



The fractal geometry of fitness landscapes at the local optima level

Sarah L. Thomson¹ · Gabriela Ochoa¹ · Sébastien Verel²

Accepted: 3 December 2020 / Published online: 19 December 2020
© The Author(s) 2020

Abstract

A *local optima network* (LON) encodes local optima connectivity in the fitness landscape of a combinatorial optimisation problem. Recently, LONs have been studied for their *fractal dimension*. Fractal dimension is a complexity index where a non-integer dimension can be assigned to a pattern. This paper investigates the fractal nature of LONs and how that nature relates to metaheuristic performance on the underlying problem. We use visual analysis, correlation analysis, and machine learning techniques to demonstrate that relationships exist and that fractal features of LONs can contribute to explaining and predicting algorithm performance. The results show that the extent of multifractality and high fractal dimensions in the LON can contribute in this way when placed in regression models with other predictors. Features are also individually correlated with search performance, and visual analysis of LONs shows insight into this relationship.

Keywords Fitness landscapes · Fractal analysis · Local optima networks

1 Introduction

Fractals are patterns which contain parts resembling the whole (Mandelbrot 1972). Under this definition fractals are ubiquitous in the complex simplicity of nature, from microscopic blood vessel networks to the macroscopic pattern of the rings of Saturn. Nature and evolution seem to favour fractal design: using a pattern repeatedly allows replicability with very few instructions. The *fractal dimension* (Mandelbrot 1975) is a complexity index capturing how the detail in a pattern changes when one views it using a different resolution or scale. Fractal dimension analysis has been used in diagnostic imaging (detecting colon cancer (Esgiar et al. 2002); characterising images in mammography (Caldwell et al. 1990); characterising

leukaemia cells (Mashiah et al. 2008), search and rescue [analysing the layout of victim location after building collapses (Saeedi and Sorensen, 2009)], and in engineering for the design of antenna (Werner and Ganguly 2003), among innumerable others. Fractal geometry can also facilitate vast amounts of information being embedded in a comparatively small space. Indeed, branching structures inside human lungs fill space in a fractal way; because of this, the equivalent surface area of a tennis court is compacted within the volume of the lungs. The fractal dimension of human lungs has been measured at approximately 2.88 (Uahabi and Atounti 2017), which indicates high spatial complexity and convolution.

Fitness landscapes of some combinatorial optimisation problems have been viewed under a fractal lens (Weinberger and Stadler 1993). Fitness landscapes are both a lucid metaphor and a mathematical object; they contain the set of solutions to an optimisation problem, the fitnesses of those solutions (these can be visualised as the heights), and a function for measuring adjacency between solutions. The study of fitness landscape architecture provides insight about reactions between metaheuristic algorithms and problems. This can serve as a springboard for more informed algorithm design or selection.

The first study to conduct fractal analysis on fitness landscapes (Weinberger and Stadler 1993) stipulated that for certain problems, landscape ruggedness scales at

✉ Sarah L. Thomson
s.l.thomson@stir.ac.uk

Gabriela Ochoa
gabriela.ochoa@cs.stir.ac.uk

Sébastien Verel
verel@univ-littoral.fr

¹ Computing Science and Mathematics, University of Stirling, Stirling FK9 4LA, UK

² Laboratoire d'Informatique Signal et Image de la Côte d'Opale, Univ. Littoral Côte d'Opale, F-62100 Calais, France

different levels of abstraction and that this indicated fractal structure. Subsequent studies have reported similar findings (Zelinka et al. 2014; Locatelli 2005; Richter 2018) and some have emphasised the potential lying dormant in the largely untapped field of fractal analysis for landscapes.

A *local optima network* (LON) (Ochoa et al. 2008) models local optima and their connectivity in a fitness landscape. That is, the nodes are local optima and the edges are metaheuristic search transitions between two local optima under a chosen search operation. There is a significant body of evidence suggesting that features of LONs can correlate to, explain, or predict metaheuristic algorithm performance on the underlying combinatorial problem (Daolio et al. 2010; Daolio et al. 2011; Verel et al. 2011; Herrmann et al. 2016; Ochoa and Veerapen 2018; Ochoa et al. 2017).

Little is known about the fractal complexity in LONs and how their fractal nature relates to metaheuristic algorithm performance. Preliminary work has indicated that the *fractal dimension* might have a connection to search (Thomson et al. 2018a; Thomson et al. 2018). That being said, the latter study considers only small problem instances (size $N = 18$ for a binary-encoded problem, NK Landscapes). The first study mentioned is on the quadratic assignment problem (QAP) and they consider some benchmark instances from QAPLIB (Burkard et al. 1997) up to $N = 28$ (Thomson et al. 2018a), although only two of the library's several instance classes for this problem size range are included; consequently, the fractal analysis is conducted on only 25 QAPLIB instances.

We intend to illuminate understanding of the relationships between fractal geometry in LONs and metaheuristic algorithm performance. The QAP serves as a testbed for the analysis and we use QAPLIB instances, increasing the number of instances considered threefold when compared to previous work (Thomson et al. 2018a) and raising the maximum problem size from 28 to 50. A recent and refined LON construction algorithm (Ochoa and Herrmann 2018) is used to intelligently build LONs for the QAPLIB instances. Features of the LONs, including fractal dimension features, are computed and the parallel between them and performance is investigated using visual tools, correlation analysis, and linear and random forest regression models.

The contributions of this article can summarised as follows:

1. We bring new insight into how multifractal geometry at the local optima level can help explain and predict algorithm performance
2. A significant expansion of the data-set used for fractal analysis in LONs (using more than 3x the previous number of QAPLIB instances and raising $N \leq 28$ to N

≤ 50 , as well as deploying a recent refined and tested sampling algorithm for constructing the LONs)

3. Enhanced statistical techniques for properly validating the use of LON fractal analysis for algorithm explanation and prediction (random forest to model non-linearities; random repeated subsampling cross-validation; using intelligible predictors such as the extent of multifractality and the median fractal dimension).

The article is structured as follows: Sect. 2 contains the necessary background information to render this article self-contained; Sect. 3 details aspects of the methodology used; Sect. 4 gives the experimental setup, with Sect. 5 presenting the results; finally, Sect. 6 finishes the article with conclusions and directions for future work.

