Plotting Impossible? Surveying Visualization Methods for Continuous Multi-Objective Benchmark Problems

Lennart Schäpermeier, Christian Grimme, and Pascal Kerschke *Member, IEEE,*

***Abstract*—Traditionally, visualizing benchmark problems is an integral task in the domain of evolutionary algorithms development. Researchers get inspired for new search heuristics by challenges observed in functional landscapes. Moreover, landscape characteristics, features, and even terminology to describe them are derived from visualizations. And most importantly, benchmark designers need visualizations for identifying diverse problems that potentially challenge different aspects of optimization algorithms. As easy as it is to visualize single-objective problems, until recently there were hardly any approaches for gaining similar insights for multi-objective problems. Also, there have been no seamlessly accessible tools to support such visualizations.**

This paper presents a comprehensive overview of the available visualization techniques from literature, including two interactive techniques to visualize three-dimensional problems, as well as two novel techniques which are suitable to scale some visualization properties to even higher-dimensional spaces. All presented techniques are integrated into a single tool, the moPLOT-dashboard, which enables users to perform landscape analyses in an interactive manner. Finally, the value of the tool and the visualizations is demonstrated in a series of usage scenarios on well-known benchmark problems.

***Index Terms*—Multi-Objective Optimization, Visualization, Multimodal Optimization, Benchmarks, Theory, Algorithms.**

I. Introduction

The goal of developing benchmarks for optimization algorithms is centrally focused on providing a portfolio of challenging objective functions that is as heterogeneous as possible. By means of the different challenges induced by these problems, weak points of algorithms shall be identified on the one hand and performance comparisons between algorithms in different situations shall be enabled on the other hand.

Traditionally, in continuous optimization, benchmarks are composed of problems that exhibit specific characteristics or structural properties. These characteristics are usually based on a fundamental notion of so-called functional or fitness landscapes [1]. In low-dimensional search spaces, we tend to even name them as mountains, valleys, plateaus, ridges, or by other names from the familiar three-dimensional world to give the topological structures an imaginable visual equivalent [2]. It is not surprising, then, that visualization is an important aspect of benchmark development – just as it is central to the development of algorithms.

Lennart Schäpermeier and Pascal Kerschke are with the TU Dresden, Germany and Christian Grimme is with the University of Münster, Germany.

Manuscript received xxx xx, xxx; revised xxx xx, xxx.

While we are familiar with and intuitively use the descriptive representations of function landscapes from optimization textbooks in the case of a single-objective function, for a long time there existed no analogous notion for multi-objective optimization. This is mostly due to the simple fact that the objective space for single-objective optimization is one-dimensional. As such, the objective value can be interpreted as "height" in the function landscape. However, each solution for a multi-objective problem (MOP) $f : \mathbb{R}^p \rightarrow \mathbb{R}^k$ needs to be optimal with respect to a k-dimensional objective space vector [3], [4]. Thus, a direct and intuitive visualization and interpretation like in the single-objective case is difficult – even in low-dimensional search spaces.

Studying a MOP's Pareto set – i.e., its global optimum as will be defined in Section II – has nonetheless been the focus of visual investigations for years. Yet, due to a lack of suitable methods, the main technique for visualizing solutions and discussing properties of MOPs has been a simple scatter plot of approximated Pareto-optimal solutions (i.e., the Pareto set) and their image in objective space (called Pareto front).

Obviously, this approach corresponds to plotting a single or few (evaluated) optimal solution(s) in case of a single-objective problem (SOP) – an approach that is usually not used for SOPs because it provides too little information. In fact, it puts objective space representation into focus and hence does not tell much about the problem's structural characteristics in the decision space. As such, it does not help at all in identifying challenging properties for benchmark development, as well as for algorithm design. Moreover, this decades-long lack of visualization methods could possibly also be the cause for the absence of a common notion of *multimodality* and *localness* in the multi-objective optimization community [5].

For benchmark development, the lack of visualization techniques for MOO function landscapes resulted in a very limited idea of the structures in decision space. Not only early descriptions of benchmark problems [6], [7] but also nowadays prevalent benchmark suites like ZDT [8], DTLZ [9], WFG [10], ED [11], or CEC benchmarks [12], [13] mainly rely on presenting scatter plots of Pareto front and Pareto set approximations. The rest of the search space is considered (at most) by randomly placed points in the corresponding scatter plots.

Although techniques for visualizing solutions for multi-objective problems have been proposed over the years [14], [15], they have not been considered extensively for benchmark visualization. Only few benchmark developers try to provide additional information on the search space [16], using, for instance, line cuts through decision space.[1] Certainly, one reason for this observation is that until recently, almost no alternative visualization technique was available. Apart from some early, though little used global techniques [17], it is only since 2016 that new approaches for visualizing multi-objective landscapes have been proposed [18]–[23]. However, until now, a survey that focuses on these recently developed visualization methods, as well as an accessible toolbox that enables the visualization and analysis of both existing as well as new benchmark problems have been missing.

Within this work, we present a web-based (hence platform-independent) and interactive tool[2] originally introduced in [24] for the visualization of benchmarks based on the `R`-package `moPLOT` [21]. It allows to take a deeper look into decision space structures by integrating multiple graphical representations of search space and objective space into a single, accessible application. Besides two-dimensional representations, we also provide two visual approaches for three-dimensional decision spaces. These allow to identify structures beyond the globally efficient set (Pareto set). Local structures and basins of attraction surrounding locally efficient sets can be visualized for standard benchmarks and individually designed problems. A new graph-based representation of basins of attraction (and their interactions) is introduced, as well as an adaptation of the well known Parallel Coordinates Plot visualization for multimodal problems. These approaches could potentially address even higher-dimensional search spaces, if appropriate methods, e.g., for computing the necessary data for all efficient sets or local fronts, would be available.

The remainder of this work is structured as follows: Section II introduces some mathematical fundamentals and Section III gives an overview of the available visualization techniques for continuous MOPs. Then, Section IV describes the architecture and features of the `moPLOT` dashboard. This is followed by several usage scenarios demonstrating the landscape analyses enabled by state-of-the-art visualizations in Section V. Finally, Section VI presents the conclusion.

II. Background

In this section, we will first provide mathematical notions and concepts which are useful for understanding and discussing the differences of the visualization techniques shown within this work. This is followed by a more detailed discussion of the notion of multi-objective gradients and basins of attraction.

Throughout this work, we assume box-constrained continuous MOPs $f : \mathcal{X} \to \mathbb{R}^k$ that w.l.o.g. shall be minimized. Moreover, the considered search space $\mathcal{X}$ is assumed to be bounded by lower and upper bounds $\mathbf{l}, \mathbf{u} \in \mathbb{R}^p$, respectively. Note that the vast majority of the visualizations used within our dashboard (see Sections III and IV for further details) are able to depict search and objective spaces that are two- or three-dimensional, i.e., $k, p \in \{2, 3\}$.

[1] http://numbbo.github.io/coco-doc/bbob-biobj/functions/

[2] https://schaepermeier.shinyapps.io/moPLOT

Potential regions of interest within a MOP's landscape are, among others, locally and/or globally efficient sets (i.e., the multi-objective counterparts of local and global optima). The *globally efficient set* $\mathcal{X}^* \subseteq \mathcal{X}$, better known as *Pareto set*, refers to the set of globally non-dominated solutions. Further, a point $\mathbf{x}^* \in \mathcal{X}$ is said to be a *non-dominated solution*, if there exists no other solution $\mathbf{x} \in \mathcal{X}$ with $f_i(\mathbf{x}) \leq f_i(\mathbf{x}^*)$, $i = 1, \ldots, k$, and $f_i(\mathbf{x}) < f_i(\mathbf{x}^*)$ for at least one i. The image $f(\mathcal{X}^*)$ of the respective Pareto set is called the *Pareto front*. In analogy to the single-objective case, local optimality in the multi-objective case is also defined by means of the dominance relation in the local neighborhood. Hence, a solution $\mathbf{x}$ is said to be a *locally efficient point*, if it is not dominated by any other solution $\mathbf{y} \in B_\varepsilon(\mathbf{x})$ within its surrounding ε-ball $B_\varepsilon(\mathbf{x}) \subseteq \mathcal{X}$ (for a sufficiently small $\varepsilon > 0$). A connected set of locally efficient points then forms a *locally efficient set* $\mathcal{X}_{LE}$, i.e., the multi-objective version of a local optimum [5], [25].

An important landscape property that affects both algorithm engineers and benchmark designers are the so-called *attraction basins*. Prior to formally defining these basins of attraction, we mention on a qualitative level that they constitute regions within the search space, where a gradient-based local search strategy will – independent of its starting position within that basin – always converge to (and potentially get trapped in) the same locally efficient set. Thus, we need to introduce the notion of a *multi-objective gradient* (MOG), which is a vector that is oriented towards a common descent direction of the MOP's single-objective components. If the MOG is computed in a locally efficient point, no common descent direction exists and hence the MOG is of length zero. As elaborated in [26], the multi-objective gradient can be defined as shown in Equation (1):

$$\nabla f(\mathbf{x}) = \sum_{i=1}^{k} \alpha_i^* \nabla f_i(\mathbf{x}). \tag{1}$$

Therein, α^* is a weight vector chosen according to Equation (2), such that the resulting MOG is the shortest vector in the convex hull of the single-objective gradients.

$$\alpha^* = \arg\min_{\alpha} \left\{ \left\| \sum_{i=1}^{k} \alpha_i \nabla f_i(\mathbf{x}) \right\| \;\middle|\; \alpha_i \geq 0,\ \sum_{i=1}^{k} \alpha_i = 1 \right\} \tag{2}$$

A potential bias induced by the individual lengths of the gradients can be compensated by normalizing them before the MOG computation [27]. We use the MOG definition with normalized single-objective gradients for the remainder of this work. In the bi-objective case, this simplifies to the arithmetic mean of the normalized single-objective gradients [24].

Related to the notion of MOGs (as one descent strategy) and according to the general notion introduced by Törn and Zilinskas [28] for any local descent procedure, we define the *region of attraction* of a locally efficient point $x_L^* \in \mathcal{X}_{LE}$ of a MOP as follows: $\text{attr}(x_L^*) \subseteq \mathcal{X}$ is the largest set of points such that for any starting point a MO steepest descent procedure with infinitely small steps will converge on x_L^*. Accordingly, the basin of attraction of a (locally) efficient set $\mathcal{X}_{LE}$ is defined as the union of non-overlapping connected regions of attraction of the elements of $\mathcal{X}_{LE}$. Clearly, these basins of attraction are of major interest for the notion of a functional landscape and an interpretation of algorithmic behavior on such landscapes.

Finally we consider the borders between different basins of attraction. These so-called *ridges* are the set of points that separate adjacent basins of attraction. All points belonging to such a ridge share the property of having multiple neighboring solutions that belong to different basins of attraction.

Next, we present a variety of techniques suitable for visualizing MOPs and also discuss their respective strengths and weaknesses using the concepts and terminology defined above.

III. Visualization Approaches

Available visualization approaches for MOPs can be classified into (1) methods for depicting the relation of objective values of multiple dominating or dominated solutions, and (2) methods for highlighting structural properties of the decision space.

