



HAL
open science

Funnels in multi-objective fitness landscapes

Gabriela Ochoa, Arnaud Liefooghe, Sébastien Verel

► **To cite this version:**

Gabriela Ochoa, Arnaud Liefooghe, Sébastien Verel. Funnels in multi-objective fitness landscapes. PPSN 2024 – Parallel Problem Solving from Nature, Sep 2024, Hagenberg, Austria. pp.343-359, 10.1007/978-3-031-70055-2_21 . hal-04692819

HAL Id: hal-04692819

<https://ulco.hal.science/hal-04692819v1>

Submitted on 10 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Funnels in Multi-objective Fitness Landscapes

Gabriela Ochoa¹[0000–0001–7649–5669], Arnaud Liefooghe²[0000–0003–3283–3122],
and Sébastien Verel²[0000–0003–1661–4093]

¹ University of Stirling, Scotland, United Kingdom
`gabriela.ochoa@stir.ac.uk`

² LISIC, Université du Littoral Côte d’Opale, Calais, France
`{arnaud.liefooghe,sebastien.verel}@univ-littoral.fr`

Abstract. Funnels are related to the big-valley hypothesis in combinatorial fitness landscapes. It suggests that local optima are not randomly distributed but are instead clustered around the global optimum, forming a coarse-grained global structure. Multi-funnel structures emerge when more than one cluster of local optima is present, some surrounding sub-optimal solutions. These multi-funnel landscapes can be challenging to search, as the optimisation process may get trapped in a sub-optimal funnel. We propose a characterisation of funnels in multi-objective combinatorial landscapes based on the solution ranks using non-dominated sorting, and a variation of the recent graph model of multi-objective landscapes: the compressed Pareto local optimal solution network (C-PLOS-net). Using a set of ρ mnk-landscapes, we construct and visualise monotonic C-PLOS-nets, and introduce a set of metrics to characterise the landscapes’ funnel structure. The proposed metrics are found to capture the landscape global structure, to correlate with benchmark parameters, and to explain the performance of well-established multi-objective local search and evolutionary algorithms.

Keywords: Multi-objective combinatorial optimisation, landscape analysis, local optima networks, LON, PLOS-net, multi-objective NK landscapes, funnel.

1 Introduction

The study of fitness landscapes has a long tradition in single-objective optimisation [1–3]. The aim is to understand the structure of search spaces in relation to the fitness function and how this structure impacts the performance of optimisation algorithms. Some landscape analysis tools and metrics have been extended to multi-objective optimisation [4–6], and recently included into algorithm selection and recommendation studies [7, 8]. However, to our knowledge, no studies have defined and examined funnels in multi-objective optimisation. Our goal is to bring the concept of funnels from single-objective to multi-objective combinatorial landscapes. More specifically, our contributions are as follows:

- We propose a definition of funnels for multi-objective combinatorial landscapes (Section 3);

- We introduce new features and visualisations that capture the funnel structure of multi-objective landscapes (Sections 3 and 4);
- We explore how the proposed funnel features correlate with algorithm performance (Section 5).

2 Background

2.1 What is a Funnel?

An energy landscape [9] is a mapping of all possible conformations of a molecular entity (clusters, glasses, proteins) to their respective energy levels. The term ‘funnel’ was introduced in this context to refer to “a region of configuration space that can be described in terms of a set of downhill pathways that converge on a single low-energy structure or a set of closely related low-energy structures” [10]. The energy landscape of proteins is characterised by a single deep funnel, accounting for their ability to fold to their native state. Conversely, certain shorter polymer chains that misfold are expected to have additional secondary funnels that can act as traps in the folding process.

Funnels have also been studied in evolutionary computation and fitness landscapes. In continuous optimisation, the *dispersion* and *nearest-better-clustering* metrics [11, 12] were introduced to detect the landscape’s funnel structure. They have subsequently been used in exploratory landscape analysis [13]. In combinatorial optimisation, funnels are related to the “big valley” hypothesis [14], which suggests that local optima in travelling salesperson problems are not distributed randomly. Instead, they are clustered around one central global optimum. The idea of a single valley was later challenged, revealing that complex landscapes are characterised by multiple valleys or funnels [15, 16]. A characterisation of funnels in combinatorial optimisation using local optima networks (LONs) [17] was proposed in [18]. The idea is to consider *monotonic sequences* [19] of local optima, that is, sequences of local minima where fitness is always decreasing. This led to the development of the monotonic (compressed) LON model, where deteriorating edges are removed, and funnels are depicted as the collection of paths to a single *sink* node (a node without outgoing edges). Although global optima are sinks with the best fitness, landscapes can also have sinks with sub-optimal fitness values.

2.2 PLOS-net and Compressed PLOS-net

A multi-objective landscape can be defined as a triplet (X, \mathcal{N}, f) such that X is the solution space, $\mathcal{N}: X \mapsto 2^X$ is a neighbourhood relation, and $f: X \rightarrow Z$ is an objective function vector. A solution $x \in X$ is a *Pareto local optimal solution* (PLOS) if it is not dominated by any of its neighbours [20]; i.e. $\forall x' \in \mathcal{N}(x)$, x' does not dominate x . For $m = 1$, we remark that this is equivalent to the conventional definition of a single-objective local optimum.