2 Preliminaries

2.1 Fitness landscapes

A fitness landscape (Stadler 2002) is composed of three parts, (S, N, f) : S is the full solution set; $N : S \rightarrow 2^S$ is known as the neighbourhood function and assigns a set of adjacent solutions $N(s)$ to every $s \in S$; and f is a fitness function $f : S \rightarrow \mathbb{R}$ that provides a mapping from solution to associated fitness. That fitness can be conceptualised as the solution *height* within the landscape metaphor.

The analysis of fitness landscape objects can provide an intense understanding of optimisation problems and their reactions with metaheuristic algorithms (Pitzer and Affenzeller 2012). Indeed, landscapes have been used to facilitate algorithm selection (Hoos et al. 2004), operator selection (Merz and Freisleben 2000), and parameter tuning (Hutter et al. 2007).

2.2 Local optima networks

The *local optima network* (LON) model (Ochoa et al. 2008) was introduced as a tool for studying the connectivity of local optima in a fitness landscape, and has subsequently shown proficiency in helping with explaining metaheuristic search dynamics (Chicano et al. 2012; Herrmann et al. 2016; McMenemy et al. 2018). We define the components of a LON before describing the model as a whole.

Nodes. The set of nodes LO are the local optima, meaning that a node lo_i has superior fitness with respect to the entire neighbourhood. Formally: $\forall n \in N(lo_i) : f(lo_i) \leq f(n)$ (assuming minimisation) where $N(lo_i)$ is the neighbourhood and n is a neighbour.

Edges. An edge is delineated between two nodes if the probability of "escape" from the source local optimum to the destination is greater than zero. The "escape" is defined with respect to a chosen search operation (or sequence of operations). The edge is weighted with the probability as w_{ij} . Formally local optima lo_i and lo_j comprise the source and destination of an edge respectively iff $w_{ij} > 0$. In this work, sampling is used; as a result, nodes are not necessarily associated with their complete set of potential edges.

Local optima network (LON). A local optima network, $LON = (LO, E)$, consists of nodes $lo_i \in LO$ which are the local optima, and edges $e_{ij} \in E$ between pairs of nodes lo_i and lo_j with weight w_{ij} iff $w_{ij} > 0$. We remark here that w_{ij} may be different than w_{ji} ; it follows that two weights are needed and that a LON is an oriented and weighted graph.

2.3 Fractal dimension

The notion of a *fractal dimension* for patterns was conceived by Mandelbrot (1975) and is defined as a complexity index which captures how the detail in the pattern changes with resolution used to measure it. The fractal dimension can be computed as the ratio between the logarithm of the detail and the logarithm of the scale used:

$$\text{fractal dimension} = \frac{\log(\text{detail})}{\log(\text{scale})} \quad (1)$$

To understand what the fractal dimension of a shape means we can begin by revisiting the familiar shapes associated with the topological dimension: a one-dimensional line; a two-dimensional square; a three-dimensional cube.

We can observe the relationship between scale and detail for a square in Fig. 1. Looking first at Fig. 1a where the length scale m used to measure is one (the length of one side of the square) the detail measured is precisely one square. Moving onto Fig. 1b we observe that a length scale of $m = \frac{1}{2}$ is used here (this is a *scaling factor* of two because the resolution is twice as fine). That results in the

measurement of four smaller copies of the larger square. The scale is two and the detail is four. Similarly, when m is one-quarter of the length of a side of the square (scaling factor of four; see Fig. 1c) this results in sixteen copies of the larger square being measured, giving a scale of four and detail of sixteen. The relationship $4^x = 16$ where x is dimension can be transformed into $\frac{\log(16)}{\log(4)}$, i.e. the ratio between detail and scale which is two in this case. The square is two-dimensional because for any *scale* the *detail* observed will be $scale^2$.

For some patterns the exponent x is not an integer but rather somewhere else on the real number line. In this case, the way detail changes with resolution cannot be captured with topological dimension. An illustrative example of this can be seen in Fig. 2 with the Sierpinski Triangle.

Figure 2a shows that when a scaling factor m of one is used we accordingly measure the complete pattern. If we increase the resolution twofold as in Fig. 2b *three* smaller copies of the large triangle are now measured. Recalling that fractal dimension can be obtained by solving for x the equation $scale^x = detail$ we observe that x is not an integer here. The equation is $2^x = 3$ which results in a fractal dimension of $x = \sim 1.585$.

Fractal dimensions can have efficacy in obtaining spatial and geometric information about real-world systems. They have been used, for example, in engineering for detecting cracks in plate structures (Hadjileontiadis and Douka 2007); in biology for characterising the tortuosity of animal trails (Dicke and Burrough 1988); and also in medicine for characterising mammographic patterns (Caldwell et al. 1990) and detecting colon cancer (Esgiar et al. 2002).

In our study we are computing fractal dimensions on LONs to obtain spatial complexity information about fitness landscapes. A widely-used method to estimate fractal dimension for a complex network is the "box counting" algorithm (Song et al. 2005). This "boxes" together nodes which are within m network edges of each other, aiming to describe the network using as few "boxes" as possible. The parameter m is the scale of measurement used and serves as

Fig. 1 The relationship between detail and scale for a square; m is the length of scale used for measurement