Certainly, one of the most intuitive approaches of visually representing solutions of a MOP is a scatter plot of the Pareto front and (if needed) also of dominated points. It aims for directly representing the quality relation of different solutions in objective space for up to three dimensions. Much effort has been put into advanced visualization techniques for enabling visualization of more than three objectives at the same time. Early visualization techniques simply reduce the objective space by measuring and plotting the distance of each solution to an approximated Pareto front or to neighboring solutions, respectively [30]. Other approaches [31], [32] employ more sophisticated dimension reduction techniques like self-organizing maps [33] for denoting similarity of solutions in lower dimensions. Further techniques from this subclass of approaches use Radial coordinate visualization [34] or allow arbitrary dimension reduction approaches (like PCA [35] or MDS [36]) for creating projection views of so-called Skylines that allow a first comparison and clustering of solutions before selecting detailed views in a dashboard-oriented decision support tool [37]. In contrast, other authors propose heatmap-based techniques [38] or other lossless projections [39] to provide decision makers with a notion of solution quality and distribution (in objective space and partly sometimes also in decision space). Finally, we want to mention a comprehensive survey on approaches from multi-criteria decision making [40], which provides a specific discussion on reduction techniques with a focus on the decision-making process.

The availability of visualization techniques for decision space structures – the landscape of MOPs – was, however, very limited for a long time. In fact, for almost two decades, the so-called *cost landscape* proposed by Fonseca [17] (see Section III-A1) has been the only visualization technique that provided a landscape-like illustration of a MOP's fitness function. In the beginning of the 2010s, the PhD thesis of Tušar [14], as well as the subsequent work by Tušar and Filipič [15] brought new life into the discussion of this research topic. Yet, as acknowledged by the authors themselves, the methods presented in their work (scatter plot matrices, sammon mapping, principal components, etc.) were insufficient for investigating landscape structures of MOPs. They thus proposed a new visualization technique, called prosection method, which iteratively reduced a MOP's dimensionality by means of subsequent hyperplane cuts. However, these cuts produce a severe information loss and hence make this technique impracticable for visual inspections of MOP landscapes. In recent years, several visualization techniques have been developed, which (a) are capable of illustrating MOP landscapes in search and/or objective spaces while also accounting for interaction effects, thereby (b) reveal challenging, as well as (to a certain degree) efficiently exploitable structural properties in these landscapes, and (c) help to study search dynamics of (evolutionary) optimization algorithms.

In the following, we will summarize the general idea of all the relevant visualization approaches from the literature. To this end, Section III-A will give an overview of several visualization techniques that combine the information of a MOP's search and objective space within a single plot. Subsequently, in Section III-B, we will complement the visualization techniques for two- and three-dimensional problems in decision and/or objective space introduced before by methods which even support the illustration of high-dimensional MOPs and many-objective problems. For illustrative purposes, as well as for a simplified comparison of the methods presented below, we depict their respective visualizations in Figures 1, 2 and 3, respectively. The strengths and weaknesses of all presented visualization methods are discussed in Section III-C, and their properties are summarized in Table I. At last, Section III-D presents two interactive 3D visualizations and elaborates on the opportunities that may arise from integrating static visualization procedures into interactive dashboards.

A. Illustrating MOPs via Heatmaps

All the images shown in this subsection are based on the same exemplary bi-objective multiple peaks problem[3]. For the methods illustrated in Figure 1, the plots are shown in the search space (top row), as well as in the corresponding objective space (bottom).

It should be emphasized that all of the methods discussed below are based exclusively on a grid of evaluated points, making them completely independent of specific function definitions. Due to this high flexibility, any of the methods can easily be applied to novel benchmark problems, true black-box functions, or even problems from real-world applications – as long as the said grid of evaluated points can be provided.

1) Cost Landscape: As stated in the beginning of this section, the cost landscapes by Fonseca [17] (left-most top and bottom images in Figure 1) has been one of the first methods to illustrate a MOP's structural relationships in the search space. For each point in the landscape, the color (which can be interpreted as height) depicts the number of solutions that dominate the respective solution (i.e., its Pareto rank). This method primarily focuses on the global dominance relationship

[3]The two single-objective components were produced with the MPM2-function generator available within the R-package `smoof` [41]. Both objectives were configured with three peaks with random topology, two-dimensional search space and random seeds 4 and 8, respectively.

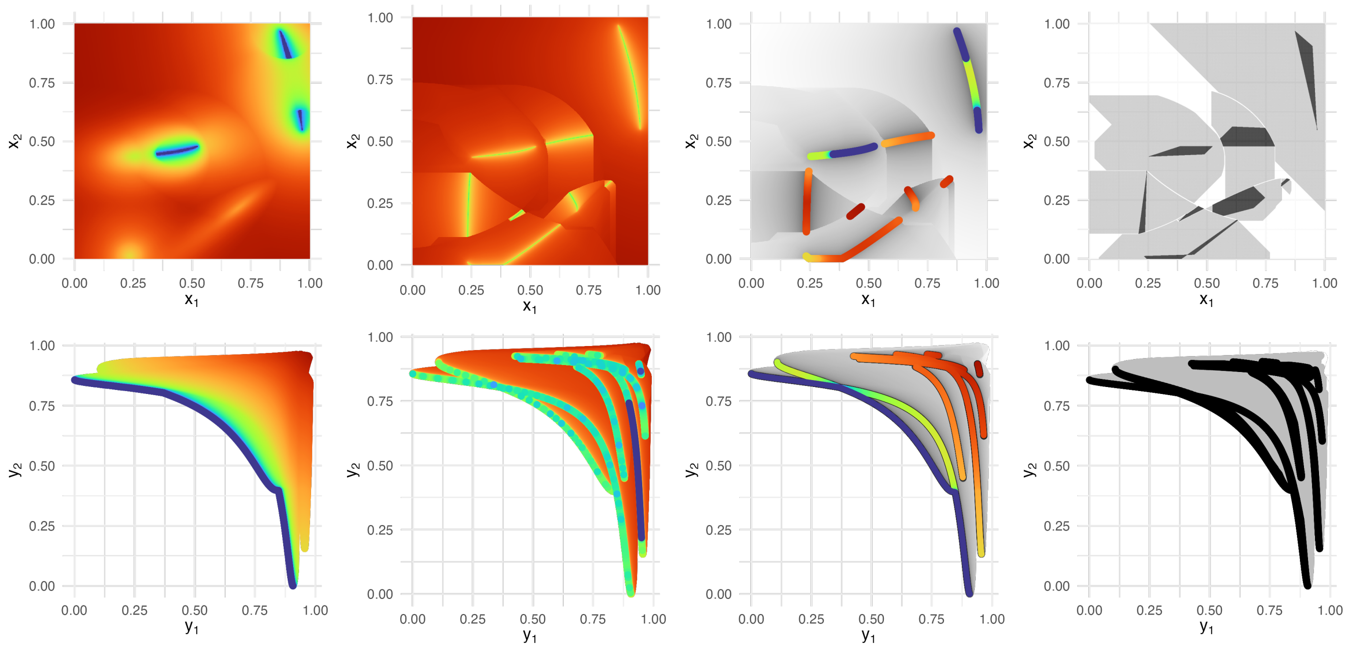

Fig. 1. Overview of decision space visualization techniques with corresponding visualizations in objective space. Visualizations from left to right: Cost landscape [17], gradient field heatmap [19], PLOT [21] and local dominance landscape [29]. The top row shows the decision, and the bottom row the corresponding objective space. All visualizations use a grid with a resolution of 1000 evenly-spaced points per dimension, except for the cost landscape, which is rendered using a resolution of 500 given the significant computational overhead. The gray color scale denotes more (locally) optimal points with darker colors, while the other scale denotes better values (w.r.t. the given visualization) with blue, and worse with red. The color scales are log-scaled.

– as can be seen very well in the continuous color gradient in the objective space – and thus is very suitable for highlighting a MOP's globally efficient sets (i.e., its Pareto sets). However, challenging structures like locally efficient sets (along with their basins of attraction) are hardly visible.

2) Multi-Contour Plot: In order to facilitate investigations of a MOP's local structures (caused by the mutual interactions of the optima of all objectives), Kerschke et al. [18] used contour plots to visually combine the landscapes of the MOP's individual single-objective components into a single image. As shown in Figure 2 (left), each contour plot reveals the positions of its objectives' local optima, which in turn are an indication for locally efficient points in the MOP landscape. By augmenting contour plots with sets of non-dominated points, i.e., Pareto sets, this visualization approach can provide some initial indications of both a MOP's local and global landscape structure. However, contour plots consider MOPs primarily from a single-objective perspective, so interaction effects can at most be anticipated. Moreover, there is no native counterpart for visualizing the respective information in the objective space.

3) Gradient Field Heatmap: An extension of the aforementioned contour plots – although with a much stronger focus on the local relationships – has been proposed by Kerschke and Grimme [19]. Their gradient field heatmaps (GFHs) not only depict the MOP's locally efficient sets, but also reveal the corresponding (surrounding) attraction basins. That is, similar to basins of attraction in single-objective optimization, starting in any point of the search space, a gradient-based local search strategy will descend to one of the locally efficient points enclosed within the respective basin (see second column of

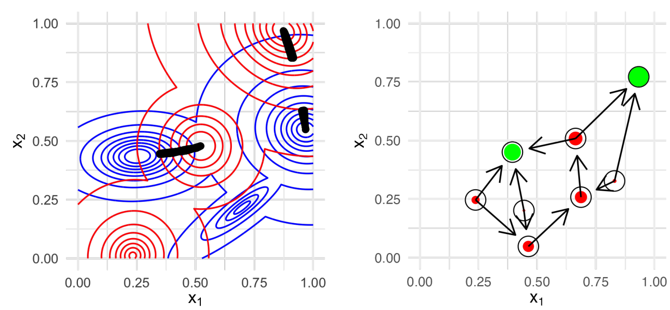

Fig. 2. Overview of additional visualization techniques. In contrast to the techniques used in Figure 1, the methods used for producing the images above are limited to illustrations of the search space. Left: Multi-contour plot (with overlaid Pareto set approximation) [18]. Right: The herein proposed set transition graph.

Figure 1). The gradients are originally evaluated directly on the test function using finite difference approximations, but are approximated here using their respective grid neighbors, instead, as in [21]. The illustrated color (or height) of a point from the GFH corresponds to the cumulated length of a path of normalized multi-objective gradients from the respective point to its attracting locally efficient point. It should be noted that although the depicted efficient sets also contain the Pareto sets, it is very difficult to identify them among the numerous locally and globally efficient sets. Therefore, these plots are less suitable for studying global relationships.

4) PLOT: While the cost landscapes were well suited for visualizing global relationships, the GFHs were particularly helpful in depicting local structures. Schäpermeier et al. [21]

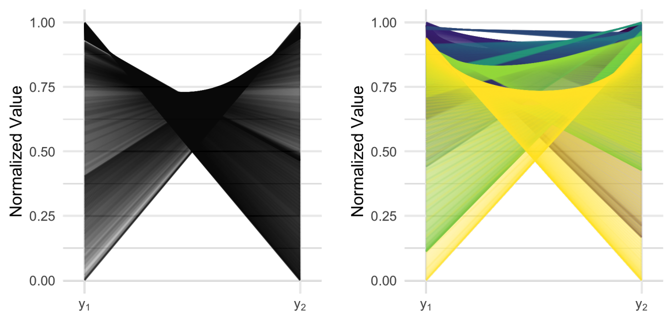

Fig. 3. The two PCP visualization techniques discussed in this work (objectives are individually scaled to the $[0, 1]$ range). Left: The global PCP visualization for nondominated evaluated points only. Right: The local PCP, including all (previously approximated) locally efficient sets, denoted by the different colors.

thus tried to combine the best of both methods within their PLOT (Plot of Landscapes with Optimal Trade-offs) method, which is illustrated in the third column of Figure 1. The gray-colored background shows the attraction basins known from the GFH approach. On top of those basins, once again, all locally and globally efficient sets are depicted (in colors). However, in contrast to GFHs, the coloring now corresponds to the Pareto ranking of the efficient points. It thus enables visual comparison of the efficient sets on a global level. In summary, PLOT is the first visualization technique that enables the mutual investigation of a MOP's global *and* local structures.