Given a multi-objective landscape, the Pareto local optimal solutions network (PLOS-net) model [21] is constructed as an unweighted, directed graph

$G = (N, E)$. The set of nodes N represents the PLOS, and there is an edge $e_{ij} \in E$ from x^i to x^j if $x^j \in \mathcal{N}(x^i)$.

The compressed PLOS-net (C-PLOS-net) [22] adds a numerical attribute to the PLOS-net nodes, which gives the rank of the corresponding solution in the landscape. All solutions from the search space are sorted into different layers of mutually non-dominated solutions, following the non-dominated sorting procedure [23] used, e.g., in NSGA-II [24]. The rank of a solution corresponds to the layer it belongs to, with a lower rank being better and a Pareto optimal solution having a rank of 1. The C-PLOS-net is constructed by compressing the nodes (i) that are connected and (ii) that have the same rank. Therefore, a C-PLOS-net is a weighted, directed graph $G' = (N', E')$ such that:

- The set of nodes N' are connected components of each PLOS-net’s sub-graph induced by the nodes with the same rank r , with $r \in \{1, 2, \dots\}$.
- There is an edge $e'_{ij} \in E'$ if a PLOS within the compressed node i has a neighbour in the compressed node j .

3 Multi-objective Funnels: Definition and Metrics

In order to define funnels in the context of multi-objective optimisation, we determine the quality of Pareto local optimal solutions by considering their ranks.

3.1 Rank-Distance Correlation

Taking inspiration from the well-known fitness-distance correlation (fdc) landscape metric from single-objective optimisation [25], we start by introducing the *rank-distance correlation coefficient* (rdc) as the Spearman correlation between the ranks and the Hamming distances of Pareto local optimal solutions to Pareto optimal solutions. Specifically, for each PLO solution i in the search space we have a pair (r_i, d_i) , where r_i is the rank of PLO i and d_i is the Hamming distance from i to its closest Pareto optimal solution. We compute this metric using the fully enumerated set of local and global Pareto-optimal solutions in the landscapes.

We note that a multi-objective fitness-distance correlation measure was introduced as the correlation between distances in the variable space and in the objective space among the Pareto global optimal solutions [26, 27]. This measure reflects the relative difficulty of moving “along” the Pareto front. However, it is crucial to understand that this is fundamentally different from rdc, which considers all Pareto local optimal solutions and the difficulty of improving from local to global optima.

3.2 Monotonic Compressed PLOS-net and Pareto Optimal Funnel

Similar to the case of single-objective optimisation [18], we define the monotonic C-PLOS-net (MC-PLOS-net for short) as a directed graph that retains improving edges only, that is, edges leading from higher-ranked (i.e. worst) nodes to

Table 1: Description of funnel metrics.

metric	description
rdc	rank distance correlation coefficient
funnel_prop	proportion of nodes in the Pareto optimal funnel
funnel_depth	longest (weighted) path length from source to Pareto optimal nodes
funnel_paths	number of pathways from source to Pareto optimal nodes
sink_num	number of sub-optimal sinks

lower-ranked (i.e. better) ones. Notice that the MC-PLOS-net does not contain neutral edges (i.e. edges between PLO solutions with the same rank) as those are compressed into single nodes. From the MC-PLOS-net, we can easily detect the *sink* and *source* nodes, as those nodes in the graph without outgoing and incoming edges, respectively.

Once the MC-PLOS-net is constructed, we define the *Pareto optimal funnel* as the aggregation of all pathways in the network that lead to Pareto optimal nodes. MC-PLOS-nets are directed graphs where all edges are improving, meaning they connect nodes with decreasing rank values. Every descending path in the network necessarily ends in a *sink*, that is a node without outgoing edges. When sink nodes have a rank of 1, they are Pareto optimal solutions. However, sinks can have a higher rank. From the MC-PLOS-net, we compute the network metrics described at the bottom of Table 1.

4 Fitness Landscape Analysis and Visualisation

4.1 Benchmark Problems

We consider ρ mnk-landscapes [28] as multi-objective multi-modal problems with objectives correlation. Candidate solutions are binary strings of length n . The neighbourhood (\mathcal{N} in section 2.2) is based on the well-established *1-bit-flip* operator: two solutions are neighbours if the Hamming distance between them is equal to one. The objective function vector $f = (f_1, \dots, f_i, \dots, f_m)$ is defined as $f: \{0, 1\}^n \rightarrow [0, 1]^m$ such that each objective f_i is to be maximised. The objective value $f_i(x)$ of a solution $x = (x_1, \dots, x_j, \dots, x_n)$ is the average value of individual contributions associated with each variable x_j . The contribution of x_j depends on its own value and on the values of $k < n$ variables other than x_j , chosen uniformly at random. By increasing k , landscapes can be gradually tuned from smooth to rugged. The contribution values follow a multivariate uniform distribution such that $\rho > \frac{-1}{m-1}$ defines the correlation among the objectives. The positive (resp. negative) correlation ρ decreases (resp. increases) the degree of conflict between the objective values. ρ mnk-landscapes show different characteristics and degrees of difficulty for multi-objective algorithms [8].