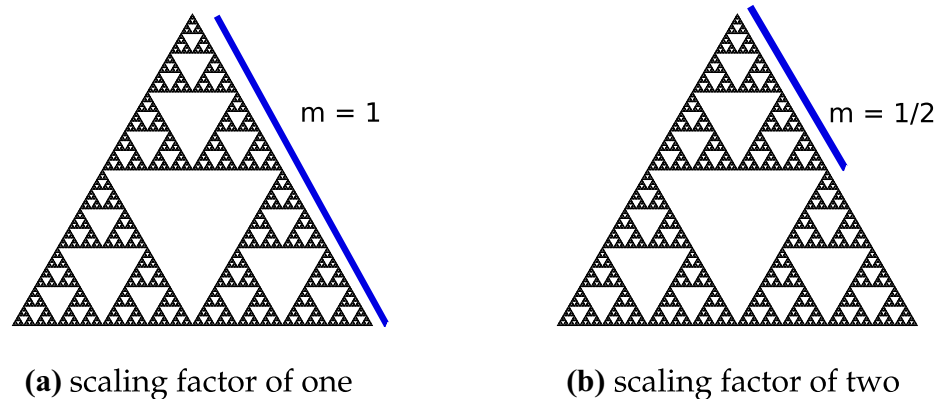


Fig. 2 The relationship between detail and scale for a fractal with topological dimension two and *fractal dimension* of ~ 1.585



the denominator in Eq. 1 to obtain fractal dimension alongside the number of “boxes” required to cover the network, which is the amount of detail observed. In stage one of the procedure, “centre” nodes are initially identified as those which are the best connected in the network. Nodes which are at a distance of no more than m edges to the centre node are then marked as “covered” and are added to the “box” associated with the centre. The process continues until all nodes are either “covered” or they are centre nodes. That means wherever a node cannot be “covered” with respect to any of the centres, it becomes a centre itself. In stage two the central distances for all nodes are calculated; this is—for each node—the closest distance to a centre. Following that, the “box” membership identity of each non-centre node is switched to that of a neighbour which is closer to a centre node. The original node is removed.

At the end of the process the number of “boxes” needed to cover the network completely, which we refer to as mb , is the number of *detail* units observed when using the resolution scale m . We can obtain the fractal dimension for the network by inserting mb and m into Eq. 1:

$$\text{fractal dimension} = \frac{\log(mb)}{\log(m)} \quad (2)$$

Fractal complexity in local optima networks has been calculated previously using box counting (Thomson et al. 2018a; Thomson et al. 2018). The box counting algorithm was altered in Thomson et al. (2018) to specialise to LONs. For two nodes to be “boxed” as a single “unit” of detail they must either be a single edge apart *or* they are within m edges of each other and they also have a fitness distance less than a set threshold ϵ .

A subsequent study proposed additional mechanisms for computing and therefore defining the fractal dimension of a LON (Thomson et al. 2018a). A box counting variant which was introduced which used LON edge weights during the process. In a LON edge weights represent the probability that a search path between the local optima will

be followed. The box counting variant used as the criteria for “boxing” that two nodes have a single edge between them which is weighted with a probability greater than β . The authors referred to values obtained using this method as *probabilistic* fractal dimensions.

In real-world complex systems a single fractal dimension can sometimes be insufficient to capture the complexity (Mandelbrot et al. 1997). Monofractal analysis such as the box counting described earlier is based on the assumption that fractal complexity is roughly uniform in the pattern. Some networks have been found to be multifractal (Song et al. 2015; Furuya and Yakubo 2011). A multifractal algorithm has been used on LONs in a prior study (Thomson et al. 2018a) and we deploy this in our experiments. The process produces a spectrum of fractal dimensions for a single pattern (LON in our case). Details and pseudo-code for the algorithm are provided later on in Sects. 3.3 and 4.3.

3 Methodology

3.1 The quadratic assignment problem

Our analysis is conducted on the much-studied quadratic assignment problem (QAP) (Lawler 1963) which is often used in fitness landscape analysis (Merz and Freisleben 2000; Merz 2004; Daolio et al. 2011; Pitzer and Affenzeller 2012; Verel et al. 2018). A QAP instance is specified with a distance matrix and a flow matrix. An entry in the distance matrix, D_{l_1, l_2} is the distance between two locations: $d(l_1, l_2)$. In the flow matrix this is the flow between two items: $f(i_1, i_2)$. Solutions are encoded as a permutation of length N , and are the allocation of N items to N locations. Fitness of a solution is the product of distances and flows between the locations and items according to the permutation and the aim is minimisation. The fitness function, g , for a solution x is then

$$g(x) = \sum_{i=1}^N \sum_{j=1}^N D_{ij} F_{ij}, \forall x \in S.$$

3.2 Constructing sampled LONs

LON sampling algorithms are generally augmented on top of an existing optimisation algorithm. We align with this trend here, opting for a recently-introduced construction algorithm which joins an ILS with LON logging for QAP (Ochoa and Herrmann 2018).

The ILS algorithm is run r times from independent random starting solutions. In the ILS process, the local optimisation is a pairwise exchange of items, with the perturbation being k pairwise exchanges. Whenever the ILS Sampling detects no improving moves from the current solution, the solution is added as a local optimum node—this is an approximation of the true structure, because the algorithm does not consider the existence of saddle points. Only improving or equal fitness local optima are accepted. It follows that local optima plateaus might be explored, although not exhaustively, during sampling. When there is a local optima plateau, this is not collapsed by default; plateaus are sometimes collapsed to facilitate the extraction of certain LON features and to assess the neutrality present. LON plateaus are collapsed in preparation for multifractal analysis in this work, and also to facilitate computation of the funnel features and the number of compressed local optima feature described in Sect. 4.5. Each local optimum encountered during search is stored in the set of nodes LO alongside its fitness, and if two optima $I1$ and $I2$ are connected by an ILS cycle (local search followed by k perturbations) during the search, an edge $e_{I1,I2}$ is stored in the LON edge-set, E . The nodes are edges logged during the r runs are joined to form a single local optima network for the problem instance. All parameters for the algorithm are stated later on in Sect. ~4.2.

3.3 Fractal analysis algorithms

As stipulated in Sect. 2.3 the standard approach for calculating and defining fractal dimension of a complex network is with a box counting algorithm (Song et al. 2005, 2006). This process iteratively "boxes" together nodes iff the distance $d(n_1, n_2) < m$, i.e. the nodes are $< m$ edges apart. The parameter m provides the scaling factor which is used to compute the fractal dimension of the network alongside the associated detail observed when using that scaling factor. Empirically the *detail* is defined as the number of "boxes" needed to completely cover the network, taken as a proportion of the network size.

We mentioned that the box counting algorithm has been specialised for the specific case of a local optima network previously (Thomson et al. 2018). In that they allowed

nodes to be "boxed" if either the distance $d(n_1, n_2) = 1$ or $d(n_1, n_2) < m$ and also $|f(n_1) - f(n_2)| < \epsilon$ where $f(n_x)$ is the fitness of node x and ϵ is the maximum fitness difference between nodes n_1 and n_2 .