5) Local Dominance Landscape: The images on the right of Figure 1 can be regarded as a discretized and hence simpler version of PLOT. In contrast to the previously discussed continuous color scales, the local dominance landscape proposed by Fieldsend et al. [29] use a categorical classification, which roughly translates into locally/globally optimal regions (black), basins of attraction (gray) and undecided regions where multiple basins are reachable by their method (white). Note that this idea is similar to the cell mapping concept from the single-objective domain [42]. Within the respective illustration, a black point indicates a point that is mutually non-dominating with all the points from its Moore neighborhood. Thus, the black regions correspond roughly to the locally efficient sets, and hence also contain the Pareto set approximation w.r.t. the grid of evaluated points. A gray-colored region indicates a basin of attraction, in which each of its points has at least one dominating neighbor and all sequential movements in direction of the dominating neighbors will converge in (one of the points of) the same locally optimal region. At last, each of the white regions is constituted of points, which have at least two dominating neighbors that lead to opposing optimal regions. Therefore, all the white regions basically form borders or ridges between adjacent attraction basins. Noticeably, this rather coarse categorization results in optimal regions that are much broader than for PLOT and co., which can be explained by artifacts that have been introduced by the Moore neighborhood. That is, there are often 45° and 90° artifacts around the locally efficient sets (as found by the other techniques), which incorrectly include a large amount of points around the locally efficient sets.

B. Approaches Capable of Visualizing Large-Scale MOPs

The methods discussed in Section III-A enable the visualization (and subsequent analysis) of MOPs whose search and objective spaces are at most three-dimensional. If there is a need to visually represent and assess higher-dimensional MOPs, these methods reach their limits. In the following, we present two methods that could potentially scale further and thereby could enable representation of high-dimensional and many-objective problems. Here, we focus on the two- and three-dimensional cases so that they can be computed within the same framework as the other visualizations. In particular, the approximation of the locally efficient sets would require a different strategy for higher-dimensional decision spaces.

1) Parallel Coordinates Plots: Apart from scatter plots, parallel coordinates plots (PCPs) have been a popular tool for analyzing MOPs for years [15], [43]. The general idea of PCPs is to arrange the MOP's objectives as parallel coordinates along one axis, then mark the corresponding objective values as 'heights' of the respective coordinates, and finally connect these markers with a line [44]. Consequently, each observation from the MOP's p-dimensional (objective) space is represented as a single line connecting the points of p successive parallel coordinates. It should be noted, though, that PCPs – the way they are commonly used in the MO context – provide very little information about a MOP's structure: usually, they only illustrate the Pareto front of the MOP. Hence, classical PCPs only offer insights into the MOP's global structure – and only based on an objective space perspective. This would be akin to describing a single-objective problem in terms of the objective value of its optimum. Next to this classical approach (which we refer to as *global PCP*), we propose a PCP visualization that is enriched by the approximated locally efficient sets extracted from the PLOT visualizations. The *local PCP* method applies the concept of colored polylines, which depict observations from the same cluster in identical colors [45]. By treating each locally efficient set as a separate set (i.e., cluster) of points, our approach enables visually distinguishing these sets from each other. Consequently, the local PCP method nicely complements its global counterpart, which focuses exclusively on the non-dominated points. Both PCP visualization approaches are demonstrated in Figure 3.

Further, previous works showed that PCPs are easily scalable to higher-dimensional problems [44]. In this context it should be noted that too many objectives (and thus PCP coordinates) could cause visual clutter. Yet, this limitation appears to be negligible for the majority of MOPs.

2) Set Transition Graph: Investigations of the previously discussed methods revealed superpositions of neighboring basins of attraction (see, e.g., gradient field heatmap and PLOT in Figure 1), which enable even rather simple local search strategies to move from one locally efficient set to an adjacent one that is dominating the current efficient set [20], [25]. The superposition relationship between the locally efficient sets lends itself to a natural representation as a graph (see right image of Figure 2), which we propose as a new visualization technique called *set transition graph*. Similar to the concept of Local Optima Networks (LONs) [23], [46] and a related combinatorial MO variant called Pareto Local Optimal Solutions Networks (PLOS-net) [47], each node in the proposed set transition graph represents a single locally efficient set and a directed edge between two sets A and B denotes that set A is superimposed by set B. In an extension of this method, the nodes within the graph are filled with one of two different colors: green nodes indicate global optima, whereas red points correspond to non-global optima. Moreover, the amount of color used for filling a node indicates the proportion of local search runs that traversed (or even converged in) the respective node (i.e., efficient set). That is, the larger the green or red area of a node is, the more likely will the respective efficient set be visited by a local search run. In the two-dimensional decision space given here, the placement of the nodes reflects that of the corresponding locally efficient set in the decision space by placing each node on the (approximate) medoid of the corresponding set. Higher-dimensional set transition graphs can be visualized, e.g., by a stress-based graph layout.

Apart from finding a suitable planar visualization of the sets and transitions, the main challenge for scaling the transition graphs to higher dimensions is an efficient computation of the locally efficient sets and the corresponding set transitions. The approach presented here, based on densely sampling a grid of points in the decision space, suffers from the curse of dimensionality and this data will need to be acquired differently, e.g., with specialized local search strategies such as the Multi-Objective Landscape Explorer (MOLE) [48], to scale the transition graphs beyond the presented three decision space dimensions. Except for this limitation, set transition graphs are well scalable since their number of nodes is independent of the dimensionality of the search or objective space. Instead, the number of nodes in the graph depends only on the MOP's number of locally efficient sets.

C. Exemplary Comparison of the Considered Visualization Techniques

Now that we have introduced the technical details of the state-of-the-art techniques for visualizing MOPs – including the two novel scalable methods that we proposed in Section III-B – we will once again take a look at Figures 1, 2 and 3. This time, however, the focus is on discussing the advantages and disadvantages of all the methods. To facilitate a comparison of the surveyed visualization approaches, their properties are summarized in Table I. Therein, we indicate for each of the examined methods whether it is capable of revealing information about a MOP's (1) globally efficient regions, (2) locally efficient sets, (3) enclosing basins of attraction, the dimensionality of the (4) search and (5) objective space that can be visualized by the respective approach, as well as any properties of the MOP's (6) objective space or its (7) single-objective components.

Looking at the decision space (of the exemplarily examined bi-objective MPM2 function) as illustrated by the gradient field heatmap, PLOT and local dominance landscape (see second to fourth column of Figure 1), and by the set transition graph (right image of Figure 2), one can observe a total of eight disjoint efficient sets, with two of them being also globally efficient and thus part of the Pareto set. Out of the four techniques mentioned, all of which are suitable candidates for visualizing a MOP's local structures, only PLOT and the set transition graph reveal the MOP's globally efficient sets or fronts as well. Yet, if the visualization's primary purpose is to gain insight into globally efficient sets (or fronts), the cost landscape approach provides a promising alternative; its color gradient, which illustrates the MOP's underlying Pareto ranking, facilitates the identification of non-dominated regions.

Interestingly, the efficient set on the top right of the MOP's decision space – as, e.g., shown in the heatmaps of Figure 1 and to a certain extent also in the multi-contour plot of Figure 2 (left image) – is partially dominated, resulting in two globally efficient parts at both ends and an inferior segment in between. Depending on the choice of the visualization method, such a trait could potentially be misleading. For instance, methods that focus exclusively on the MOP's global structures (i.e., the Pareto set approximation depicted within the multi-contour plot, as well as the cost landscapes) would highlight a total of three (disjoint) efficient sets and users would have to guess whether the sets are connected or not.

In line with the previous argumentation, it can be observed that, in general, the gradient field heatmap, local dominance landscape, and PLOT are very useful for revealing local structures like locally efficient sets and their accompanying basins of attraction. However, with the exception of PLOT, those methods have clear weaknesses when it comes to identifying global structures. Due to the fact that the coloring of each attraction basin is assigned independent of the information from the other basins, all the basins look almost alike – and the same holds true for the efficient sets. This makes it extremely difficult to identify promising regions – or even the Pareto set – in images produced by these techniques. For such a task, the cost landscape and in particular PLOT would provide the most valuable insights.

As one might notice, the set transition graph and the parallel coordinates plot take a special role. Due to the graph structure of the former, which naturally introduces a considerable compression of the MOP's landscape, it is much more difficult to envision the structural (both positive and negative) properties of the problem at hand using the set transition graph. On the positive side, the set transition graph shows in a very easy-to-understand way how easy or hard it is to reach the different basins of attraction. Moreover, both approaches (i.e., the set transition graph and PCPs) scale with the MOP's dimensionality and thus can potentially be used to characterize high-dimensional or many-objective problems.

Finally, our modified version of PCPs, i.e., local PCP, is the only visualization technique to date that provides insight into the *local* structures of MOPs with high-dimensional search or objective space. As exemplified in Section V, local PCPs indicate the location of a MOP's local fronts (and sets, if studied) per dimension of the objective (or search) space. However, due to their compact and abstract representation of the examined MOP, it is difficult to imagine the exact spatial location of the associated fronts (or sets) using PCPs. Yet, they enable the identification of the relevant areas within the MOP's objective (or search) space [50].

TABLE I

SUMMARIZED VIEW AT THE PROPERTIES OF THE PRESENTED VISUALIZATION TECHNIQUES. FOR EACH PROPERTY, IT IS INDICATED WHETHER IT IS SATISFIED (✓), PARTIALLY SATISFIED (●) OR NOT SATISFIED (✗) BY THE RESPECTIVE METHOD. NOTE THAT THE SCALING OF METHODS "PARALLEL COORDINATES PLOT" AND "SET TRANSITION GRAPH" TO HIGHER DIMENSIONS (MORE THAN 3) WOULD IN PRINCIPLE BE POSSIBLE, BUT WOULD REQUIRE ADDITIONAL ALGORITHMS THAT COLLECT THE NECESSARY DATA.