We generate 240 ρ mnk-landscapes following the parameters listed in Table 2. This allows us to investigate landscapes ranging from smooth to rugged, with

Table 2: Benchmark parameters (10 instances are randomly generated for each parameter combination).

description	values
number of variables	$n = 16$
number of interactions	$k \in \{0, 1, 2, 4\}$
number of objectives	$m \in \{2, 3\}$
objectives correlation	$\rho \in \{-0.4, 0.0, 0.4\}$ s.t. $\rho > \frac{-1}{m-1}$

two and three objectives, and conflicting, uncorrelated or correlated objectives. These are small landscapes that can be exhaustively enumerated. We then proceed to generate the PLOS-net, C-PLOS-net and MC-PLOS-net models. The metrics and visualisations reported in this article focus on the newly proposed MC-PLOS-net model.

4.2 Rank-Distance Correlation

Fig. 1 shows the rank-distance correlation (rdc) plots, along with the Spearman correlation coefficients (R) and significance levels (p -values). This is shown for six exemplary landscapes with conflicting objectives ($\rho = -0.4$) and increasing ruggedness $k \in \{0, 2, 4\}$. The top plots report measures for landscape with $m = 2$ objectives, whereas the bottom plots report results for $m = 3$ objectives. The regression lines with 95% confidence regions are also displayed. The horizontal

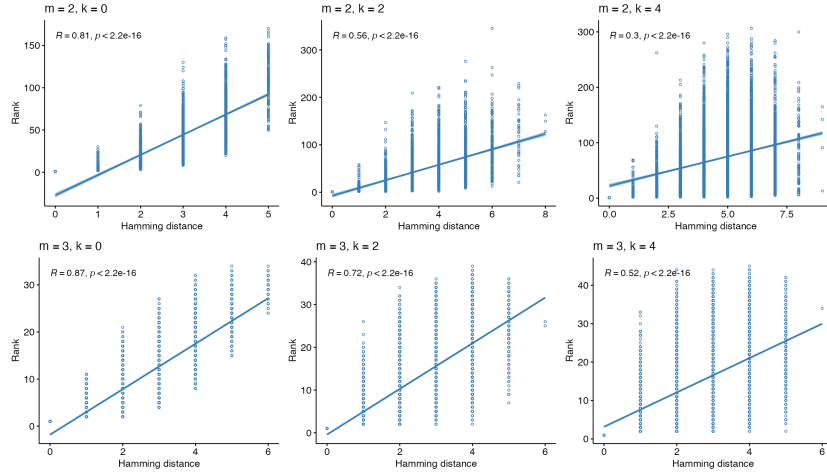


Fig. 1: Rank-distance correlation for landscapes with conflicting objectives ($\rho = -0.4$), increasing ruggedness $k \in \{0, 2, 4\}$, and two values of $m \in \{2, 3\}$. The x -axis represents the Hamming distance to the closest Pareto optimal solution. The Spearman correlation coefficient and its corresponding p -value are also displayed.

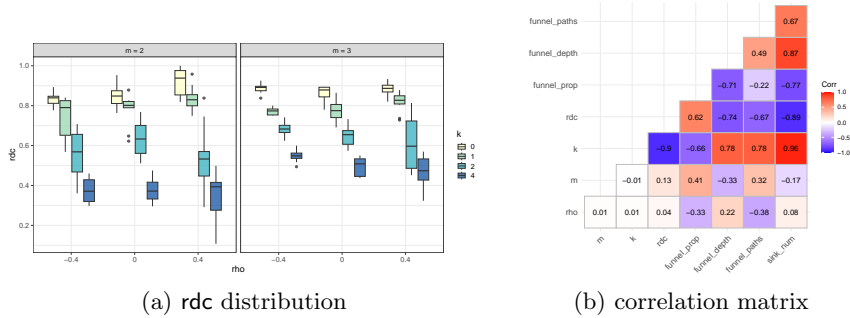


Fig. 2: Funnel metrics and correlation with benchmark parameters.

axis is for the Hamming distance between all PLO solutions and their closest Pareto optimal solution. The vertical axis is for the rank of each PLO solution. From these plots, we observe a high positive correlation between distance and rank values for smooth landscapes ($k = 0$, left plots) suggesting a rank gradient towards Pareto optimal solutions. As the ruggedness increases ($k \in \{2, 4\}$, middle and right plots), the correlation reduces but remains moderately positive. The plots reveal that solutions at the same distance to Pareto optimal solutions have a wide range of rank values (vertical set of points with the same x -coordinate). This is particularly noticeable for the largest value of $k = 4$.

The distribution of the rdc metric for the entire set of benchmark parameters can be seen in Fig. 2a. Overall, rdc values are moderate to high, suggesting that the landscapes are not overly difficult to search. They seem to have a global structure where Pareto local optima are clustered around global optima. The rdc coefficient decreases with k in all cases, which is consistent with the increasing search difficulty for rugged landscapes. The impact of the two other instance parameters, ρ and m , is less significant. The rdc values remain fairly consistent across different ρ values and generally tend to be higher for $m = 3$. The distribution of rdc values is also more compact for $m = 3$. In most cases, the coefficient is above 0.2 for $m = 2$ and above 0.4 for $m = 3$.