3.3.1 Multifractal dimension analysis

The process for calculating multifractal dimensions is different to standard box counting and it produces a spectrum of fractal dimensions for a single pattern. One approach is called the *sandbox* algorithm (Liu et al. 2015) where several nodes are randomly selected to be sandbox "centres". Members of the sandboxes are computed as nodes are r edges apart from the centre c . After that the average sandbox size is calculated. The procedure is replicated for different values of r which is the sandbox radius. To facilitate the production of a dimension spectrum the whole process is repeated for several arbitrary real-valued numbers which supply a parameter we call q . The sandbox algorithm was specialised and modified to suit LONs in a prior study (Thomson et al. 2018a) and this is the process we use for our fractal analysis experiments. In our version of the algorithm a node n can be included in the "sandbox" of a central node c if *either* the distance $d(n, c) = 1$ or $d(n, c) = r - 1$ and $|(f(n) - f(c))| < \epsilon$. Pseudocode is given in Algorithm 1.

Algorithm 1 Algorithm for Multifractal Analysis of a LON

Input: $LON, q.values, radius.values, fitness.thresholds, number.centres$
Output: *mean sandbox size*

```

1: Initialisation:
2:  $centre.nodes \leftarrow \emptyset, noncentre.nodes \leftarrow all.nodes$ 
3:  $mean.sandbox.sizes \leftarrow \emptyset$ 
4: for  $q$  in  $q.values$  do
5:   for  $rd$  in  $radius.values$  do
6:     for  $\epsilon$  in  $fitness.thresholds$  do
7:        $centre.nodes \leftarrow RANDOM.SELECTION(all.nodes, number.centres)$ 
8:        $sandbox.sizes \leftarrow \emptyset$ 
9:       for  $c$  in  $centre.nodes$  do
10:         $number.boxed \leftarrow 0$ 
11:        for  $v$  in  $all.nodes$  do
12:           $d \leftarrow DISTANCE(c, v)$ 
13:           $j \leftarrow DIFFERENCE(f(c), f(v))$ 
14:          if  $(d == 1)$  OR  $(d == rd - 1$  and  $j < \epsilon)$  then:
15:             $number.boxed \leftarrow number.boxed + 1$ 
16:          end if
17:        end for
18:         $sandbox.sizes \leftarrow sandbox.sizes \cup \{number.boxed\}$ 
19:      end for
20:       $bs \leftarrow MEAN(sandbox.sizes)$ 
21:       $mean.sandbox.sizes[q][rd][\epsilon] \leftarrow bs$ 
22:    end for
23:  end for
24: end for

```

At the end of each "sandboxing" iteration conducted with particular values for the parameters q , r and ϵ , the associated fractal dimension is calculated:

$$\text{fractal dimension} = \frac{\log(detail^{q-1})}{(q-1) * \log(scale)} \quad (3)$$

where *detail* is the average "sandbox" size (as a proportion of the network size), q is an arbitrary real-valued value, and

scale is $\frac{r}{dm}$, with r being the radius of the boxes and dm the diameter of the network. We use this process as the foundation for obtaining multifractal dimensions for LONs in this work. In addition we separately implement a modified version of it where the metaheuristic path probabilities encoded in LON edge weights are used in the calculations. Specifically, for a node n to be a member of the sandbox with centre c there must be one of two situations: either there is a single edge between n and c (of any probability; this is to guarantee boxing momentum), or there is an edge between a direct neighbour of n and c which is weighted with a probability greater than a specified threshold β . This element was implemented with the motivation that nodes which are in close proximity to a probable path towards the central node should be included in the sandbox. We remove the fitness distance check for this algorithm variant and instead of a set of values for "sandbox" radii the sandboxes are of a fixed width, $r = 2$. The rest of the algorithm remains the same and a spectrum of fractal dimensions is produced. To differentiate the results produced by this particular algorithm variant in the following Sections we refer to the fractal dimensions obtained by this method as *probabilistic* fractal dimensions. The algorithm which does *not* include the probability constraint but instead includes "sandbox" radius variation as well as fitness distance constraints produces values which we refer to as *deterministic* fractal dimensions. The parameters for both algorithms described are stated in Sect. 4.3.

4 Experimental setup

4.1 Instances used

All instances used are from the QAPLIB benchmark library for QAP, the quadratic assignment problem library (QAPLIB) (Burkard et al. 1997). We cap the maximum problem size at 50 due to the computational expense (Liu et al. 2015) associated with multifractal analysis of large networks. It follows that further study is needed in order to confirm any findings on larger problem instances. We additionally remove the "esc" instances from the group because their LONs have very few distinct fitnesses due to large amounts of neutrality present. The resultant set consists of 85 problems, with the problem sizes ranging from 12 to 50. In all cases the global optimum is known.

The nature of QAP instances can commonly be characterised into one of four classes (Stützle 2006): uniform random distances and flows; random flows based on grids; real-world; and random "real-world like", which are not real-world but mimic distance and flow patterns seen in real-world presentations of QAP. Table 1 shows the

Table 1 The QAPLIB instances used in the experiments

Class	Instance names
Uniform random	tai{12a, 15a, 17a, 20a, 25a, 30a, 35a, 40a, 50a} rou{12, 15, 20}
Random grids	had{12, 14, 16, 18, 20} nug{12, 14, 16a-b, 17, 18, 20, 21, 22, 24, 25, 27, 30} scr{12, 15, 20} sko{42, 49} tho{30,40} wil{50}
Real-world	bur{a-h} chr{12a-c, 15a-c, 18a-b, 20a-c, 22a-b, 25a} {els19} kra{30a-b, 32} lipa{20a-b, 30a-b, 40a-b, 50a-b} ste36{a-c}
Real-world like	tai{12b, 15b, 20b, 25b, 30b, 35b, 40b, 50b}

QAPLIB instances used in the experiments and present them in these four categories. Numbers which form part of the instance names indicate the problem size, i.e. number of locations and flows and the length of a permutation solution.

