermeier et al. developed PLOT, which can be seen as an improved version of GFH in combination with cost landscape. Similar to GFH, PLOT employs the cumulative length of MOGs as its fitness, but with a grey-scale colour scheme. The global and local efficient points are highlighted by the colour on the basis of the Pareto rank considered in cost landscape. As such, it is very clear to see global/local optimality structure of the problem.

3 The Proposed Pareto Landscape

Classical fitness landscape, introduced by the geneticist Sewall Wright in 1930s [31], is a 3D “mountainous” landscape in which the genotypes (decision variables) of solutions are organised in the x - y plane and the phenotype (fitness) is plotted as the height (i.e., on the z axis). Fitness landscape plots have been widely used in optimisation and evolutionary computation [16,25]. They enable an intuitive understanding of diverse problem characteristics and spatial features, encompassing peaks, valleys, plateaus, ridges, and funnels, as well as optimality-related aspects like multi-modality, level of ruggedness, and the shape/location of (local) optimal solutions along with their attraction basins.

Yet, such straightforward visualisation cannot be directly extended to the multi-objective case in which a solution corresponds to a vector consisting of values of multiple objective functions. A possible solution to this issue is to utilise the Pareto dominance relation to sort and rank the solutions and then plot their ranks as the height, instead of objective function values, in the landscape.

The non-dominated sorting procedure [12] is a well-established way to sort solutions in multi-objective optimisation, e.g., used in NSGA-II [5]. It stratifies a set of solutions into many levels based on their Pareto dominance relation, with level 1 meaning nondominated solutions of the set, level 2 meaning new nondominated solutions after removing the level-1 solutions, and so on. In the proposed Pareto landscape, the “fitness” of a solution is its level after the non-dominated sorting procedure, termed the non-domination level. Figure 1 gives the Pareto landscape[†] of a 4-objective optimisation problem called multi-point distance minimisation problems (MP-DMPs). MP-DMPs, introduced by [19] and extended and generalised by [15,9], are class of increasingly popular test functions which are capable of being embedded with a variety of features (see [9] for a survey). One prominent feature of MP-DMPs is that they inherently have a 2D

[†]In this paper, the points to plot Pareto landscape are those evenly distributed on a 500×500 grid unless specifically stated otherwise.




Fig. 1. Pareto landscape of a four-objective multi-point distance minimisation problem (MP-DMP) [15,9] whose four objectives of a point are the Euclidean distances of a point in the 2D space to the four vertices of a square, respectively. The decision space of the MP-DMP is $x_1, x_2 \in [0, 1]$, and the Pareto optimal set of the problem is the square (shaded region in Figure 1(a)).

decision space (despite easily accommodating an arbitrary number of objectives), thus allowing the Pareto optimal solutions to be visually identified.

MP-DMPs, in its basic form, minimise the Euclidean distance of a solution to the vertices of a given regular polygon, where the distance to one of these vertices is treated as an individual objective. It can be easily derived that the Pareto optimal region of an MP-DMP is the regular polygon.

Figure 1(b) plots the Pareto landscape of the MP-DMP in Figure 1(a), where the Pareto optimal solutions are highlighted in green for facilitating observation. As can be seen in the figure, Pareto landscape can well reflect the problem's landscape – points being farther away from the optimal region, gradually and symmetrically, take lower altitude (i.e., higher non-domination levels), with the four corners of the decision space having the lowest altitude touching the x_1 - x_2 plane. Next, we will use a toy example to explain the characteristics of Pareto landscape, in comparison with fitness landscape in single-objective optimisation.