Fig. 2b shows the correlations amongst all the proposed funnel metrics and the benchmark instance parameters. The rdc metric highly correlates with the ruggedness parameter k , indicating that it is a good predictor of search difficulty. rdc also shows a strong correlation with the other network-based funnel metrics, particularly with the number of sub-optimal sinks. This is noteworthy as rdc is not a network metric, and still captures aspects of the global landscape connectivity structure. A more comprehensive discussion of the network-based metrics is presented in Section 4.4.

4.3 Network Visualisation

Visualising fitness landscapes brings an intuitive support in understanding search difficulty. Node-edge diagrams are commonly used for network visualisation, where shapes (circles, squares, etc.) represent nodes and lines or curves represent edges. Attributes like colour, shape and width can highlight relevant features of nodes and edges. Another key aspect of network visualisation is the graph layout, which refers to the positioning of nodes and edges in the 2D plane.

To visualise the MC-PLOS-net models, we identify relevant features of nodes and edges that relate to the landscape’s funnel structure. Specifically, we differentiate between four types of nodes: (i) **pos**: Pareto optimal solutions, (ii) **sinks**: sub-optimal nodes without outgoing edges, (iii) **sources**: nodes without incoming edges, and (iv) **standard**: nodes that do not fit into categories (i) to (iii). Node and edge colours indicate whether they belong to the Pareto optimal funnel(s), that is, whether they belong to pathways converging to Pareto optimal solutions. The size of nodes is proportional to their incoming weighted degree (also known as incoming strength). Therefore, it indicates the extent to which nodes act as attractors in the search process. Similarly, the intensity of an edge’s colour reflects its enumeration frequency.

Regarding the graph layout, we explore three alternatives: (i) for problems with two objectives ($m=2$) a natural layout is to place nodes using their objective values (f_1, f_2) as coordinates (x, y) ; (ii) a force-directed layout (*stress majorisation* [29]), and (iii) keeping the x -coordinates suggested by the stress majorisation layout and embedding the node ranks as the y -coordinate. Force-directed layouts are based on a physical analogy where nodes are charged particles joined by strings. These algorithms strive to distribute vertices evenly in space, maintain approximately uniform edge lengths, and minimise edge crossings. Nodes which share more edges are closer to each other, providing an intuitive view of the graph connectivity. These algorithms typically involve a stochastic minimisation process. Stress majorisation [29] adapts an optimisation function from multidimensional scaling. It is a reliable choice for our purposes as it is deterministic and suitable for networks with multiple components.

Fig. 3 illustrates these three layouts (Objectives, Force-directed and Rank) for four example landscapes with $m = 2$, $\rho = -0.4$, and $k \in \{0, 1, 2, 4\}$. The plots reveal that the smooth landscape ($k = 0$, top row) contains a single Pareto optimal funnel, with all pathways leading to the single compressed node with rank = 1. In contrast, the rugged landscapes ($k > 0$, rows 2 to 4) show an increasing number of separated connected components and nodes (visualised in blue) which are not in pathways leading to Pareto local optimal nodes. Instead, these blue nodes are in pathways leading to sub-optimal sink nodes, depicted as triangles with a darker outline. The number of these sink nodes (triangles) drastically increases with the ruggedness parameter k . These plots confirm that the notion of funnel, using the monotonic sequences definition, does not correspond to connected components in the graph. Rather, pathways that do not lead to Pareto optimal solutions may belong to the connected component containing the Pareto optimal node(s).