4.2 Construction of local optima networks

For each QAP instance we construct a local optima network. As stipulated in Sect. 3.2, this is done by using an ILS algorithm which has been augmented with LON logging mechanisms. The LON logging amalgamates the unique nodes and edges from 200 ILS runs into a single network. Each run terminates after 10,000 iterations without an improvement. This is a deliberately lenient condition which was chosen with the motivation that ILS runs should converge to a natural stalling point. The parameter setting, however, means that some LONs (those associated with the larger problems in the instance set) are built over a number of hours; this computational cost is a limitation to our approach. We argue, however, that the benefit of the insight gained through our method outweighs the cost. The remaining ILS parameters and setup are detailed shortly in Sect. 4.4.

4.3 Fractal analysis

In contrast to traditional *monofractal* analysis, to generate *multifractal* dimensions for the LONs a range of arbitrary real-valued numbers is needed. We set these as q in the range [3.00, 8.90] in step sizes of 0.1. The number of "sandbox" centres in each iteration is set at 50 and the choice of these centres is randomised. As mentioned, in our *deterministic* multifractal algorithm variant fitness distance is considered in order to specialise to local optima

networks. The comparison between two fitness values is conducted through *logarithmic returns*:

$$\text{fitnessdifference} = \log(f_1/f_2) \quad (4)$$

where f_1 and f_2 are the fitnesses of two local optima at the start and end of a LON edge. We take the absolute value of the computed fitness difference because if $f_1 \leq f_2$, the result of Eq. 4 is negative. This value can then be compared with a set threshold, ϵ . A range of ten values is used for that algorithm: $\epsilon \in \{0.01, 0.19\}$ in step sizes of 0.02. Another essential element of *deterministic* multifractal analysis is the sizes (radii) for the sandboxes. For these we use values in the range $r \in \{2, \text{diameter} - 1\}$ where *diameter* is the LON diameter.

For the *probabilistic* multifractal algorithm variant the fitness constraint is not used and the sandboxes are of a fixed width, $r = 2$. The probability threshold parameter β must be chosen. Recall that β sets the minimum edge weight between two nodes. After preliminary runs it was noted that if β was set as greater than the minimum weight in the weights distribution then little-to-no "boxing" occurred. For this reason β is set as the minimum weight present in the distribution.

An important note. We note here that 32 out of the 85 LONs had only a single edge weight present throughout the network. Recall that the probability-based boxing process outlined in Sect. 3.3.1 stipulates that nodes can be boxed together when either there is a single edge between n and c (of any probability; this is to guarantee boxing momentum), or there is an edge between a direct neighbour of n and c which is weighted with a probability greater than a specified threshold β , which is set as the lowest weight in the network. As a consequence, when all weights are equivalent, then no boxing based on probability will occur at all—no pairs of nodes will pass the acceptance condition that their connecting edge has a weight greater than the minimum weight seen in the LON. This renders these particular networks ineligible for probabilistic fractal analysis under these conditions. Consequently, results which pertain to probabilistic dimensions consider the 53 eligible LONs and their features, while those pertaining to deterministic fractal dimensions consider all 85 LONs.

4.4 Metaheuristic performance

To obtain algorithm performance information with which to compare the LON features we use two search algorithms for the QAP. Stützle introduced iterated local search (ILS) variants for state-of-the-art performance on the QAP (Stützle 2006). We use his ILS configured as follows: first-improvement pairwise exchanges for local search; $\frac{3n}{4}$ exchanges for perturbation; accepting only improving local

optima; and terminating when the global optimum is found or after 100 iterations. Taillard's Robust Taboo Search (ROTS) (Taillard 1991) is also a competitive heuristic for the QAP. This is a best-improvement pairwise exchange local search with a variable-length tabu list tail. For each facility-location combination, the most recent point in the search when the facility was assigned to the location is retained. A potential move is deemed to be "tabu" (not allowed) if both facilities involved have been assigned to the prospective locations within the last s cycles. The value for s is changed randomly, but is always from the range $[0.9n, 1.1n]$. A run terminates when the global optimum is found, or after 100 iterations.

We run the ILS and the ROTS in these configurations on each QAPLIB instance 100 times from different starting solutions. As a measure for their performance we define the performance gap p as follows:

$$p = \frac{f(\text{alg})}{f(\text{opt})} \quad (5)$$

where $f(\text{alg})$ is the fitness obtained by the algorithm and $f(\text{opt})$ is the fitness of the global optimum. In this way, a "solved" run will output "1" and lower values are closer to the optimal fitness. For each QAP instance we report the mean p over 100 runs. In the results that follow, $p(\text{ILS})$ is this value for iterated local search and $p(\text{ROTS})$ is for robust tabu search.

4.5 LON features

Features are extracted from the local optima networks. Deterministic fractal dimension sets are calculated for each of the 85 LONs considered. In those sets there are $60 * (\text{diameter} - 2) * 10$ dimensions, where 60 is the number of (arbitrary) q values, *diameter* is the LON diameter (which differs between LONs), and 10 is the number of values for fitness distance threshold ϵ . As we recall from Sect. 4.3 53 of the 85 LONs are eligible for probabilistic fractal analysis. Those 53 have sets of probabilistic fractal dimensions calculated in addition to the deterministic ones. In each set there are 60 dimensions (one for each value of q). This is less than there was when using the *deterministic* multifractal algorithm and this is because the *probabilistic* variant does not consider parameter ranges for "sandbox" radius and does not include fitness distance calculations. The statistics we draw from the fractal complexity data are: the *minimum* fractal dimension, the *median* fractal dimension, the *maximum* fractal dimension, the *range* of fractal dimensions (calculated as the difference between the largest and smallest values), and the number of *distinct* fractal dimensions. The latter two capture the degree of multifractality present.