3.1 Properties and Characteristics of Pareto Landscape

Let us consider a set of 10 solutions (**a** – **j**) with one decision variable (x) and two minimisation objectives (f_1, f_2). In the format of $(x|f_1, f_2)$, these solutions are **a** = (0.1|1, 4), **b** = (0.2|1, 3), **c** = (0.3|3, 3), **d** = (0.4|3, 4), **e** = (0.5|2, 3), **f** = (0.6|6, 3), **g** = (0.7|2, 2), **h** = (0.8|3, 1), **i** = (0.9|3, 3), and **j** = (1.0|4, 1). Figure 2 gives their Pareto landscape as well as the scatter plot in the objective space. As can be seen in the figure, the 10 solutions can be divided into four non-domination levels. Solutions **b, g, h** form the first level, solutions **a, e, j** the second level, solutions **c, i** the third level, and solutions **d, f** the last level. Looking closer at the figure, one may find the following properties of Pareto landscape, in comparison with fitness landscape in single-objective optimisation.

- Like fitness landscape, optimal solutions (i.e., Pareto optimal solutions in the context of multi-objective optimisation) are always located in the highest peak(s), such as solutions **b, g, h** in Figure 2.




Fig. 2. A toy example of a set of 10 solutions (**a–j**) with one decision variable (x) and two minimisation objectives (f_1, f_2), shown by Pareto landscape. They, in the format of $(x|f_1, f_2)$, are **a** = (0.1|1, 4), **b** = (0.2|1, 3), **c** = (0.3|3, 3), **d** = (0.4|3, 4), **e** = (0.5|2, 3), **f** = (0.6|6, 3), **g** = (0.7|2, 2), **h** = (0.8|3, 1), **i** = (0.9|3, 3), **j** = (1.0|4, 1).

- Like fitness landscape, it is straightforward to see the characteristics of the problem through Pareto landscape, e.g., local optimality and connectedness of Pareto optimal solutions. For example, in Figure 2 solution **e** is a local optimal solution; optimal solution **b** is disconnected from the other two optimal solutions **g** and **h**.
- Like fitness landscape, in Pareto landscape, for any two solutions, if one is superior to the other (i.e., Pareto dominating), then the better one is always in a higher location, such as solutions **b** versus **a**, and solutions **g** versus **e** in Figure 2.
- Like fitness landscape, in Pareto landscape if two solutions have identical objective vectors, they have the same altitude, such as solutions **c** and **i** in the figure.
- Like fitness landscape, in Pareto landscape the slope indicates the degree of change of a solution to its neighbours (with respect to non-domination levels). For example, in Figure 2 the slope from solution **b** to solution **c** is steeper than that from solution **c** to solution **d** since the level difference between **b** and **c** is greater than that between **c** and **d** (i.e., $2 > 1$).

The above are properties of Pareto landscape that are analogous to fitness landscape in single-objective optimisation. Nevertheless, Pareto landscape does have its own characteristics.

- Unlike fitness landscape, higher altitude in Pareto landscape always means being superior, irrespective of a minimisation or maximisation problem.
- Since Pareto landscape is based on the Pareto dominance relation rather than objective values, solutions with significantly different objective disparities may exist within the same non-domination level provided that they are nondominated to each other. For example, in Figure 2 for the same-level solutions **d** and **f**, on the first objective, **d** is significantly better than **f**, whereas on the second objective, **f** is only slightly better than **d**.
- Solutions that are nondominated to each other *may or may not* be in the same level; it depends on other solutions in the set. For example, in Figure 2




Fig. 3. Pareto landscape of the four-objective MP-DMP problem with different sizes of the considered solution set in the space, where the solutions are uniformly chosen from 500×500 , 200×200 , 100×100 and 50×50 grids, respectively.

solutions **e** and **j** are in the same level, whereas solutions **i** and **j** are not in the same level as solution **e** (who is in the same level as **j**) dominates **i**, thus “pushing” **i** down to the next level.

The reason the proposed Pareto landscape exhibits the last two characteristics is its reliance on the Pareto dominance relation between solutions rather than their native objective values. Pareto landscape is a “relative” landscape, and unlike fitness landscape, the levels of solutions can change depending on the considered solution set. That being said, the change in levels is unlikely to affect the spatial location (i.e., height) of solutions in the landscape, provided that a set of well-distributed solutions is sampled in the space. Figure 3 gives the Pareto landscape of the four-objective MP-DMP problem with different sizes of the considered set. As can be seen in the figure, despite having different scale of non-domination levels, the four landscapes look identical.