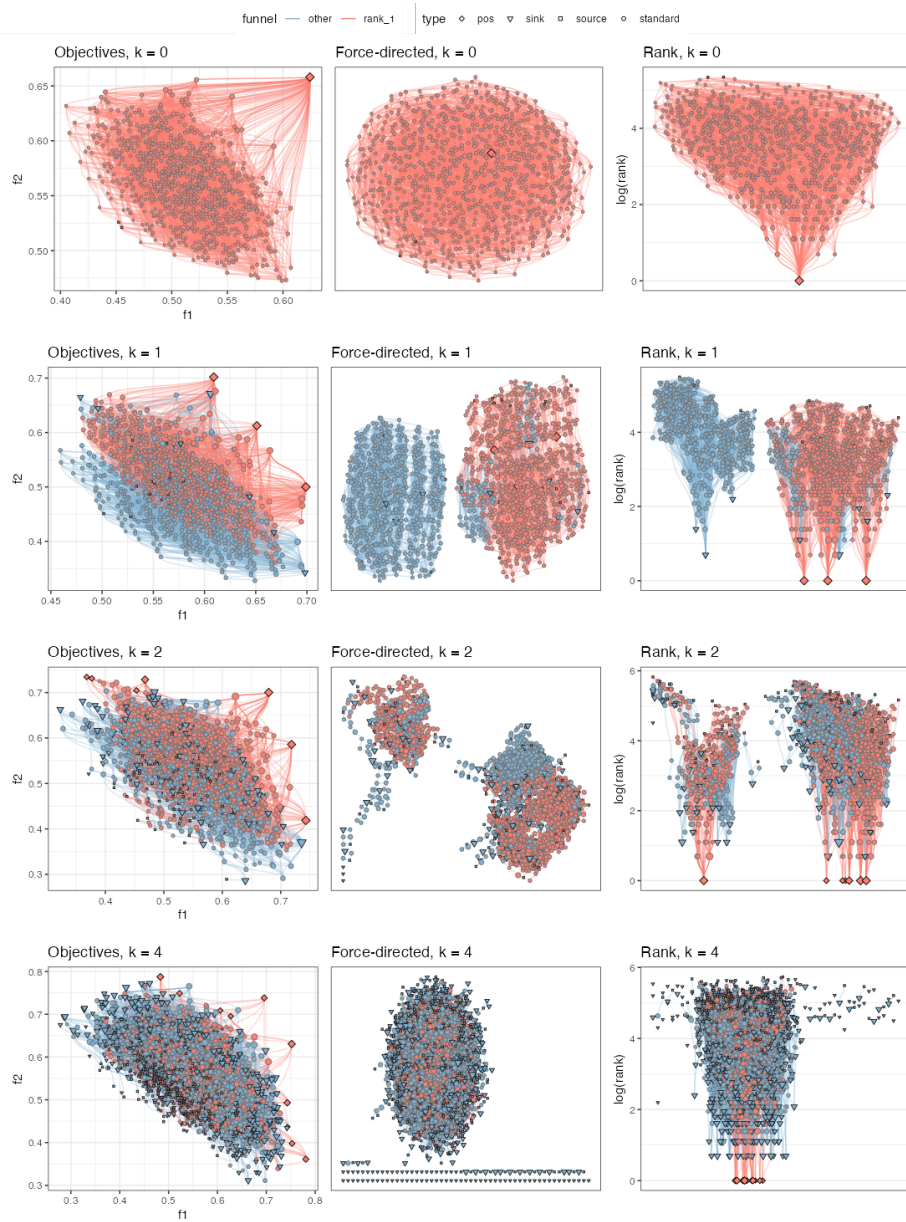


Fig. 3: Node-edge diagrams with three alternative graph layouts for four exemplary MC-PLOS-nets with $\rho = -0.4$, $m = 2$, and $k \in \{0, 1, 2, 4\}$. The smooth landscape ($k = 0$, top row) presents a single funnel that converges to Pareto optimal nodes (rank = 1). By contrast, the rugged landscapes $k \in \{1, 2, 4\}$ reveal secondary funnels and an increasing number of connected components and sink nodes.

We argue that the three considered layouts, in conjunction with the network decorations (Fig. 3), offer alternative views of the same complex networks. The objectives layout looks familiar as it is consistent with traditional Pareto front visualisations of bi-objective problems. However, it provides additional insights into the connectedness of PLO solutions as well as their grouping by funnel membership. An information that standard visualisations might overlook. The force-directed layout makes the most of the space to reveal network connectivity patterns and separate connected components. Lastly, the rank layout makes the optimisation process more tangible by showing the decreasing pathways of connected nodes converging towards optimal or sub-optimal sinks.

Fig. 3 shows MC-PLOS-nets with conflicting objectives only ($\rho = -0.4$). To examine the effect of varying the correlation of objectives on the landscape global structure, Fig. 4 reports the node-edge diagrams for networks with $\rho \in \{-0.4, 0.0, 0.4\}$, $m = 2$, and $k = 1$. We chose the objectives layout to convey the networks. The plots show a drastic reduction in both the number of nodes and the proportion of nodes in the Pareto optimal funnel when the objectives correlation increases, even though the landscapes have the same ruggedness level ($k = 1$). The connectivity between nodes (number of edges) is also markedly reduced with an increase in objectives correlation, leading to the emergence of many small-sized connected components. Figs 3 and 4 illustrate that MC-PLOS-nets capture the funnel structure of the landscape and reveal the underlying benchmark parameters. The following section explores the distribution of the network metrics and their correlation with benchmark parameters.




Fig. 4: Node-edge diagrams with the objectives graph layouts for three exemplary MC-PLOS-nets with increasing objectives correlation $\rho \in \{-0.4, 0.0, 0.4\}$, $m = 2$, and $k = 1$. The number of nodes and the proportion of nodes in the Pareto optimal funnel clearly decrease with the objectives correlation.

4.4 Network Metrics

The distributions of all the network-based funnel metrics in relation to benchmark parameters ρ , m and k are reported in Fig. 5. All metrics appear to capture




Fig. 5: Distribution of network metrics with respect to benchmark parameters.

the benchmark parameters, particularly the ruggedness k . For $k = 0$, all nodes consistently belong to the Pareto optimal funnel (i.e. `funnel_prop` equals one). `funnel_prop` then decreases as k increases, nearing zero for $k = 4$ and correlated objectives $\rho = 0.4$. `funnel_prop` decreases with the amount of correlation among the objectives and is substantially higher for 3 objectives ($m = 3$), especially for conflicting and uncorrelated objectives. The correlation matrix from Fig. 2b reveals that `funnel_prop` moderately correlates with the benchmark parameters, positively with m and negatively with ρ and k .

The depth of the Pareto optimal funnel (`funnel_depth`) measures the maximum weighted path distance between source nodes (nodes without incoming edges) and Pareto optimal nodes. As k increases, the depth of funnels also increases, as shown in the correlation matrix (Fig. 2b). For $m = 2$, `funnel_depth` values decrease with increasing objectives correlation. However, for $m = 3$, this trend seems to reverse. A striking observation is the shallow funnels for $m = 3$ and conflicting objectives $\rho = -0.4$.