Also considered in the experiments are other LON features which are not related to fractal dimension values:

1. number of *local optima*;
2. *mean fitness* of sampled local optima in the LON;
3. *fitness range*;
4. fitness of *sinks* [”sinks” are nodes which have no out-going edges, i.e. the algorithm used to construct the LON became trapped there. Sinks are the bottom of funnels, which themselves are basins of attraction at the level of local optima (Ochoa et al. (2017))];
5. extent of *meta-neutrality*, which is neutrality at the local optima level, computed as $meta\text{-}neutrality = \frac{\text{number unique fitnesses}}{\text{number local optima}}$;
6. mean *out-degree*;
7. the LON *diameter*;
8. the number of compressed local optima (after connected LON nodes of the same fitness are compressed together—labelled as *comp.opt* in the Figures);
9. the correlation between the fitnesses of neighbours in the LON (*fit.fit.corr*);
10. the number of sink nodes present (*sinks*);
11. and the sub-optimal sink strength (that is, the total incoming edge weight to any sub-optimal sink nodes—*so.strength*).

4.6 Regression model setup

We build algorithm performance models using LON features for predictors and the performance of competitive metaheuristic algorithms as the response variables. The aim is clarifying how LON features can contribute to explaining or predicting algorithm proficiency, paying particular attention to the fractal nature of the LON. In pursuit of that we conduct linear and random forest regressions. The number of observations we have is relatively small—85 for the deterministic dimensions and 53 for probabilistic—so we use *random repeated subsampling cross-validation* for obtaining model statistics. This is conducted for 10,000 iterations with a training-test split of 80-20. The random forest regression uses 500 trees. Predictors are standardised (due to different value ranges) as follows: $p = \frac{(p-E(p))}{sd(p)}$, with p being the predictor in question. The model statistics we focus on are R^2 , which captures the amount of variance in the response variable which can be explained using the predictor set, and *mean squared error*, which expresses the mean squared difference between the model-estimated values and the actual values.

For the random forest models, *variable importance* rankings and values are reported. The values are calculated as the reduction in decision tree node impurities when splitting on the variable and are averaged over all 500 trees used in the regression. Node impurities are measured with the residual sum of squares.

The non-fractal LON predictors used in the models are the mean fitness; fitness range; fitness of sinks; extent of meta-neutrality; out-degree; and the number of global optima. For the deterministic fractal dimensions, we include the minimum fractal dimension and median. In the probabilistic case, these two are replaced with the fractal dimension range and number of unique dimensions.

5 Results

5.1 Distribution analysis

In Figs. 3 and 4, box-plots convey information about the fractal dimensions calculated on the local optima networks. Each box contains values for LONs associated with a particular QAPLIB instance class—those are indicated on the x -axis labels. Only a sub-set of the instance classes which are involved in the central experimentation are considered in these plots. We chose these groups because displaying their distributions alongside each other illustrates evident visual differences between these particular classes. Also provided in the Figures as accompanying text for each box is the performance of iterated local search on the QAP instances associated with those LONs; this is the performance metric $p(ILS)$.

In Fig. 3a, b the distributions concern the *median* LON fractal dimension which is associated with using the deterministic and probabilistic methodologies, respectively. In the case of the deterministic fractal analysis, this is the median value computed over all of the dimensions produced under these conditions; each dimension is the output resulting from using a different combination of the fractal analysis parameters q , r and ϵ . The probabilistic median is computed from the spectrum of dimensions associated with the range of values for q .

In both Fig 3a, b, the ”lip” class of LONs seem to have the highest values and the ”had” group have the lowest. On both plots, the highest value belongs to the ”lip” category and the lowest to ”had”. Notice that in 3b the ”lip” and the ”nug” instances—whose LONs generally have the highest fractal dimensions in this plot—also have higher values of $p(ILS)$. As stipulated in Sect. 4.4, values like these reflect that metaheuristic performance was of lower quality. With deterministic analysis, the ”lip” group have the largest variation, while the ”had” LONs have among the smallest; with probabilistic dimensions (Fig. 3b), ”had” have the




Fig. 3 Boxplots for the distributions of median fractal dimension. Each box displays values for LONs extracted from one of four QAPLIB instance classes, as indicated in the legend. The black line within the boxes is the median value

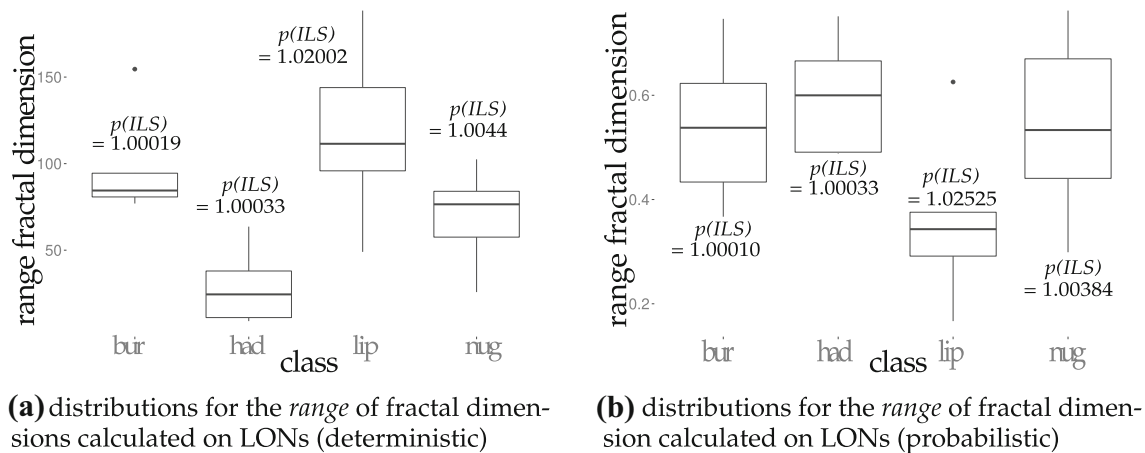


Fig. 4 Boxplots for the distributions of the *range* of fractal dimension. Here the range is the smallest dimension for that LON subtracted from the largest. Each box displays values for LONs

extracted from one of four QAPLIB instance classes, as indicated in the legend. The black line within the boxes is the median value

largest and "lip" the smallest. Deterministic fractal dimensions appear to be higher than probabilistic fractal dimensions.

Consider now the *range* of fractal dimensions in the deterministic and probabilistic spectra calculated for the LONs, which are given in Fig. 4a, b.