Visualization Technique	**Global Optimality**	**Locally Efficient Sets**	**Attraction Basins**	**Dimensionality** $\mathcal{X}$	$\mathcal{Y}$	**Properties of** **Obj. Space**	**SO components**
Cost Landscape [17]	✓	✗	✗	2-3	2-3	✓	✗
Multi-Contour Plot [49]	● If overlaid	● If overlaid	✗	2	≥ 2	✗	✓
Gradient Field Heatmap [19]	✗	✓	✓	2-3	2-3	✓	● Pos. of local optima
PLOT [21]	✓	✓	✓	2-3	2-3	✓	● Pos. of local optima
Local Dominance Landscape [29]	✗	● Approx.	● Approx.	2-3	2-3	● Approx.	✗
Parallel Coordinates Plot	✓	✗ (Global PCP) ✓ (Local PCP)	✗	≥ 2	≥ 2	● Limited Extend	✗
Set Transition Graph	● Per LE set	● Schematically	● Size, Interactions	≥ 2	≥ 2	● Limited Extend	✗
3D Visualizations							
Scatter Plot of Nondominated Points	✓	✗	✗	3	2-3	✓	✗
MRI Scan [24]	✓	✓	✓	3	2-3	✓	● Pos. of local optima
Onion Layers [24]	✓	✓	✓	3	2-3	✓	● Pos. of local optima

D. Interactive 3D Visualization

Another limiting factor for visualizing MOPs has been the restriction to static, planar visualizations. In such cases, plotting MOPs with three dimensions in search or objective space was hardly imaginable as the points along the boundaries of the respective space would cover all inner points. However, by using the interactive functionalities of the dashboard presented below, we enhanced all the techniques illustrated in Figure 1 (i.e., the ones described in Section III-A, except for the multi-contour plot) such that they can also be used for visualizing three-dimensional spaces. To achieve this, we use two approaches – called "MRI Scan" and "Onion Layers" – that were recently introduced in related work [24]. Note that within this work, we extend their application to the local dominance landscapes proposed by Fieldsend et al. [29]. By computing the necessary data, such as neighborhood dominance relationships, in three dimensions and applying one of the following 3D visualization techniques, we extend their technique to three decision space dimensions for the first time.

The first 3D technique uses an MRI-like procedure that enables a scan of the volume of interest (e.g., the PLOT visualization of a MOP with a three-dimensional search space) by interactively slicing the respective 3D cube. An exemplary MRI scan based on PLOT is illustrated for the three-dimensional Aspar function [24] in Figure 4.

As alternative to the MRI scan, we provide an interactive visualization based on 3D isosurfaces. These allow to interactively add and/or remove a MOP's domination layers similar to 'peeling off' the layers of an onion – hence, we denoted it the 'Onion Layer' approach. An exemplary illustration can be seen in Figure 7, which depicts a three-dimensional instance of the SGK problem [21] using this Onion Layer method.

The properties of the two interactive visualization methods discussed herein are also listed in Table I. It should be noted, though, that their ability to show global and local structures essentially depends on the choice of the underlying heatmap method. That is, an MRI scan of a three-dimensional local dominance landscape visualization is still unable to reveal global optimality. Similarly, three-dimensional interactive versions of cost landscapes remain unsuitable for examining locally efficient sets or basins of attraction. However, using a suitable approach for the desired purpose, these interactive heatmap cubes enable investigations of the respective global and local structures for three-dimensional MOPs.

Finally, we should also mention the "classical" 3D visualization, in which (only) the nondominated solutions of the sample are displayed through scatter plots in the search and objective space, respectively. We collected and implemented all of the discussed visualization techniques in our visualization dashboard, which we will discuss next.

IV. THE MOPLOT DASHBOARD

The previous section presented a multitude of visualization techniques for continuous MO landscapes. However, the study of existing benchmark problems or application of the visual-

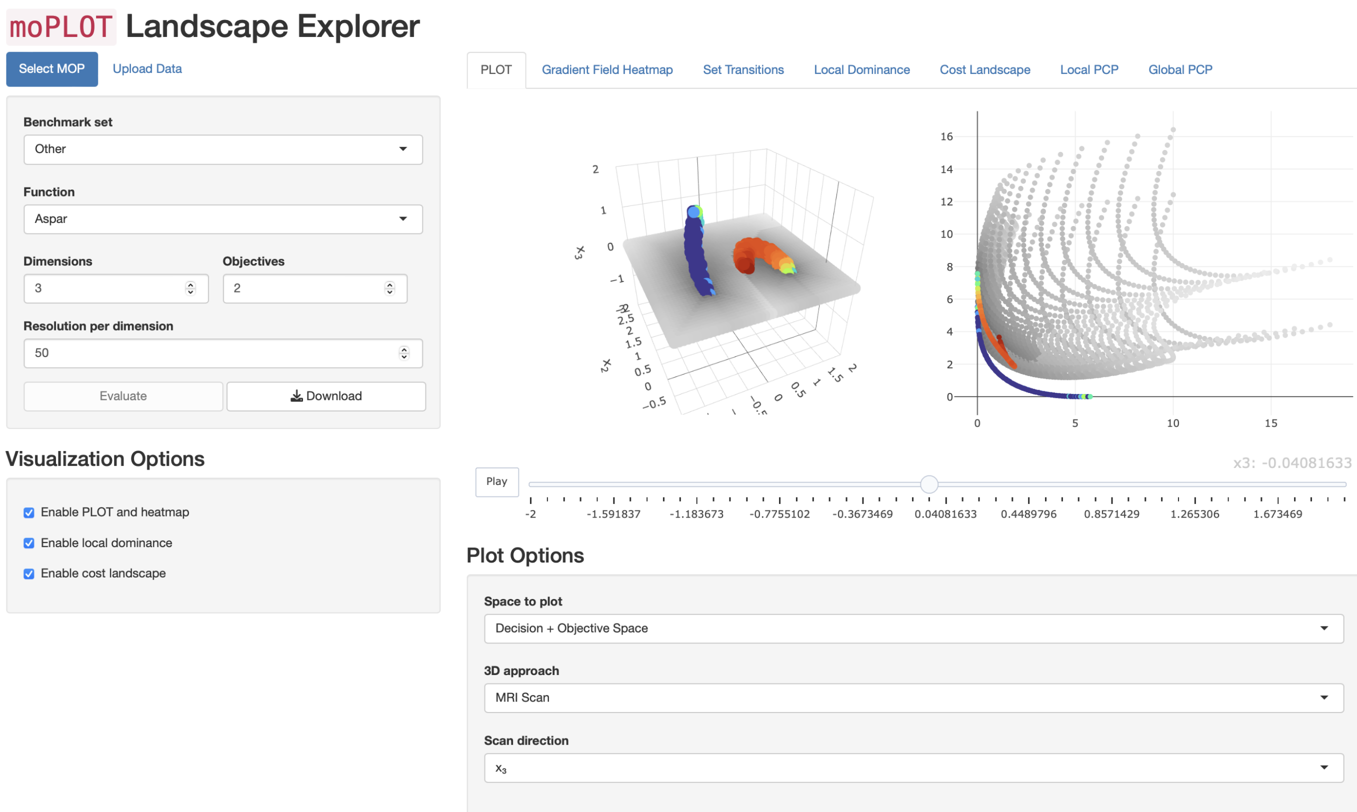

Fig. 4. A screenshot of the moPLOT dashboard demonstrating the PLOT visualization with the *MRI Scan* 3D approach. The shown function is the 3-parameter 2-objective Aspar function [24]. While the problem has two locally efficient sets, which are shown permanently, the non-optimal set intersects the attraction basin of the globally optimal one. This is shown by the selected x_3-slice of the gray-scaled gradient field heatmap, which can be interactively moved by the user.

izations on novel problems, e.g., from practice, requires good software support.

In particular, all visualization techniques should be implemented in the same software package to enable a comparable application of them on the given test problems. Further, a wide range of benchmark suites should be available for direct visualization to enable a thorough study of known benchmark problems, with a possibility to incorporate further test problems by extension mechanisms. Finally, it should enable also novice users to create and study visualizations in a seamlessly accessible environment.

So far, individual implementations of the visualizations, like accompanying code together with their original publications, rely on a range of software environments (such as R and Matlab). Others need to be implemented by users themselves as there is no standard implementation. This is the case for the contour and cost landscape visualizations. This limits their applicability for novice users and hinders the comparison of landscape structures in the context of benchmarking.

For a few dedicated benchmarks, however, there exist websites[4], which provide visualizations of their test problems. For the bi-objective BBOB an almost comprehensive overview of available decision space visualizations is available, plus some visualization techniques focused on dimensionality reduction. However, the visualizations are limited to the bi-objective BBOB problems and visual comparisons to further benchmark suites cannot be drawn. Additionally, novel problems cannot be visualized interactively, which limits the scope of this approach.

The here presented moPLOT dashboard aims at providing a solution for this problem. An initial version of this software has already been published within a conference paper [24]. Therein, we have implemented all the visualization methods that were published at that time (including the interactive components) and embedded them into an interactive web application. It enables (i) the selection of a test function from a variety of benchmark problems and (ii) the upload of external test data for visualizing novel problems. We herein refined this implementation of moPLOT by extending it, among other things, the new methods presented in Section III-B.

Common benchmark suites such as the *Extended* Bi-objective BBOB [51], DTLZ [9], ZDT [8], MMF (CEC 2019) [13], MOP [52] and combinations of MPM2 [53] functions are integrated directly into the dashboard. In addition, a variety of individual test functions which are commonly used in the literature or which are useful for demonstrating particular landscape properties, such as the Aspar [21], Kursawe [6], SGK [21], Viennet [7] and Bi-Rosenbrock (see Section V) functions, are available.

A publicly accessible version of the dashboard is hosted online on https://schaepermeier.shinyapps.io/moPLOT. For im-

[4] https://numbbo.github.io/bbob-biobj/

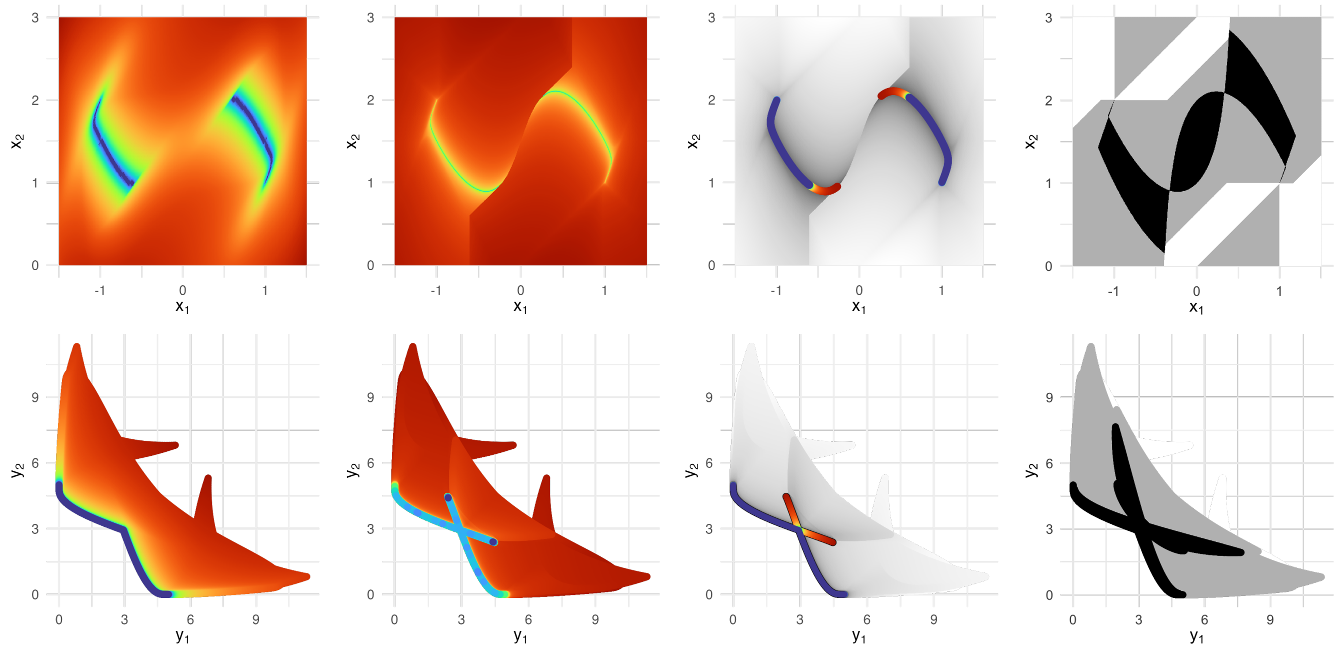

Fig. 5. Overview of decision space visualization techniques (top) with corresponding visualizations in objective space (bottom) for the Bi-Rosenbrock problem. Visualizations from left to right: Cost landscape [17], gradient field heatmap [19], PLOT [21] and local dominance landscape [29].

plementation details, we refer to [24] and the source code, which is available at https://github.com/kerschke/moPLOT.