The total number of alternative pathways from source nodes to Pareto optimal nodes is captured by the `funnel_paths` metric. This metric strongly correlates with the ruggedness parameter k , as evidenced by the log scale on the y -axis and the correlation coefficient from Fig. 2b. The number of paths decreases with increasing objectives correlation, and is higher for $m = 3$ than for $m = 2$.

Finally, the number of sub-optimal sinks (`sink_num`) strongly correlates with k , as revealed by both the logarithmic y -scale in Fig. 5 and the correlation coefficient (0.96, see Fig. 2b). The number of sinks is always zero for $k = 0$ and then rapidly increases for larger k values. Conversely, `sink_num` does not show any correlation with the other two benchmark parameters ρ and m .

5 Funnel Features vs. Search Performance

This section examines the impact of funnel features on both algorithm performance and algorithm selection.

5.1 Algorithms and Parameter Settings

We consider the following multi-objective algorithms, commonly applied to ρ mnk-landscapes [7, 8].

PLS: Pareto local search (PLS) [20] is a multi-objective local search method that maintains an unbounded archive A of mutually non-dominated solutions. This archive is initialised with a randomly chosen solution. In every iteration, a solution is randomly picked from the archive $x \in A$, and its neighbours are evaluated. For ρ mnk-landscapes, the neighbours $\mathcal{N}(x)$ are solutions that are a Hamming distance 1 away from x . Solutions that are dominated are filtered out, while non-dominated solutions from $A \cup \mathcal{N}(x)$ are saved in the archive. The current solution x is then tagged as *visited* to avoid re-evaluating its neighbours in future iterations. The entire process naturally stops once all solutions in the archive are marked as *visited*.

G-SEMO: The global simple evolutionary multi-objective optimiser (G-SEMO) is an elitist, steady-state, multi-objective evolutionary algorithm [30]. Like PLS, G-SEMO maintains an unbounded archive A and selects one solution $x \in A$ randomly in each iteration. It then generates a single offspring x' , where each binary variable of x is independently flipped at a rate of $1/n$. The archive is then updated with non-dominated solutions from $A \cup \{x'\}$. This process repeats until a stopping condition is met. Unlike PLS, which explores the entire 1-bit-flip neighbourhood, G-SEMO uses a standard stochastic bit-flip mutation so that there is a non-zero probability of reaching any solution from the search space.

NSGA-II: The non-dominated sorting genetic algorithm II (NSGA-II) is a dominance-based multi-objective evolutionary algorithm [24]. At a given iteration t , the current population P_t combines with its offspring Q_t , and is divided into non-dominated fronts $\{F_1, F_2, \dots\}$ according to the non-dominated sorting procedure [23]. Offspring generation follows a standard setting with uniform crossover and stochastic bit-flip mutation. The front that a solution belongs to determines its rank within the population — this corresponds to the ranking used as a node attribute in C-PLOS-nets. Selection is based on rank-values, with crowding distance serving as a tie breaker. Survival selection involves populating P_{t+1} with the best-ranked solutions. We record all non-dominated solutions found by NSGA-II and return them at the end of the search process.

Each algorithm is run independently 30 times per instance. PLS terminates when it naturally reaches a Pareto local optimum set [20]. We set the stopping condition of G-SEMO and NSGA-II at 10 000 evaluations. NSGA-II uses a population




Fig. 6: Spearman correlation between features and algorithm performance.

of size 100. For every algorithm, we evaluate the quality of all non-dominated solutions found in each trial. Algorithm performance is measured as the proportion of identified Pareto optimal solutions, referred to as the *Pareto resolution* (*reso*). We also report the hypervolume [31], and more specifically, the *relative hypervolume* (*hv*) covered by the final archive with respect to the exact Pareto front. A higher *hv* value is better, and $hv = 1$ means that the exact Pareto front was identified. The hypervolume reference point is set to the origin.

5.2 Experimental Results and Discussion

We begin by measuring how the benchmark parameters and funnel features correlate with algorithm performance in Fig. 6. Specifically, we look at the number of evaluations (*eval*) for PLS, as well as the Pareto resolution (*reso*) and relative hypervolume (*hv*) achieved by PLS, G-SEMO, and NSGA-II. Performance measures are to be maximised, so a positive correlation indicates a positive effect on approximation quality. The only exception is for *eval*, where a positive correlation would actually indicate a negative impact on runtime.

The runtime of PLS increases with the proportion of funnels (*funnel_prop*) — as it does with the number of objectives and their degree of conflict. This result on a positive impact on both *reso* and *hv*, similar to the effect of correlated ranks and distances (*rdc*). By contrast, the depth of funnels (*funnel_depth*) and the number of sinks (*sink_num*) negatively affect approximation quality. The impact of *rdc* and *sink_num* is consistent across algorithms, while the influence of the number of pathways from the source to Pareto optimal nodes (*funnel_paths*) is more noticeable for NSGA-II and G-SEMO.

Let us now study the collective impact of funnel features on algorithm performance. Using the benchmark parameters and funnel features as predictors, we construct a regression model for predicting the approximation quality of the different algorithms and the runtime of PLS. Our model is based on the well-established random forests [32, 33] with default parameters, using all considered 240 ρ mnk-landscape instances for training. We conduct 30 independent runs of random forests for each performance measure, and we report average values below. The variance explained by the model (R^2) and the relative importance of predictors according to random forests are provided in Fig. 7. The measure




Fig. 7: Relative importance of features in predicting algorithm performance.

of importance is the standard mean decrease in prediction accuracy [32, 33]: A higher value indicates a more important predictor.