The range of fractal dimensions for a LON is a way to quantify the extent of multifractality present and is calculated as *maximumvalue* - *minimumvalue* with respect to the complete set of fractal dimensions produced using either the deterministic or probabilistic paradigm. Also provided is the average ILS performance, $p(ILS)$, for the QAP instances included in the classes.

Looking at the two plots and noting the different scales used for them, it seems clear that the probabilistic dimension calculation process leads to more compact ranges. This is intuitive: the conditions are stricter for measuring

"boxes" during the dimension calculation process. Let us consider in both plots the levels of the black lines (which indicate the distribution median). The "had" group has the lowest in 4a and the "lip" group has the highest. That hints that the degree of multifractality in the "lip" group is the most pronounced among the four, and it is the least pronounced in the "had" group. The previous plots told us that "lip" LONs had the highest dimensions, and "had" showed the lowest. It follows that the degree of deterministic multifractality might be associated with lower fractal dimensions. For 4b though, "had" LONs have the highest ranges of dimension and "lip" have the lowest—the opposite trend to the deterministic dimensions. With respect to algorithm performance, we can that the "lip" LONs, associated to problems with the lowest metaheuristic performance ($p(ILS)$), appear to have a higher extent of deterministic multifractality and a lower extent of

probabilistic multifractality. In Fig. 4b, the two problem groups with the best ILS performance have the widest ranges of values for dimension (i.e. amount of multifractality) of the four categories.

5.2 Visualisation

Visual analysis of LONs provides valuable insight into algorithm performance and problem structure, and can augment more empirical or statistical findings (Ochoa and Veerapen 2016). We begin with visualisation before moving onto correlation analysis (Sect. 5.3) and machine learning models (Sect. 5.4) thereafter.

Figure 5 shows two partial LONs, each for a different QAPLIB instance. Only the fittest 10% of local optima are plotted for visual clarity. Global optima are red squares and all other nodes are grey circles. The node sizes are proportional to the *incoming strength* to that node, which is the weighted incoming degree. These two LONs were selected from the "had" and "lip" instance classes because the former have lower fractal dimensions and also a lower degree of deterministic multifractality than the latter. These two instances chosen have the same problem size, $N = 20$, and similar numbers of local optima.


In accordance with the higher fractal dimensions, the algorithm performance is lower on the "lip" group of problems. Using as a performance measure the obtained fitness (as a proportion of the global fitness), robust tabu search averaged 1.096 on the "lip" instances. For the "had" group this was 1.011. Our task in this Section of the results is to seek explanation in the networks concerning the algorithm performance differences while also paying particular attention to how their fractal nature relates to what is visually seen in the structure.

The median fractal dimension for "had20", plotted in Fig. 5a, is 2.975; for "lipa20b" it is 4.015. The range of


fractal dimensions for "had20" is around 63, and is around 49 for "lipa20b". An evident difference in the two Figures is the number and connectivity of global optima—Fig. 5a shows that the "had20" LON has many, and they appear to be densely connected to other nodes. Contrarily, the "lipa20b" LON in Fig. 5b has a single global optimum, which seems to be more sparsely connected within its network. Also noteworthy is the relative sizes of the non-optimal (grey) nodes. In Fig. 5a there are many large nodes which are sub-optimal and they have access to the global optima. Figure 5b is not the same; in fact, many of the nodes which are one step from the global optimum are very small indeed. That tells us that these nodes have small incoming degree which might hinder ascension through fitness levels during optimisation. These grey nodes are also not well-connected to each other. The opposite is true for the other network. In the "had20" LON (Fig. 5a), connectivity is so dense in the promising local optima region that visually tracking paths is impossible.

Let us now view the Figures using an algorithm performance explanation lens. Of course, the number of global optima matters and so does the accessibility of them. The "lipa20b" global optimum has many incoming edges but most of these are sourced from nodes which have low incoming degree themselves. It follows that the global optimum is less accessible. The "had20" LON, which is highly populated with edges in this promising landscape region, is probably easily solvable in part because when an algorithm reaches one of the large grey nodes (this should be likely because they have high incoming degree) there is an abundance of paths to a global optimum. The same trends are present when comparing the two networks in Figs. 6a, b.

These are the partial LONs of the "had18" and "nug16b" QAPLIB instances. Figure 6a shows "had18", which has a lower median fractal dimension (3.175) and




(a) partial LON of "had20" QAPLIB instance which has median fractal dimension 2.975 and associated tabu search performance 1.011



(b) partial LON of "lipa20b" QAPLIB instance which has median fractal dimension 4.015 and associated tabu search performance 1.154

Fig. 5 Partial local optima networks for two QAPLIB instances; only local optima which are in the fittest 10% are shown. Global optima are square and red; all others are grey circles. The size of the nodes

captures the incoming "strength" to the node in the LON, i.e. the weighted in-degree. (Color figure online)



(a) partial LON of "had18" QAPLIB instance which has median fractal dimension 3.175 and associated tabu search performance 1.011

(b) partial LON of "nug16b" QAPLIB instance which has median fractal dimension 4.090 and associated tabu search performance 1.055

Fig. 6 partial local optima networks for two QAPLIB instances; only local optima which are in the fittest 15% are shown. Global optima are square and red; all others are grey circles. The size of the nodes

captures the incoming "strength" to the node in the LON, i.e. the weighted in-degree. (Color figure online)

better tabu search performance on the underlying problem (1.011) when compared with "nug16b" shown in Fig. 6b, which has a median fractal dimension of 4.090 and tabu search performance of 1.055. Surveying the two figures, we can again visually account for the discrepancy in fractal dimension and algorithm performance by looking at the spatial complexity. Although the LON of "nug16b" has more global optima (in red), edges appear less uniformly distributed in their vicinity when compared to the LON of "had18". In addition we notice that some nodes which are one step from a global optimum in Fig. 6b are small in size. This tells us that they have low incoming degree and that the probability of search paths reaching them is small. As a consequence potential routes towards the global optima may be missed by algorithms.