Figure 4 demonstrates the user interface of the `moPLOT` dashboard with an example of the 3D *MRI Scan* visualization. In the upper left drop-down menu, a test function can be selected (and parameterized in many cases). Subsequently it will be evaluated on a grid of points to generate the necessary data for the MOP's visualization. Alternatively, the grid data may be uploaded in CSV format in the second tab. For this purpose, the user can evaluate a real-world problem, novel benchmark or surrogate function on a self-defined grid of points outside of the dashboard. The obtained data set of decision and objective space values can then be uploaded to the dashboard for visualizing the underlying landscape. A corresponding download functionality is also included for expensive to evaluate functions. Once evaluated and plotted, the results can be exported and later reloaded when needed for another visualization. Then, on the lower part of the left side panel, individual visualization groups may be enabled. More importantly, they can also be disabled, e.g., if they are too costly to compute (as is the case for the cost landscape for high resolutions) or not of interest in an individual use case.

The right hand side then provides access to all implemented (and enabled) visualizations. Here, the user can, for instance, select whether to show the decision or objective space (or both), if supported by the chosen visualization. Further options may also be presented underneath the visualization area. These most notably comprise options regarding the 3D approach in the case of a three-dimensional decision space.

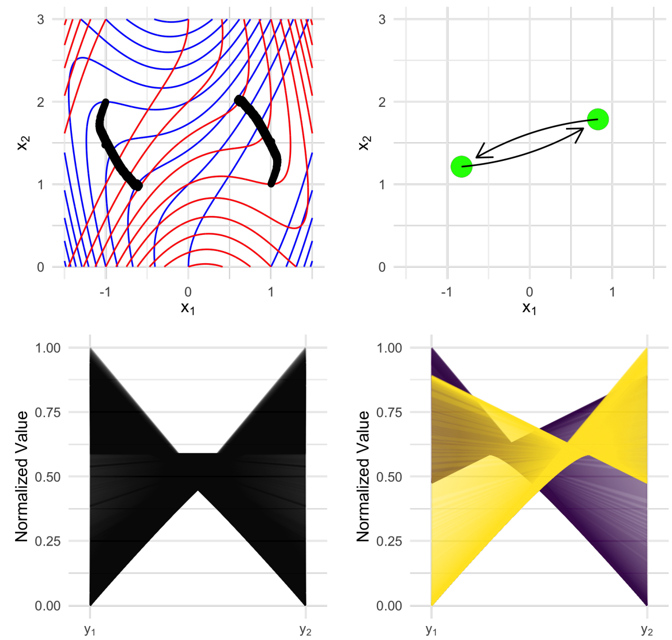

Fig. 6. Additional visualizations on the Bi-Rosenbrock problem. Top row: Multi-contour plot with superposed Pareto set (left) and set transition graph (right). Bottom row: Global PCP (left) and local PCP (right). The set transition graph (upper right) clearly shows the mutual superposition of the two locally efficient sets.

V. Usage Scenarios

This section presents four usage scenarios [54], which demonstrate how the visualizations (see Section III) and dashboard (see Section IV) can be used for comparison, exploration and analysis of multi-objective optimization problems.

A. A Bi-Rosenbrock Problem

In this first scenario, we consider a particular two-dimensional Bi-Rosenbrock problem, whose component func-

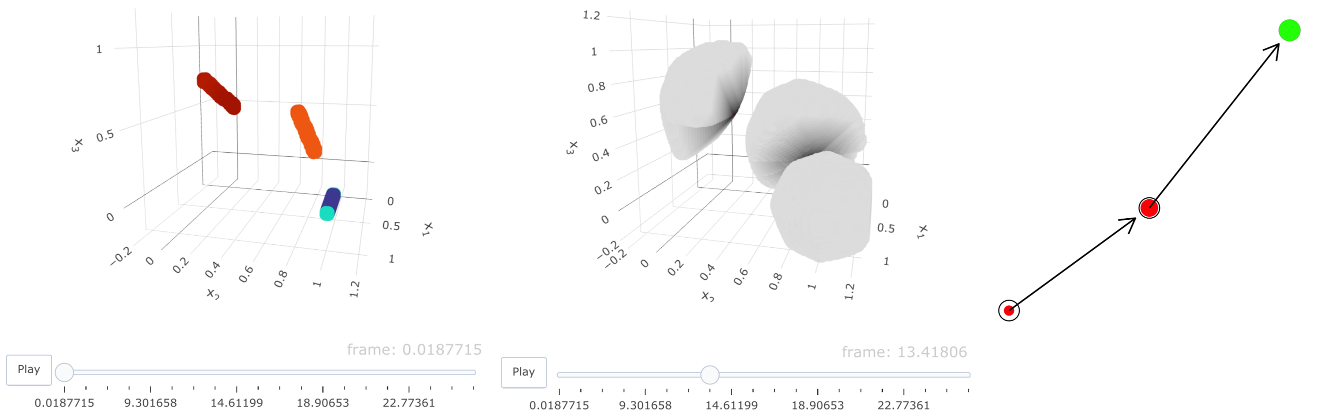

Fig. 7. Selected visualizations of the three-dimensional SGK problem. Left: 3D visualization of locally efficient sets in decision space. Middle: One step of the interactive *Onion Layer* visualization of attraction basins in decision space from the same viewpoint. Right: Set transition graph. The two basin intersections modeled by the set transition graph are clearly visible in the middle visualization. The fully filled circle corresponding to the globally efficient set indicates that all local search runs would be successful on this problem, even though it contains three locally efficient sets.

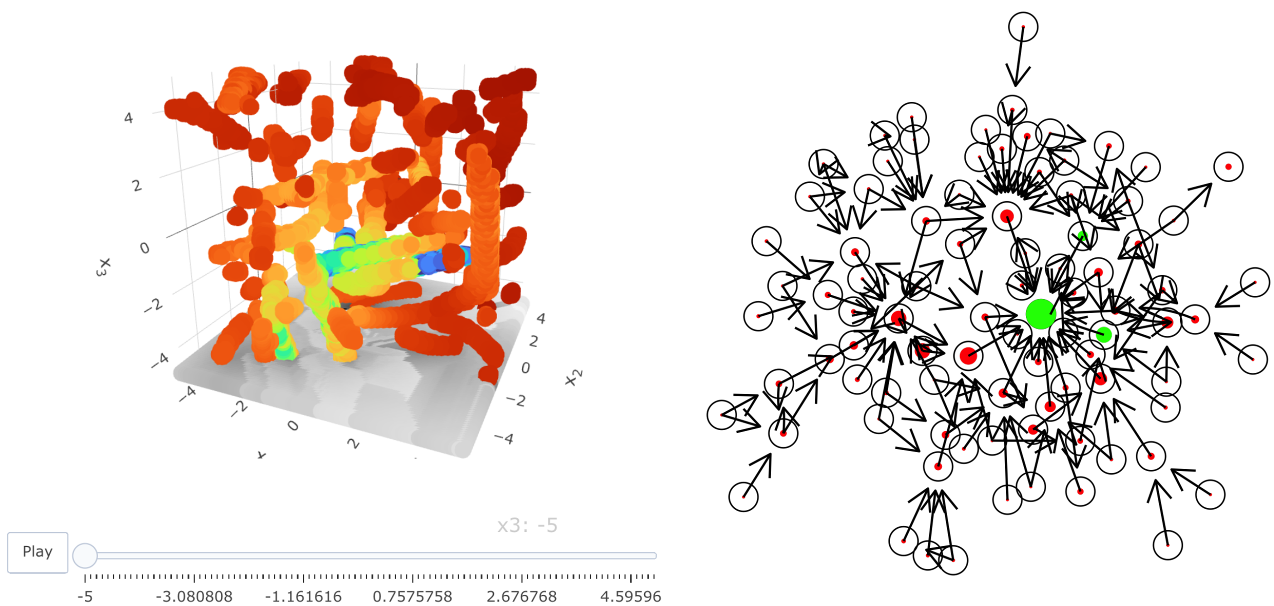

Fig. 8. PLOT visualization with the *MRI Scan* technique and set transition graph of a three-dimensional problem from the bi-objective BBOB (FID: 10, IID: 12). While the analysis of the basin intersections is essentially impossible due to the overload of information in the PLOT visualization, some information can still be gathered from the set interaction graph: There are three locally efficient sets contributing to the Pareto set, where one is reachable from any starting point while the others might require multiple restarts to find.

tions are given by $f_1(x_1, x_2) = (1 - x_1)^2 + (x_2 - x_1^2)^2$ and $f_2(x_1, x_2) = (1+x_1)^2 + (-(x_2-3) - x_1^2)^2$, respectively. This essentially results in two copies of the same landscape rotated by 180°. The previously discussed landscape visualizations for this problem are given in Figures 5 and 6.

We observe, particularly in the GFH and PLOT visualizations, that there are two locally efficient sets that intersect each other. Assuming a local search algorithm which first discovers only one efficient set by simple multi-objective gradient descent, it can subsequently explore along that connected set and finally cross the ridge into the second basin of attraction. There it finds the second locally efficient set and thus also the second part of the global Pareto set during this process. That is, the given Bi-Rosenbock problem features multiple locally efficient sets and a disconnected Pareto set, but can be solved optimally from any starting point by a local search algorithm. Note that this property is only shown in the more recent visualization techniques. In fact, the cost landscape, the contour plot, and in this case even the local dominance landscape would not lead to the mentioned conclusions. Finally, while the PCP visualizations may show some properties of the objective space – or the local PCP even properties of the individual locally efficient fronts – they are unable to demonstrate local search dynamics at all.