The variance explained by the model is consistently higher than 0.65. This suggests that the prediction model accounts for over 65% of the variance in predicted values across all problems. This is comparable to previous work [22], despite we here consider significantly fewer features. Interestingly, each funnel feature is among the most important ones for at least one performance measure: `rdc` for the hypervolume of PLS, `funnel_prop` for all PLS performance measures, `funnel_depth` for PLS and the resolution of NSGA-II, `funnel_paths` for NSGA-II and G-SEMO, and `sink_num` for all measures except the runtime of PLS and the resolution of NSGA-II.

Our study concludes with a simple CART decision tree [34, 35] for algorithm selection, depicted in Fig. 8. The tree recommends the best algorithm for a given problem, using benchmark parameters and funnel features as predictors. The algorithm with the best average hypervolume among PLS, G-SEMO, and NSGA-II is considered the correct class for a given instance. Therefore, this decision tree addresses a classification task with the three considered algorithms as classes. Out of the 240 ρ mnk-landscape instances, 53 were omitted because no algorithm was superior, leaving 187 instances for training. The numbers beneath each tree node indicate the instances where G-SEMO, NSGA-II, and PLS are the best, respectively, followed by the proportion of instances covered by the node. The cross-validated classification accuracy is 88%, slightly higher than [22]. This is significantly better than always selecting NSGA-II, which is superior in 44% of instances from this dataset.

Interestingly, the predictors that appear in the tree, deemed as the most important for algorithm selection, are all funnel features. This implies that funnel features provide valuable information to make an informed decision on which algorithm to choose from the portfolio. We observe that the classifier recommends PLS over NSGA-II when more than 98% of nodes are in the Pareto optimal funnel. This makes sense, given that there is no need to escape from local optima when the vast majority of those lead to the Pareto set. Among instances with less than 98% of nodes in the Pareto optimal funnel, NSGA-II clearly outperforms G-SEMO when the number of sub-optimal sinks is relatively large (20 or more), highlighting the benefit of crossover in such cases.




Fig. 8: CART decision tree for algorithm selection (classification accuracy: 0.88). The numbers beneath each tree node indicate the instances where G-SEMO, NSGA-II, and PLS are the best, respectively, followed by the proportion of instances covered by the node.

6 Conclusions

We proposed a characterisation of funnels in combinatorial multi-objective landscapes based on the notion of solution ranks (layers of non-dominated solutions). We adapted the PLOS-net landscape model with two main changes: node-compression and monotonic (rank-decreasing) edges. We named this new model monotonic compressed PLOS-net (MC-PLOS-net). We then proposed visualisations and network metrics that characterise the landscape’s global structure. Additionally, we proposed the rank-distance correlation metric, *rdc*, which uses the set of Pareto local and global optimal solutions. All the proposed funnel metrics proved to accurately capture the landscape’s global structure, to correlate with benchmark parameters, and to explain the performance of well-established multi-objective local search and evolutionary algorithms.

Our empirical results were derived from small, fully enumerated combinatorial landscapes. In future work, we plan to consider continuous optimisation, as well as scale the proposed model to larger problems. The rank-distance correlation metric can be generalised to larger landscapes by sampling Pareto local optimal solutions and the best known solutions. Similarly, the proposed funnel network features could be extracted from approximated MC-PLOS-net models.

We also plan to integrate the newly-proposed features with the set of existing multi-objective landscape features in order to conduct a more comprehensive study on performance prediction and algorithm selection for multi-objective optimisation.