In the following section, we will look at more examples of the MP-DMP problem to understand how Pareto landscape reflects other characteristics of multi-objective problems, such as multi-modality, neutrality, and scalability.

3.2 On Different MP-DMP Problem Instances

The MP-DMP problem has been frequently used to examine the behaviours of evolutionary algorithms in the area [29,20,14,23], due to its flexibility of equipping with diverse features [9], such as (global) multi-modality [15] and neutrality [8]. Multi-modality means multiple Pareto optimal solutions in the decision space mapping to the same point in the objective space, and neutrality means solutions in a particular region having the same objective values. The multi-modal MP-DMP is constructed by multiple polygons having the same shape and size, where an objective is the minimisation of distances of a solution to a set of counterpart vertices in the multiple polygons. Figure 4 gives a four-objective multi-modal MP-DMP instance, where the four congruent squares are Pareto optimal regions, each mapping to the whole Pareto front. As shown in the figure, Pareto landscape can well reflect this multi-modality feature, with the four Pareto optimal regions sitting on the four squares.

The neutrality feature of the MP-DMP problem can be easily generated by assigning a constant value to all solutions in a region [8]. Figure 5 presents




Fig. 4. An example of the four-objective multi-modal MP-DMP with four Pareto optimal regions (shaded).




Fig. 5. An example of the four-objective MP-DMP with a Pareto optimal region (shaded) and a neutral region (dashed) of $x \in [0.7, 0.9]^2$, in which $f_1(x) = f_2(x) = f_3(x) = f_4(x) = 2$.

an MP-DMP instance with a neutral region of $x \in [0.7, 0.9]^2$, in which $f_1(x) = f_2(x) = f_3(x) = f_4(x) = 2$. Clearly, Pareto landscape can reflect this feature. All solutions in the neutral region have the same worst level as they are dominated by any solution in the rest of the decision space.

Another important feature of MP-DMP is its scalability in terms of the number of objectives, which is determined by the number of the vertices of the polygon. Figure 6 gives the Pareto landscape of three MP-DMP instances with three, four and six objectives. As can be seen in the figure, although different MP-DMP problems have different scales of non-domination levels, their landscapes look very similar, indicating that the level of difficulty they pose for search algorithms may be very similar.

4 Comparison with Existing Methods

In this section, we compare Pareto landscape with three landscape visualisation methods, gradient field heatmap (GFH) [17], PLOT [26] and cost landscape [11]. The first two are recently proposed state-of-the-arts using a search algorithm to locate global/local efficient sets. Their plots were obtained directly from the moPLOT platform [27] developed by their authors. The third one is similar to ours – the only difference is that cost landscape considers Pareto rank (i.e., the




Fig. 6. Pareto landscape of the MP-DMP instances with different numbers of objectives.

number of solutions that dominate a solution), whereas our method considers the non-domination level a solution is located in. Here, due to their similarity, we use the same colour to represent their “fitness”. Note that since existing methods are usually presented in a 2D space, for facilitating comparison, we present Pareto landscape in a 2D space as well (i.e., the height being removed).

A landscape visualisation method is deemed helpful if it can accurately reflect some important features in multi-objective optimisation. For example, we may hope that it can accurately identify Pareto optimal solutions of the problem. We may also hope that it is aligned with the Pareto dominance relation. That is, if two solutions (especially when they sit closely) are comparable in terms of dominance relation, it is desirable they are visibly distinguishable (i.e., having distinct colours/height). Equally, if two solutions are nondominated to each other, it is desirable that they have the same or similar colour/height.

4.1 Comparison with Search-based Visualisation Methods

As search-based methods, GFH [17] and PLOT [26] first sample a large number of points uniformly in a grid (e.g., a grid of 500×500 cells), and then for each cell, it considers its adjacent cells according to the steepest multiobjective gradient (MOG) direction (see Section 2.1) until the MOG of the cell’s point is zero. However, in practice, in many cases locating a point whose MOG is precisely zero may not be possible. As such, one needs to set a small number as a threshold δ . In GFH [17] and PLOT [26], this parameter was set to $\delta = 1.0e - 4$. On the other side, there may be non-optimal points with a very small MOG that is less than the threshold. Therefore, GFH and PLOT may not be able to accurately identify the Pareto optimal set.