This demonstrates that complex multi-objective benchmarking landscapes with very specific problem structure may even arise from two rather simple, unimodal single-objective functions. Their interaction, characteristics, and challenges (or

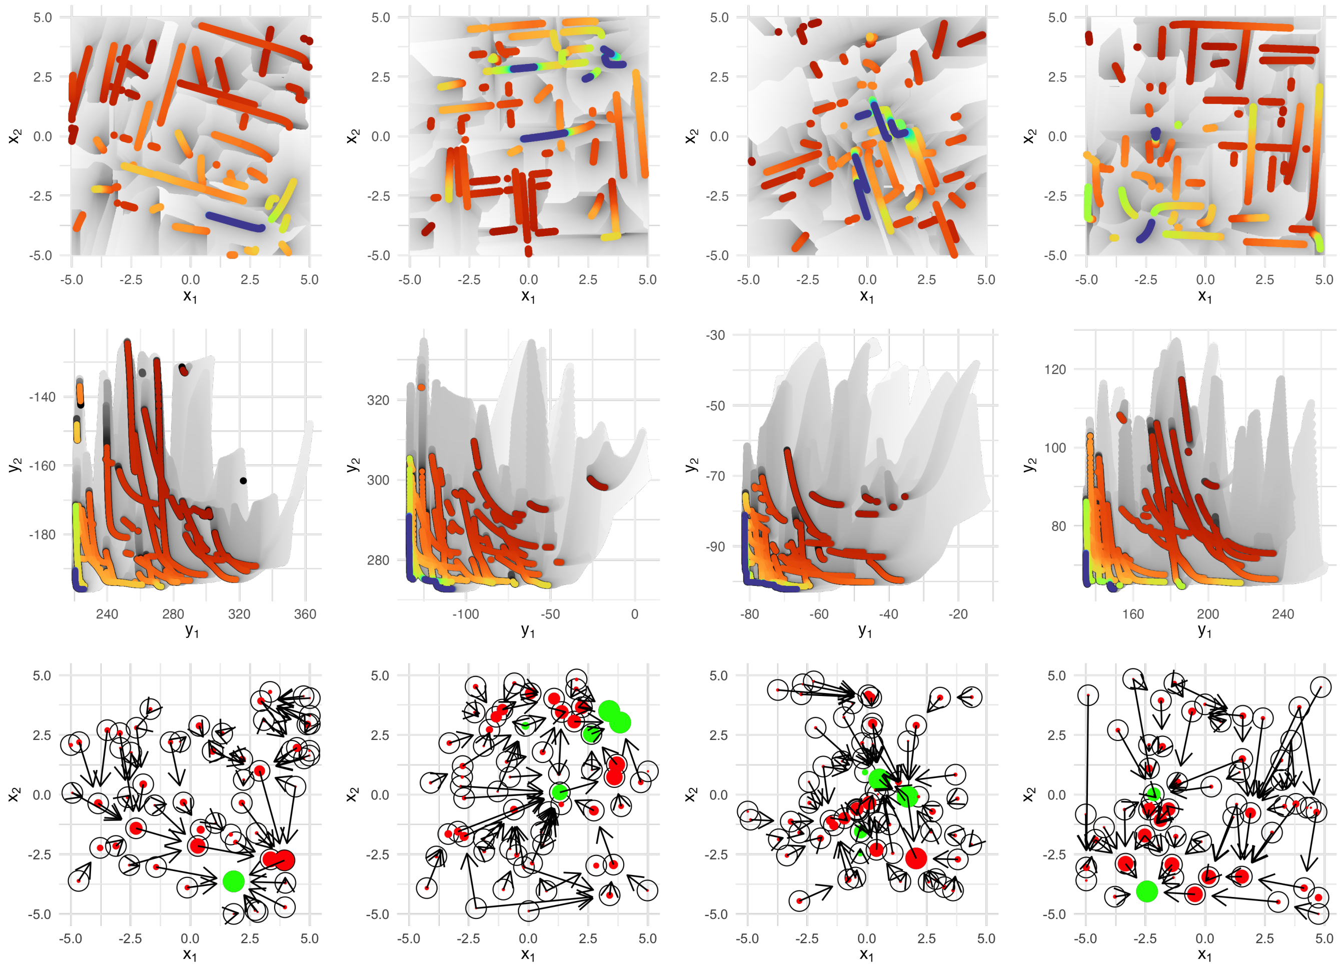

Fig. 9. Visual comparison of different instances (IIDs left to right: 5, 10, 13 and 15) of the 10-th function from the Bi-Objective BBOB in the PLOT visualization in decision and objective space, as well as the set transition graph. While the different instances of a function in the BBOB should have comparable properties, and they are all shown to contain many locally efficient sets, the visualizations reveal that other properties can vary greatly: The Pareto set may be connected (IID 3), or disconnected in many (IIDs 10, 13) or only two sets (IID 15). In addition, the Pareto front may be connected (IIDs 1, 13) or disconnected (IIDs 10, 15). The set transition graph finally shows that some Pareto set components may be hard to reach by local search as well (IIDs 10, 13).

simplicity) for certain types of solvers are only revealed, if we apply an adequate visualization technique.

B. 3D Visualization

As a second example, the dashboard's suitability for analyzing a three-dimensional MOP are highlighted. The three-dimensional SGK problem [21] is given by two objectives, one of them being unimodal, the other one being trimodal. Selected visualizations are shown in Figure 7.

Each of the locally efficient sets is superposed by the attraction basin of a dominating set, leading a local solver to reliably discover the globally efficient set out of the three sets. This property is reflected perfectly by the set transition graph, which uses a stress-based graph layout due to the higher search space dimensionality, and shows a fully filled green node for the globally efficient set.

Finally, Figure 8 shows an example with a more complex three-dimensional landscape. Due to the large amount of locally efficient sets, visually analyzing the resulting landscape in the MRI scan for PLOT (left-hand side) becomes a very tedious task. In this case, the set transition graph allows some further and simplified analyses, including the observation of three locally efficient sets (represented by green nodes), which contribute to the Pareto set. We also observe that one of the sets is reachable from any starting point, as denoted by the completely filled green node in the graph.

C. Analyzing Bi-Objective BBOB Functions

Two-dimensional visualizations can also aid in the design and analysis of a benchmark suite with poorly-understood or unknown properties. We demonstrate this exemplarily for different instances of the 10th function from the two-dimensional Bi-Objective BBOB [16] using the PLOT and set transition graph visualizations, which are depicted in Figure 9. This test problem consists of a sphere function for the first and a (rotated) Gallagher's 101 peaks function, which is multimodal without global structure, for the second objective.

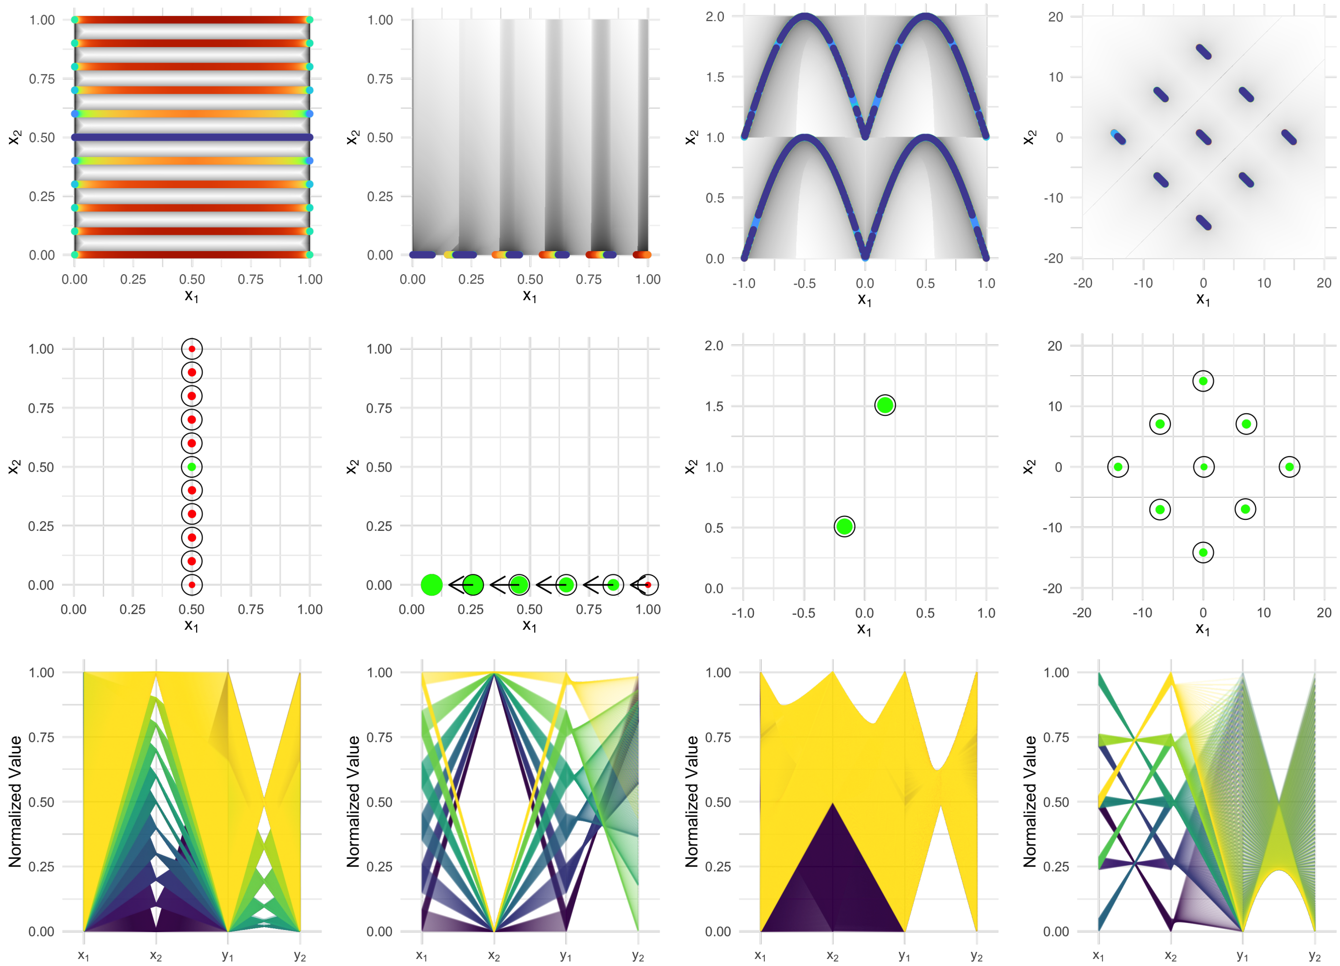

Fig. 10. Visual comparison of multimodal benchmark functions from literature: The benchmark functions (left to right) DTLZ1 [9], ZDT3 [8], MMF4 [13] and SYMPART-rotated [55] are shown with the (top to bottom) PLOT, set transition graph, and local PCP visualizations.

While the test functions in this benchmark are constructed from well-understood component functions from the BBOB suite, the resulting instances may have properties that are unknown beforehand and which can vary greatly across instances. As demonstrated in Figure 9, the shown instances are all multimodal by featuring many locally efficient sets. At the same time, the Pareto set and front may be disconnected or connected, and those may be either easy or hard to find by local search runs. Understanding these properties by means of (combinations of) appropriate visualization techniques (e.g., PLOT combined with transition graphs) may help to better understand algorithm behavior and performance on different instances, supporting benchmark designers to construct benchmark suites with well-understood properties and algorithm designers to get a better picture of (local) search dynamics of the problem instances within a benchmark. In the end, this could pave the way to developing novel algorithmic concepts.

D. Multimodality in Classical MO Benchmarks

In the final usage scenario, we show how visualizations can support the comparison of functions across different benchmark problems. Here, we focus on well-known multimodal functions from literature, in particular the DTLZ1 [9], ZDT3 [8], MMF4 [13] and SYMPART-rotated [55] functions (see Figure 10).