References

1. Pitzer, E., Affenzeller, M.: A Comprehensive Survey on Fitness Landscape Analysis, pp. 161–191. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
2. Richter, H., Engelbrecht, A. (eds.): Recent Advances in the Theory and Application of Fitness Landscapes. Emergence, Complexity and Computation, Springer (2014)
3. Malan, K.M.: A survey of advances in landscape analysis for optimisation. *Algorithms* 14(2), 40 (2021)
4. Knowles, J., Corne, D.: Towards landscape analyses to inform the design of a hybrid local search for the multiobjective quadratic assignment problem. In: *Soft Computing Systems*. vol. 2002, pp. 271–279 (2002)
5. Garrett, D., Dasgupta, D.: Multiobjective landscape analysis and the generalized assignment problem. In: Maniezzo, V., Battiti, R., Watson, J.P. (eds.) *Learning and Intelligent Optimization*. pp. 110–124. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
6. Kerschke, P., Wang, H., Preuss, M., Grimme, C., Deutz, A.H., Trautmann, H., Emmerich, M.: Towards analyzing multimodality of continuous multiobjective landscapes. In: *PPSN, LNCS*, vol. 9921, pp. 962–972. Springer (2016)
7. Daolio, F., Liefoghe, A., Verel, S., Aguirre, H.E., Tanaka, K.: Problem features versus algorithm performance on rugged multiobjective combinatorial fitness landscapes. *Evolutionary Computation* 25(4) (2017)
8. Liefoghe, A., Daolio, F., Verel, S., Derbel, B., Aguirre, H., Tanaka, K.: Landscape-aware performance prediction for evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation* 24(6), 1063–1077 (2020)
9. Wales, D.J.: Energy landscapes and properties of biomolecules. *Physical Biology* 2(4), S86 (nov 2005), <https://dx.doi.org/10.1088/1478-3975/2/4/S02>
10. Doye, J.P.K., Miller, M.A., Wales, D.J.: The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *Journal of Chemical Physics* 110(14), 6896–6906 (1999)
11. Lunacek, M., Whitley, D.: The dispersion metric and the cma evolution strategy. In: *Genetic and Evolutionary Computation Conference (GECCO)*. p. 477–484. Association for Computing Machinery, New York, NY, USA (2006)
12. Preuss, M.: Improved topological niching for real-valued global optimization. In: *Applications of Evolutionary Computation*. pp. 386–395. Springer (2012)
13. Kerschke, P., Preuss, M., Wessing, S., Trautmann, H.: Detecting funnel structures by means of exploratory landscape analysis. In: *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. pp. 265–272. GECCO '15, ACM, New York, NY, USA (2015)
14. Boese, K.D., Kahng, A.B., Muddu, S.: A new adaptive multi-start technique for combinatorial global optimizations. *Operations Research Letters* 16(2), 101–113 (1994)
15. Hains, D.R., Whitley, L.D., Howe, A.E.: Revisiting the big valley search space structure in the TSP. *Journal of the Operational Research Society* 62(2), 305–312 (2011)
16. Ochoa, G., Veerapen, N.: Deconstructing the big valley search space hypothesis. In: *Evolutionary Computation in Combinatorial Optimization, EvoCOP*. Lecture Notes in Computer Science, vol. 9595, pp. 58–73. Springer (2016)
17. Ochoa, G., Tomassini, M., Verel, S., Darabos, C.: A study of nk landscapes' basins and local optima networks. In: *Genetic and Evolutionary Computation Conference, GECCO 2008*. pp. 555–562. ACM, New York, NY (2008)

18. Ochoa, G., Veerapen, N., Daolio, F., Tomassini, M.: Understanding phase transitions with local optima networks: Number partitioning as a case study. In: *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2017. Lecture Notes in Computer Science*, vol. 10197, pp. 233–248 (2017)
19. Berry, R.S., Kunz, R.E.: Topography and dynamics of multidimensional interatomic potential surfaces. *Phys. Rev. Lett.* 74, 3951–3954 (1995)
20. Paquete, L., Schiavinotto, T., Stützle, T.: On local optima in multiobjective combinatorial optimization problems. *Annals of Operations Research* 156(1), 83–97 (2007)
21. Liefoghe, A., Derbel, B., Verel, S., López-Ibáñez, M., Aguirre, H., Tanaka, K.: On pareto local optimal solutions networks. In: *Parallel Problem Solving from Nature, PPSN XV*. pp. 232–244. Springer International Publishing, Cham (2018)
22. Liefoghe, A., Ochoa, G., Vérel, S., Derbel, B.: Pareto local optimal solutions networks with compression, enhanced visualization and expressiveness. In: Silva, S., Paquete, L. (eds.) *Genetic and Evolutionary Computation Conference, GECCO*. pp. 713–721. ACM (2023), <https://doi.org/10.1145/3583131.3590474>
23. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Boston, MA, USA (1989)
24. 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)
25. Jones, T., Forrest, S.: Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In: *International Conference on Genetic Algorithms*. pp. 184–192. Morgan Kaufmann (1995)
26. Knowles, J., Corne, D.: Instance generators and test suites for the multiobjective quadratic assignment problem. In: *Proceedings of the Evolutionary Multi-Criterion Optimization Conference (EMO 2003)*. pp. 295–310. No. 2632 in LNCS, Springer (2003)
27. Garrett, D., Dasgupta, D.: Multiobjective landscape analysis and the generalized assignment problem. In: *Learning and Intelligent OptimizatioN (LION 2). Lecture Notes in Computer Science*, vol. 5313, pp. 110–124. Springer, Trento, Italy (2007)
28. Verel, S., Liefoghe, A., Jourdan, L., Dhaenens, C.: On the structure of multiobjective combinatorial search space: MNK-landscapes with correlated objectives. *European Journal of Operational Research* 227(2), 331–342 (2013)
29. Gansner, E.R., Koren, Y., North, S.: Graph drawing by stress majorization. In: Pach, J. (ed.) *Graph Drawing*. pp. 239–250. Springer Berlin Heidelberg, Berlin, Heidelberg (2005)
30. Laumanns, M., Thiele, L., Zitzler, E.: Running time analysis of evolutionary algorithms on a simplified multiobjective knapsack problem. *Natural Computing* 3(1), 37–51 (2004)
31. Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., Da Fonseca, V.G.: Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation* 7(2), 117–132 (2003)
32. Breiman, L.: Random forests. *Machine Learning* 45(1), 5–32 (2001)
33. Liaw, A., Wiener, M.: Classification and regression by randomforest. *R News* 2(3), 18–22 (2002)
34. Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A.: *Classification and regression trees*. Taylor & Francis, Andover, UK (1984)
35. Therneau, T., Atkinson, B.: *rpart: Recursive Partitioning and Regression Trees* (2023), <https://CRAN.R-project.org/package=rpart>, r package version 4.1.23