Figure 7 shows the result of GFH and PLOT, along with the two non-search-based visualisation methods cost landscape and Pareto landscape, on the 2D bbob-biobj test function 11 [2], denoted by BBOB11. The Pareto optimal set of BBOB11 is a rectangle in the middle of the decision space. As can be seen in the figure, GFH and PLOT fail to identify the Pareto optimal set of the problem; the blue colour, which is supposed to represent Pareto optimal solutions, misses the target. In contrast, cost landscape and Pareto landscape can accurately reflect the Pareto optimal set (represented by the green points).




Fig. 7. GFH, PLOT, cost landscape and Pareto landscape of function 11 in bbob-biobj [2].




Fig. 8. GFH, PLOT, cost landscape and Pareto landscape of ZDT1 [32].

In addition, the way that search-based methods perform the search also plays a role in their accuracy. In GFH and PLOT, the search is conducted on the pixel of the grid. For each step, the move can only be taken in 8 possible directions: up, down, left, right, and the 4 diagonals, which may not be aligned with the actual steepest direction of MOG. Moreover, the fitness of a point is calculated based on the length of the path to a local efficient set. Points which have the same quality may have different path lengths. For example, Figure 8 shows the result of GFH, PLOT, cost landscape and Pareto landscape on the well-known problem ZDT1 [32]. Looking closely at Figure 8(a) and (b), one may see that there is a diagonal line with inferior fitness than points on either side. This does not mean the points in the diagonal line are worse than others in quality, but means they have longer length than others. As such, the search-based methods, which well reflect the gradient direction, may not be very accurate to reflect the quality of solutions.

In Figure 8, one may also find differences between cost landscape and the proposed Pareto landscape, which is not like the example in Figure 7 where the plots obtained by the two methods are virtually the same. For example, in cost landscape (Figure 8(c)), the worst area is the top-left, whereas in Pareto landscape (Figure 8(d)), the worst is the top. In the next section, we will discuss differences between cost landscape and Pareto landscape.

4.2 Comparison with Cost Landscape

Cost landscape [11] measures the fitness of a point based on the number of points in the whole set that dominate the point. Like Pareto landscape, it provides “relative” fitness of the set considered. However, cost landscape is often much more sensitive than Pareto landscape to points in the set considered. Adding




Fig. 9. 3D cost landscape and Pareto landscape of ZDT1, where the coordinates $(X1, X2)$ of points **A–D** are **A**(0, 0.2), **B**(0, 0.4), **C**(0.2, 0) and **D**(0.4, 0).




Fig. 10. 3D cost landscape and Pareto landscape of DTLZ2, where the coordinates $(X1, X2)$ of points **A–D** are **A**(0, 0.6), **B**(0, 0.8), **C**(0.2, 1) and **D**(0.4, 1).

every new point that dominates the concerned point will change its fitness in cost landscape. This, however, is not the case for Pareto landscape since it considers the level of a Pareto optimal front. For example, in the example of Figure 2, adding any nondominated points between $\mathbf{g}=(2,2)$ and $\mathbf{h}=(3,1)$ will not affect the level of any point in the set. However, for cost landscape, adding such a point will affect the fitness of points **d**, **c**, **i** and **f**. This characteristic of Pareto landscape can make a point's fitness more steady and robust to other points.