The PLOT and set transition graph visualizations highlight that many of the shown test functions do not feature set interactions at all. As only exception, ZDT3 has a very simple set interaction graph, in which the sets are essentially chained to each other. This chaining leads to a globally efficient set that is guaranteed to be found (at $x_1 \in [0.00, 0.10]$), while the two rightmost sets – the globally efficient set at $x_1 \in [0.75, 0.85]$, and the locally efficient set at $x_1 \in [0.95, 1.00]$ (constituting the start of the chain) – may be hard to find. Another trend in benchmark design is to design problems, which are purely multi-global [5], such as the shown SYMPART-rotated [55] problem. These problems only comprise locally (and simultaneously globally) efficient sets that are isolated from each other, and do not contain local search traps that are not globally efficient. The focus of such problems is on multimodal MO optimization, i.e., searching for multiple well-distributed solutions (in the search space) that result in comparable values in the objective space. However, these properties are distinctly different from the local search properties of, e.g., the Bi-

Objective BBOB and MPM2 functions visualized above. One should thus be careful to generalize results of algorithms evaluated only on a limited set of problems.

Finally, the local PCP visualizations here are shown using the decision *and* objective space values. The usefulness for multimodal problems is best shown for the DTLZ1 and SYMPART-rotated problems. Next to the objective space properties, local PCP can illustrate the different placements of the locally efficient sets in decision space.

VI. Conclusion

In this work, we presented a comprehensive overview of visualization techniques for multi-objective optimization problems (MOPs). Beyond techniques that enable the depiction of MOPs with two-dimensional decision and objective spaces, we (a) presented two three-dimensional interactive visualization techniques and (b) introduced two novel approaches that adapt and refine the concepts of established visualization methods (PCPs, as well as LONs or PLOS-nets) and thereby pave the way towards visualizing high-dimensional and many-objective problems. In addition, we presented an updated version of an integrated and web-based dashboard, named `moPLOT`, for seamlessly accessing the visualizations and enabling interactive analyses of test problems – even for users without a programming background. Finally, the usefulness and complementarity of the presented visualization techniques, and the analyses enabled by them, were demonstrated in a series of usage scenarios in which several challenges in the context of benchmarking MOPs were discussed on the one hand and revealed in the corresponding landscape visualizations on the other hand.

In addition, we have demonstrated that the benefits of the different visualization techniques are manifold for researchers of the benchmarking community and beyond:

1) A wide selection of visualization techniques provides a variety of perspectives on benchmark problems for their designers. Visual analyses enable benchmark designers to describe the challenges of their problems and to classify these problems adequately.
2) The various visualization methods – made accessible using the dashboard – enable the comparison of already existing benchmark problems from established benchmark suites. Also, the updated dashboard provides the functionalities to visualize custom and real-world problems and visually compare them to existing benchmark suites without the need for specific function definitions.
3) Beyond that, decision space-focused visualizations of MOPs enable more profound analyses of landscape properties for algorithm designers, helping them to gain new insights into the characteristics and the challenges of their (benchmark) problems. These insights enable the development of new search principles – almost as intuitive as in the single-objective case. For first examples, see [22], [25], [48], [56].

Finally, the improved understanding of landscape properties may boost research in benchmarking towards understanding search dynamics [5], [57], development of informative landscape features [58]–[60], as well as automated algorithm selection [61] in future work. A further perspective for future work is provided by the challenge of adapting the novel scalable visualization techniques to higher-dimensional spaces – including the visualization of many-objective problems [62]. While the set transition graph and local PCP are, in principle, scalable to an arbitrary dimensionality, the currently implemented framework relies on a densely evaluated grid of points, whose scalability is limited by the curse of dimensionality. Using these visualizations on real-world problems that only allow for sparse evaluations presents another future research direction. This may be accomplished by a number of different techniques, such as surrogate models built on the evaluated points [63]. Currently, pre-processing on the evaluated points is required to ensure that a sufficiently dense grid is sampled for the approaches presented in this work.

Acknowledgment

The authors acknowledge support by the *European Research Center for Information Systems (ERCIS)*. Further, L. Schäpermeier and P. Kerschke acknowledge support by the *Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig*.

References

[1] S. Wright, "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," in *Proceedings of the 6th International Congress of Genetics*, vol. 1, no. 8, 1932, pp. 356 – 366.

[2] M. Preuss, *Multimodal Optimization by Means of Evolutionary Algorithms*. Springer, 2015.

[3] C. A. Coello Coello, D. A. van Veldhuizen, and G. B. Lamont, *Evolutionary Algorithms for Solving Multi-Objective Problems*, 2nd ed. Springer, 2007.

[4] K. Miettinen, *Nonlinear Multiobjective Optimization*, ser. International Series in Operations Research & Management Science. Springer, 1998, vol. 12. [Online]. Available: https://link.springer.com/book/10.1007/978-1-4615-5563-6

[5] C. Grimme, P. Kerschke, P. Aspar, H. Trautmann, M. Preuss, A. H. Deutz, H. Wang, and M. Emmerich, "Peeking Beyond Peaks: Challenges and Research Potentials of Continuous Multimodal Multi-objective Optimization," *Computers & Operations Research*, vol. 136, p. 105489, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0305054821002367

[6] F. Kursawe, "A Variant of Evolution Strategies for Vector Optimization," in *Proceedings of the 1st International Conference on Parallel Problem Solving from Nature (PPSN I)*, ser. Lecture Notes in Computer Science (LNCS), H.-P. Schwefel and R. Männer, Eds., vol. 496. Springer, October 1990, pp. 193 – 197. [Online]. Available: https://link.springer.com/chapter/10.1007/BFb0029752

[7] R. Viennet, C. Fonteix, and I. Marc, "Multicriteria Optimization Using a Genetic Algorithm for Determining a Pareto Set," *International Journal of Systems Science*, vol. 27, no. 2, pp. 255 – 260, 1996. [Online]. Available: https://www.tandfonline.com/doi/abs/10.1080/00207729608929211

[8] E. Zitzler, K. Deb, and L. Thiele, "Comparison of Multiobjective Evolutionary Algorithms: Empirical Results," *Evolutionary Computation Journal (ECJ)*, vol. 8, no. 2, pp. 173 – 195, June 2000. [Online]. Available: http://www.mitpressjournals.org/doi/10.1162/106365600568202

[9] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable Test Problems for Evolutionary Multiobjective Optimization," in *Evolutionary Multiobjective Optimization*, ser. Advanced Information and Knowledge Processing (AI & KP), A. Abraham, L. Jain, and R. Goldberg, Eds. Springer, 2005, pp. 105 – 145. [Online]. Available: https://link.springer.com/chapter/10.1007/1-84628-137-7_6

[10] S. Huband, P. Hingston, L. Barone, and L. While, "A Review of Multiobjective Test Problems and a Scalable Test Problem Toolkit," *IEEE Transactions on Evolutionary Computation (TEVC)*, vol. 10, no. 5, pp. 477 – 506, 2006. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/1705400/

[11] M. T. M. Emmerich and A. H. Deutz, "Test Problems Based on Lamé Superspheres," in *Proceedings of the 4th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*, ser. Lecture Notes in Computer Science (LNCS), vol. 4403. Springer, 2007, pp. 922 – 936. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-540-70928-2_68

[12] Q. Zhang, A. Zhou, S. Zhao, P. N. Suganthan, W. Liu, and S. Tiwari, "Multiobjective Optimization Test Instances for the CEC 2009 Special Session and Competition," University of Essex, Colchester, UK and Nanyang Technological University, Singapore, Tech. Rep., 2008, special Session on Performance Assessment of Multi-Objective Optimization Algorithms. [Online]. Available: http://dces.essex.ac.uk/staff/zhang/MOEAcompetition/cec09testproblem0904.pdf

[13] C. Yue, B. Qu, K. Yu, J. Liang, and X. Li, "A Novel Scalable Test Problem Suite for Multimodal Multiobjective Optimization," *Swarm and Evolutionary Computation*, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S2210650218307387

[14] T. Tušar, "Visualizing Solution Sets in Multiobjective Optimization," Ph.D. dissertation, Jožef Stefan International Postgraduate School, June 2014.

[15] T. Tušar and B. Filipič, "Visualization of Pareto Front Approximations in Evolutionary Multiobjective Optimization: A Critical Review and the Prosection Method," *IEEE Transactions in Evolutionary Computation (TEVC)*, vol. 19, no. 2, pp. 225 – 245, April 2015. [Online]. Available: https://ieeexplore.ieee.org/document/6777535/

[16] T. Tušar, D. Brockhoff, N. Hansen, and A. Auger, "COCO: The Bi-Objective Black Box Optimization Benchmarking (bbob-biobj) Test Suite," *arXiv preprint*, vol. abs/1604.00359, May 2016. [Online]. Available: https://arxiv.org/abs/1604.00359

[17] C. M. M. d. Fonseca, "Multiobjective Genetic Algorithms with Application to Control Engineering Problems," Ph.D. dissertation, Department of Automatic Control and Systems Engineering, University of Sheffield, September 1995. [Online]. Available: https://eden.dei.uc.pt/~cmfonsec/fonseca-phd.pdf

[18] P. Kerschke, H. Wang, M. Preuss, C. Grimme, A. H. Deutz, H. Trautmann, and M. T. M. Emmerich, "Towards Analyzing Multimodality of Multiobjective Landscapes," in *Proceedings of the 14th International Conference on Parallel Problem Solving from Nature (PPSN XIV)*, ser. Lecture Notes in Computer Science (LNCS), J. Handl, E. Hart, P. R. Lewis, M. López-Ibánez, G. Ochoa, and B. Paechter, Eds., vol. 9921. Springer, September 2016, pp. 962 – 972. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-45823-6_90

[19] P. Kerschke and C. Grimme, "An Expedition to Multimodal Multi-Objective Optimization Landscapes," in *Proceedings of the 9th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*, ser. Lecture Notes in Computer Science (LNCS), H. Trautmann, G. Rudolph, K. Kathrin, O. Schütze, M. Wiecek, Y. Jin, and C. Grimme, Eds., vol. 10173. Springer, March 2017, pp. 329 – 343. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-319-54157-0_23

[20] C. Grimme, P. Kerschke, and H. Trautmann, "Multimodality in Multi-Objective Optimization — More Boon than Bane?" in *Evolutionary Multi-Criterion Optimization*, ser. Lecture Notes in Computer Science (LNCS), K. Deb, E. Goodman, C. A. Coello Coello, K. Klamroth, K. Miettinen, S. Mostaghim, and P. Reed, Eds., vol. 11411. Springer, 2019, pp. 126 – 138. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-12598-1_11

[21] L. Schäpermeier, C. Grimme, and P. Kerschke, "One PLOT to Show Them All: Visualization of Efficient Sets in Multi-objective Landscapes," in *International Conference on Parallel Problem Solving from Nature (PPSN)*. Springer, 2020, pp. 154 – 167.

[22] P. Kerschke and C. Grimme, "Lifting the multimodality-fog in continuous multi-objective optimization," in *Metaheuristics for Finding Multiple Solutions*. Springer, 2021, pp. 89 – 111.

[23] J. E. Fieldsend and K. Alyahya, "Visualising the Landscape of Multi-Objective Problems Using Local Optima Networks," in *Proceedings of the 21st Annual Conference on Genetic and Evolutionary Computation (GECCO)*, 2019, pp. 1421 – 1429.