Figure 9 gives such an example, where the 3D plots of cost landscape and Pareto landscape on ZDT1 are displayed. For ZDT1, when it has only two variables, $X1$ and $X2$, one can easily derive two properties. That is, 1) for a solution, with $X1$ being fixed, increasing its value on $X2$ will always generate a new solution that is dominated by it; and 2) with $X2$ being fixed, changing its value on $X1$ will always generate a new solution that is nondominated to it (the proofs can be found in <https://github.com/zxc990/Pareto-landscape>). However, cost landscape cannot reflect that. As can be seen from Figure 9(a), solutions **A** and **B** have the same value on $X1$ (hence **A** dominating solution **B**); however, they are virtually on the same height. On the other hand, solutions **C** and **D** have the same value on $X2$ (hence being non-dominated to each other); however, solution **C** is significantly higher than **D** by cost landscape. The same misleading result happens for the problem DTLZ2 [6], shown in Figure 10(a). From the figure, it




Fig. 11. Pareto landscape of ZDT3, ZDT4, DTLZ1 and DTLZ7.

is clear that, by cost landscape, solution **A** dominates **B** but they have the same height; solutions **C** and **D** are nondominated to each other but their heights are highly different. In contrast, Pareto landscape is in line with the problem's properties (Figure 9(b) and Figure 10(b)) – solution **A** is always higher than **B**, and solutions **C** and **D** are always on the same level.

5 Pareto Landscape on Commonly Used Problems

In this section, we use Pareto landscape to plot several widely used problem suites in the domain. They are ZDT [32], DTLZ [6], WFG [13] and BBOB (i.e., bbo-biobj) [2]. For each BBOB problem, there are multiple instances; here the first instance was considered. Due to space limitations, we do not show all test problems for a suite but several representatives. For visualisation, all the problems were set to be in a 2D decision space and with two objectives.

5.1 ZDT and DTLZ

Figure 11(a) and (b) give the Pareto landscape of two ZDT problems, ZDT3 and ZDT4. A characteristic of the ZDT suite is that the Pareto optimal solutions are located in a region (or regions) where all decision variables except x_1 are zero [32], as can be seen from the Pareto landscape in Figure 11(a) and (b). As shown, ZDT3 has five disconnected Pareto optimal regions; note that the problem is not multi-modal, these disconnected regions corresponding to different parts of the Pareto front in the objective space. This feature poses somewhat of a challenge for search algorithms to locate all the optimal regions, but the challenge is mild since the overall landscape is smooth and there are no local optimal regions that can trap search algorithms. ZDT4 is a challenging problem in the suite. From Figure 11(b), we can see that it has many local optima, which can certainly trap the algorithms if their search step is not appropriate (e.g., too small).

The DTLZ suite [6] generally has similar problem features and behaviours to the ZDT suite. Figure 11(c) and (d) give the Pareto landscape of two DTLZ problems, DTLZ1 and DTLZ7. As can be seen from Figure 11(c), DTLZ1, similar to ZDT4, has a number of equally-spaced local optimal regions (despite with deeper valleys). DTLZ7 is similar to ZDT3 and has disconnected Pareto optimal regions when $x_2 = 0$ (see Figure 11(d)), but the gap between the optimal regions is bigger, hence more challenging for search algorithms to jump to the other when they find one of the local optimal regions.




Fig. 12. Pareto landscape of several WFG problems.

5.2 WFG

Figure 12 gives the Pareto landscape of four WFG problems, WFG1, WFG4, WFG7 and WFG8. As can be seen from the figure, a difference of WFG1 from the previous problems is that it has a neural area (around $x_2 = 0.2$), hence posing challenges for search algorithms, particularly for local search ones, to make progress. For WFG4, there are many local optimal peaks along both x_1 and x_2 directions, which contrasts those only along x_2 direction on ZDT4 and DTLZ1. As shown in Figures 12(c) and (d), WFG7 and WFG8 have two piecewise Pareto optimal lines since both contain a transition function where one of the decision variables is mapped onto two regions [13]. It is worth mentioning that the variables of WFG7 on the two Pareto optimal lines can be optimised separately, whereas the variables of WFG8 on one Pareto optimal line are non-separable.

5.3 BBOB

A common feature amongst the previously considered problem suites ZDT, DTLZ and WFG is that they all inherently have some “man-made patterns” (so that the Pareto optimal region is known and controllable). This includes Pareto optimal solutions clustered at regions where certain variables are constants, and local optimal regions that are equally spaced. This may make them easily exploitable by search algorithms; for example, by separately searching for optimum on different variables or employing specific search step-sizes during the search process. This is not the case in the BBOB suite (i.e., bbo-biobj), where the two objectives are made up of well-known single-objective functions (so that the Pareto optimal region is not known nor controllable).

Figure 13 shows the Pareto landscape of eight bbo-biobj problems. As can be seen from the figure, their landscapes have a great variety. There are problems of uni-modality (BBOB1 and BBOB29), high ruggedness (BBOB10, BBOB47 and BBOB55), global structure (BBOBF47), big mountain-like local optima (BBOB10 and BBOB55), and small spike-like local optima (BBOB18, BBOB39 and BBOB53). Moreover, the Pareto optimal regions of the problems are rather irregular. They include the optimal region as a simple straight line (BBOB1), a broken line (BBOB29), several separate lines (BBOB10 and BBOB55), disconnected points (BBOB47), and a mix of lines and points (BBOB18, BBOB39 and BBOB53).

Note that the disconnectedness of Pareto optimal regions in many BBOB problems represents an important feature that challenges search algorithms. To




Fig. 13. Pareto landscape of several BBOB (namely bbob-biojb) problems.

find the whole Pareto optimal region, search algorithms need to maintain diversity during the search – they cannot first search along one direction and, upon reaching an optimal solution, then move around within the neighbourhood. Unfortunately, in contrast to BBOB, most ZDT, DTLZ and WFG problems do not have this feature, making search algorithms easier to find the whole Pareto optimal region.

6 Conclusions

In this work, we attempted to develop a tool to display the “fitness” of solutions in multi-objective optimisation, termed Pareto landscape. We presented the similarity and dissimilarity of the Pareto landscape compared with fitness landscape in single-objective optimisation. Pareto landscape is different from conventional fitness landscape in the sense that it provides “relative” fitness of solutions in the set considered, but like fitness landscape, it can effectively capture a range of problem characteristics (e.g., optimality, multi-modality, 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 also compared Pareto landscape with other popular landscape visualisation methods and presented that it can reflect important features that other methods fail to. Note that like other methods which directly display the decision space, the proposed method can only work on problems with less than three variables.

Lastly, We employed Pareto landscape to visualise commonly used multi-objective benchmark problem suites in the domain and discussed their features and characteristics. We found that the ZDT, DTLZ and WFG problems exhibit some man-made features which can be easily exploited by specific search configurations, whereas the BBOB problems may better represent challenging real-world optimisation scenarios with diverse, natural features.

References

1. Brockhoff, D., Auger, A., Hansen, N., Tušar, T.: BBOB biobj Visualizations (2021), <https://numbbo.github.io/bbob-biobj/vis>
2. Brockhoff, D., Auger, A., Hansen, N., Tušar, T.: Using well-understood single-objective functions in multiobjective black-box optimization test suites. *Evolutionary computation* **30**(2), 165–193 (2022)
3. Brockhoff, D., Tušar, T.: Gecco 2023 tutorial on benchmarking multiobjective optimizers 2.0. In: *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*. pp. 1183–1212 (2023)
4. Custódio, A.L., Madeira, J.F.A.: MultiGLODS: global and local multiobjective optimization using direct search. *Journal of Global Optimization* **72**(2), 323–345 (2018)
5. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation* **6**(2), 182–197 (2002)
6. Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable test problems for evolutionary multiobjective optimization. In: Abraham, A., Jain, L., Goldberg, R. (eds.) *Evolutionary Multiobjective Optimization. Theoretical Advances and Applications*, pp. 105–145. Berlin, Germany: Springer (2005)
7. 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*. pp. 1421–1429 (2019)
8. 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*. pp. 541–549 (2019)
9. Fieldsend, J.E., Chugh, T., Allmendinger, R., Miettinen, K.: A visualizable test problem generator for many-objective optimization. *IEEE Transactions on Evolutionary Computation* **26**(1), 1–11 (2021)
10. Fliege, J., Svaiter, B.F.: Steepest descent methods for multicriteria optimization. *Mathematical Methods of Operations Research* **51**(3), 479–494 (Aug 2000)
11. Fonseca, C.M.M.d.: Multiobjective genetic algorithms with application to control engineering problems. Ph.D. thesis, University of Sheffield (1995)
12. Goldberg, D.: *Genetic algorithms in search, optimization, and machine learning*. Addison-wesley (1989)
13. Huband, S., Hingston, P., Barone, L., While, L.: A review of multiobjective test problems and a scalable test problem toolkit. *IEEE Transactions on Evolutionary Computation* **10**(5), 477–506 (2006)
14. Ishibuchi, H., Akedo, N., Ohyanagi, H., Nojima, Y.: Behavior of EMO algorithms on many-objective optimization problems with correlated objectives. In: *Proc. IEEE Congr. Evol. Comput.* pp. 1465–1472 (2011)
15. 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)*, pp. 91–100 (2010)
16. Jones, T., et al.: *Evolutionary algorithms, fitness landscapes and search*. Ph.D. thesis, Citeseer (1995)
17. Kerschke, P., Grimme, C.: An expedition to multimodal multi-objective optimization landscapes. In: *the 9th International Conference Evolutionary Multi-Criterion Optimization (EMO)*. pp. 329–343. Springer (2017)