[24] L. Schäpermeier, C. Grimme, and P. Kerschke, "To Boldly Show What No One Has Seen Before: A Dashboard for Visualizing Multi-objective Landscapes," in *Evolutionary Multi-Criterion Optimization*, H. Ishibuchi, Q. Zhang, R. Cheng, K. Li, H. Li, H. Wang, and A. Zhou, Eds. Cham: Springer International Publishing, 2021, pp. 632–644.

[25] C. Grimme, P. Kerschke, M. T. M. Emmerich, M. Preuss, A. H. Deutz, and H. Trautmann, "Sliding to the Global Optimum: How to Benefit from Non-Global Optima in Multimodal Multi-Objective Optimization," in *AIP Conference Proceedings*. AIP Publishing, 2019, pp. 020 052–1–020 052–4. [Online]. Available: https://aip.scitation.org/doi/pdf/10.1063/1.5090019?class=pdf

[26] J. Fliege and B. F. Svaiter, "Steepest Descent Methods for Multicriteria Optimization," *Mathematical Methods of Operations Research*, vol. 51, no. 3, pp. 479 – 494, 2000.

[27] J.-A. Désidéri, "Multiple-Gradient Descent Algorithm (MGDA) for Multiobjective Optimization," *Comptes Rendus Mathematique*, vol. 350, no. 5-6, pp. 313 – 318, 2012.

[28] A. A. Törn and A. Zilinskas, "Global optimization," in *Lecture Notes in Computer Science*, vol. 350. Heidelberg: Springer Verlag, 1989.

[29] J. E. Fieldsend, T. Chugh, R. Allmendinger, and K. Miettinen, "A Feature Rich Distance-Based Many-Objective Visualisable Test Problem Generator," in *Proceedings of the 2019 Genetic and Evolutionary Computation Conference (GECCO)*. ACM, 2019, pp. 541 – 549.

[30] K. H. Ang, G. Chong, and Y. Li, "Visualization technique for analyzing non-dominated set comparison," in *Proceedings of the 4th Asia-Pacific Conference on Simulated Evolution and Learning (SEAL'02)*, vol. 1, 2002, p. 36.

[31] S. Obayashi and D. Sasaki, "Visualization and data mining of pareto solutions using self-organizing map," in *Proceedings of the 2nd International Conference on Evolutionary Multi-Criterion Optimization*. Berlin, Heidelberg: Springer-Verlag, 2003, pp. 796–809.

[32] S. Chen, D. Amid, O. M. Shir, L. Limonad, D. Boaz, A. Anaby-Tavor, and T. Schreck, "Self-organizing maps for multi-objective pareto frontiers," in *2013 IEEE Pacific Visualization Symposium (PacificVis)*, 2013, pp. 153–160.

[33] T. Kohonen, *Self-organizing maps*, 3rd ed., ser. Springer series in information sciences, 30. Berlin: Springer, 2001.

[34] A. Ibrahim, S. Rahnamayan, M. V. Martin, and K. Deb, "3d-radvis: Visualization of pareto front in many-objective optimization," in *2016 IEEE Congress on Evolutionary Computation (CEC)*, 2016, pp. 736–745.

[35] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine Series 1*, vol. 2, pp. 559–572, 1901.

[36] J. B. Kruskal and M. Wish, *Multidimensional scaling*. Sage, 1978, no. 11.

[37] X. Zhao, Y. Wu, W. Cui, X. Du, Y. Chen, Y. Wang, D. L. Lee, and H. Qu, "Skylens: Visual analysis of skyline on multi-dimensional data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 246–255, 2018.

[38] A. Pryke, S. Mostaghim, and A. Nazemi, "Heatmap visualization of population based multi objective algorithms," in *Proceedings of the 4th International Conference on Evolutionary Multi-Criterion Optimization*. Springer, 2007, pp. 361—-375.

[39] L. Cibulski, H. Mitterhofer, T. May, and J. Kohlhammer, "Paved: Pareto front visualization for engineering design," *Computer Graphics Forum*, vol. 39, no. 3, pp. 405–416, 2020.

[40] A. V. Lotov and K. Miettinen, *Visualizing the Pareto Frontier*. Springer, 2008, pp. 213–243.

[41] J. Bossek, "smoof: Single- and Multi-Objective Optimization Test Functions," *The R Journal*, 2017. [Online]. Available: https://journal.r-project.org/archive/2017/RJ-2017-004/index.html

[42] P. Kerschke, M. Preuss, C. Hernández, O. Schütze, J.-Q. Sun, C. Grimme, G. Rudolph, B. Bischl, and H. Trautmann, "Cell Mapping Techniques for Exploratory Landscape Analysis," in *EVOLVE — A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation V*, ser. Advances in Intelligent Systems and Computing, A.-A. Tantar, E. Tantar, J.-Q. Sun, W. Zhang, Q. Ding, O. Schütze, M. T. M. Emmerich, P. Legrand, P. Del Moral, and C. A. Coello Coello, Eds. Springer, 2014, vol. 288, pp. 115–131. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-319-07494-8_9

[43] A. Inselberg and B. Dimsdale, "Parallel coordinates: A tool for visualizing multi-dimensional geometry," in *Proceedings of the First IEEE Conference on Visualization*. IEEE, 1990, pp. 361 – 378.

[44] J. Johansson and C. Forsell, "Evaluation of Parallel Coordinates: Overview, Categorization and Guidelines for Future Research," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 1, pp. 579 – 588, 2015.

[45] D. Holten and J. J. Van Wijk, "Evaluation of Cluster Identification Performance for Different PCP Variants," in *Computer Graphics Forum*, vol. 29, no. 3. Wiley, 2010, pp. 793 – 802.

[46] G. Ochoa, M. Tomassini, S. Vérel, and C. Darabos, "A Study of NK Landscapes' Basins and Local Optima Networks," in *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation*, 2008, pp. 555 – 562.

[47] A. Liefooghe, B. Derbel, S. Verel, M. López-Ibáñez, H. Aguirre, and K. Tanaka, "On pareto local optimal solutions networks," in *International Conference on Parallel Problem Solving from Nature*. Springer, 2018, pp. 232 – 244.

[48] L. Schäpermeier, C. Grimme, and P. Kerschke, "MOLE: Digging Tunnels Through Multimodal Multi-Objective Landscapes," in *Proceedings of the 24th Annual Conference on Genetic and Evolutionary Computation (GECCO)*. ACM, 2022, accepted.

[49] P. Kerschke, M. Preuss, S. Wessing, and H. Trautmann, "Low-Budget Exploratory Landscape Analysis on Multiple Peaks Models," in *Proceedings of the 18th Annual Conference on Genetic and Evolutionary Computation (GECCO)*. ACM, July 2016, pp. 229 – 236. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2908812.2908845

[50] M. Li, L. Zhen, and X. Yao, "How to read many-objective solution sets in parallel coordinates," *IEEE Computational Intelligence Magazine*, vol. 12, no. 4, pp. 88 – 100, 2017.

[51] D. Brockhoff, T. Tušar, A. Auger, and N. Hansen, "Using well-understood single-objective functions in multiobjective black-box optimization test suites," *arXiv preprint arXiv:1604.00359*, 2016.

[52] D. A. van Veldhuizen, "Multiobjective Evolutionary Algorithms: Classifications, Analyzes, and New Innovations," Ph.D. dissertation, Faculty of the Graduate School of Engineering of the Air Force Institute of Technology, Air University, June 1999. [Online]. Available: http://neo.lcc.uma.es/emoo/veldhuizen99a.ps.gz

[53] S. Wessing, "Two-Stage Methods for Multimodal Optimization," Ph.D. dissertation, Technische Universität Dortmund, 2015. [Online]. Available: http://hdl.handle.net/2003/34148

[54] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller, "A systematic review on the practice of evaluating visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2818–2827, 2013.

[55] G. Rudolph, B. Naujoks, and M. Preuss, "Capabilities of EMOA to detect and preserve equivalent pareto subsets," in *Proceedings of the 4th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*, ser. Lecture Notes in Computer Science (LNCS), vol. 4403. Springer, 2007, pp. 36 – 50. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-540-70928-2_7

[56] P. Aspar, P. Kerschke, V. Steinhoff, H. Trautmann, and C. Grimme, "Multi^3: Optimizing Multimodal Single-Objective Continuous Problems in the Multi-Objective Space by Means of Multiobjectivization," in *Proceedings of the 11th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*. Springer, 2021, pp. 311 – 322.

[57] R. Tanabe and H. Ishibuchi, "A review of evolutionary multimodal multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 24, no. 1, pp. 193 – 200, 2019.

[58] P. Kerschke and H. Trautmann, "The R-Package FLACCO for Exploratory Landscape Analysis with Applications to Multi-Objective Optimization Problems," in *Proceedings of the IEEE Congress on Evolutionary Computation (CEC)*. IEEE, July 2016, pp. 5262 – 5269. [Online]. Available: http://ieeexplore.ieee.org/document/7748359/

[59] P. Kerschke, H. Wang, M. Preuss, C. Grimme, A. H. Deutz, H. Trautmann, and M. T. M. Emmerich, "Search Dynamics on Multimodal Multi-Objective Problems," *Evolutionary Computation (ECJ)*, vol. 0, pp. 1–30, 2018. [Online]. Available: https://www.mitpressjournals.org/doi/abs/10.1162/evco_a_00234

[60] A. Liefooghe, S. Verel, B. Lacroix, A.-C. Zăvoianu, and J. McCall, "Landscape features and automated algorithm selection for multi-objective interpolated continuous optimisation problems," in *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 2021, pp. 421 – 429.

[61] P. Kerschke, H. H. Hoos, F. Neumann, and H. Trautmann, "Automated Algorithm Selection: Survey and Perspectives," *Evolutionary Computation (ECJ)*, vol. 27, pp. 3 – 45, 2019.

[62] M. Li, C. Grosan, S. Yang, X. Liu, and X. Yao, "Multiline distance minimization: A visualized many-objective test problem suite," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 1, pp. 61 – 78, 2017.

[63] W. Berger, H. Piringer, P. Filzmoser, and E. Gröller, "Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction," in *Computer Graphics Forum*, vol. 30, no. 3. Wiley, 2011, pp. 911 – 920.

Lennart Schäpermeier is a research assistant at the group of *Big Data Analytics in Transportation* at TU Dresden, Germany. He previously received his Bachelor's and Master's degrees in *Information Systems* at the University of Münster, Germany. His main research interests lie in understanding continuous (multi-objective) optimization landscapes, Exploratory Landscape Analysis, and algorithm selection.

Christian Grimme is an associate professor at the *Group for Statistics and Optimization* at the University of Münster, Germany. He received his diploma degree (2006) and PhD (2012) in *Computer Science* from the TU Dortmund University, Germany. In 2019, he obtained his habilitation (venia legendi) in *Information Systems* at the University of Münster, Germany. One of his main research topics is evolutionary MOO with a special focus on fundamentals in MOO and their application in algorithm development.

Pascal Kerschke is professor of *Big Data Analytics in Transportation* at TU Dresden, Germany. His research interests cover a wide range of topics in the context of benchmarking, data science, machine learning, and optimization – including automated algorithm selection, Exploratory Landscape Analysis, as well as continuous single- and multi-objective optimization. Moreover, he is the main developer of the R-package `flacco`, co-authored further R-packages such as `smoof` and `moPLOT`, co-organized numerous tutorials and workshops in the context of *Exploratory Landscape Analysis* and/or benchmarking, and is an active member of the Benchmarking Network and the COSEAL group.