18. Kerschke, P., Wang, H., Preuss, M., Grimme, C., Deutz, A., Trautmann, H., Emerich, M.: Towards analyzing multimodality of continuous multiobjective landscapes. In: *Parallel Problem Solving from Nature–PPSN XIV: 14th International Conference*, Edinburgh, UK, September 17–21, 2016, *Proceedings 14*. pp. 962–972. Springer (2016)
19. Köppen, M., Yoshida, K.: Substitute distance assignments in NSGA-II for handling many-objective optimization problems. In: *Evolutionary Multi-Criterion Optimization*, pp. 727–741. Springer Berlin / Heidelberg (2007)
20. 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)*. pp. 261–275 (2013)
21. Liefoghe, A., Derbel, B., Verel, S., López-Ibáñez, M., Aguirre, H., Tanaka, K.: On pareto local optimal solutions networks. In: *International Conference on Parallel Problem Solving from Nature*. pp. 232–244. Springer (2018)
22. 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*. pp. 713–721 (2023)
23. Liu, Y., Ishibuchi, H., Nojima, Y., Masuyama, N., Shang, K.: A double-niched evolutionary algorithm and its behavior on polygon-based problems. In: *the 15th International Conference on Parallel Problem Solving from Nature (PPSN)*. pp. 262–273. Springer (2018)
24. Ochoa, G., Liefoghe, A., Lavinas, Y., Aranha, C.: Decision/objective space trajectory networks for multi-objective combinatorial optimisation. In: *European Conference on Evolutionary Computation in Combinatorial Optimization (Part of EvoStar)*. pp. 211–226. Springer (2023)
25. Ochoa, G., Malan, K.: Recent advances in fitness landscape analysis. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. pp. 1077–1094 (2019)
26. Schäpermeier, L., Grimme, C., Kerschke, P.: One plot to show them all: visualization of efficient sets in multi-objective landscapes. In: *International Conference on Parallel Problem Solving from Nature*. pp. 154–167. Springer (2020)
27. Schäpermeier, L., Grimme, C., Kerschke, P.: moPLOT website (2021), <https://schaepermeier.shinyapps.io/moPLOT/>
28. Schäpermeier, L., Grimme, C., Kerschke, P.: Plotting impossible? surveying visualization methods for continuous multi-objective benchmark problems. *IEEE Transactions on Evolutionary Computation* **26**(6), 1306–1320 (2022)
29. 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*, pp. 401–410. Springer Berlin / Heidelberg (2008)
30. 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*. pp. 647–655 (2019)
31. Wright, S., et al.: The roles of mutation, inbreeding, crossbreeding, and selection in evolution pp. 356–366 (1932)
32. Zitzler, E., Deb, K., Thiele, L.: Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation* **8**(2), 173–195 (Jun 2000)