5.3 Correlation analysis


Figures 7 and 8 show pairwise correlations between variables. These are Spearman rank coefficients, which are more appropriate to use where variables are not linearly related. Included are $p(ILS)$ and $p(ROTS)$ on the QAP instances, alongside the proposed fractal dimension features (*unique FDs*, *range FD*, *median FD*, *max FD* and *min FD*)—in Fig. 7, these concern deterministic dimensions; in Fig. 8, they are probabilistic. Also shown are LON features which are not associated with fractal complexity—these were introduced in Sect. 4.5. In addition, fractal dimensions excerpted from arbitrary points on the multifractal spectra are considered as features—these are *arb.dfd1* and *arb.dfd2* in Figure 7 and *arb.pfd1* and *arb.pfd2* in Fig. 8. The approach of taking an arbitrary excerpt from the spectrum and using it as a feature was taken in the previous

work on multifractality in LONs—its inclusion here facilitates a comparison between previous features and the proposed ones.

In particular we are interested in the correlation between fractal features of the LON and algorithm performance variance on the associated combinatorial problem. The intersections between the $p(ILS)$ column and the fractal feature rows in Fig. 7 reveal moderate positive correlations between them in the case of the fractal dimension range, median, maximum, and minimum—as well as the two dimension excerpts, *arb.dfd1* and *arb.dfd2*. For all of these the associated p -value is less than 0.001. We notice that the correlations are stronger than the $p(ILS)$ correlations with other LON features such as *mean.fitness*, *fitness.sinks*, *LO.neutrality*, *edges*, *diameter*, *comp.opt*, *optima*, *fit.fit.corr*, and *assortativity*. They are also slightly stronger than the correlations between $p(ILS)$ and *sinks* and *so.strength*. The correlations with dimensional summary statistics such as *med FD* appear slightly larger than *arb.dfd1* and *arb.dfd2*, which are the fractal dimension features calculated using the approach of previous work on multifractality in LONs (Thomson et al. 2018a), although the difference is not pronounced. In the $p(ROTS)$ column, there are only weak positive correlations with the fractal dimension variables, and indeed there are no strong correlations with any of the fitness landscape variables included.

Next we will consider the correlation plot which includes probabilistic fractal dimension variables in Fig. 8. Only two of the proposed fractal features show a moderate negative relationship with $p(ILS)$ —these are those related to the extent of multifractality, i.e., *unique FDs* and *range FD*. Those correlations have associated p -value of less than

Fig. 7 Spearman correlations between pairs of variables including $p(ILS)$ and $p(ROTS)$, fractal dimension metrics for the LONs, and other landscape features



0.001. These two also have a weak negative correlation with $p(ROTS)$ —for *unique FDs*, the p -value is 0.0118; for *range FD*, it is 0.0776. Again, the correlations between the other LON features and $p(ILS)$ and $p(ROTS)$ are diminutive, with the exception of *out-degree* and *clusteringcoef*. Observe that the fractal dimensions which were arbitrarily excerpted from the multifractal spectra, *arb.pfd1* and *arb.pfd2*, are far less correlated to $p(ILS)$ than two of the fractal features proposed in this work (that is, *unique FDs* and *range FD*).

5.4 Algorithm performance regression models

5.4.1 Deterministic fractal dimension features

Table 2 contains regression model statistics whose values are estimated over 10,000 random repeated subsampling iterations. Each row represents a particular model setup. The response variable is shown in the second column. The R^2 and mean squared error are given.

We can see from the R^2 values that random forest regression produces a stronger model fit. This is likely because random forest trees are adept at considering nonlinearities between variables. The amount of variance in

the iterated local search and tabu search performance which can be explained using the predictors is higher in the random forest models. The mean squared error is very low in the case of the random forest regression which is explaining $p(ILS)$. The strongest model in terms of R^2 is using random forest regression with $p(ROTS)$ as the response, with around 61% of variance being explained using the landscape features. Less variance in $p(ILS)$ response, around 48%, is explained using the same type of regression. This model setup does, however, have a much lower error rate than the associated $p(ROTS)$ model.

Now let us look at the random forest predictor importance rankings, which are provided in Fig. 9. The values are averaged over 10,000 iterations of random repeated subsampling cross-validation. For explaining $p(ILS)$, the fitness of sinks is most important. The *median fractal dimension* is second most important. These two predictors have importance values noticeably higher than the rest, although even the lowest predictors, *LO neutrality* and *minimum fractal dimension*, still have importance values around 0.009, which is around half the value of the highest. *Fitness of sinks* is again the most important factor in Figure 9b, which is the tabu search response model setup. This is followed by other fitness-based features in second, third,

Fig. 8 Spearman correlations between pairs of variables including $p(ILS)$ and $p(ROTS)$, fractal dimension metrics for the LONs, and other landscape features




Table 2 Summary statistics averaged with 10,000 iterations of random repeated subsampling cross-validation for explaining $p(ILS)$ and $p(ROTS)$. Predictors include deterministic fractal dimension LON

statistics, as well as other landscape features such as number of local optima and fitness distribution measures

Type of regression	Response variable	R^2	Mean squared error
Linear	Iterated local search performance	0.160	0.001
Linear	Robust tabu search performance	0.161	0.027
Random forest	Iterated local search performance	0.482	0.000
Random forest	Robust tabu search performance	0.611	0.011

and fourth place—*LO neutrality*, *mean fitness*, and *fitness range*. The *median fractal dimension* contributes moderately well, ranking fourth out of eight features. The position of *minimum fractal dimension* is last, but even so, it does contribute to the model.

5.4.2 Probabilistic fractal dimension features

In Table 3 is model statistics where the predictor set includes probabilistic fractal dimension features instead of the deterministic ones seen in Table 2. This is followed by the associated random forest predictor rankings in Fig. 10.

The random forest $p(ILS)$ model setup is rather weak with respect to the R^2 estimate. Indeed, in this measure it is weaker than the equivalent model which used deterministic dimensions. It should be reiterated at this point that the data-set is composed of fewer observations here than in the previous models (Tables 2, 9). There are 53 observations here, compared with a previous 85. This might impact the formulation of a well-fitting model. Nonetheless both setups with $p(ILS)$ as the response variable have markedly lower mean squared errors than their ROTs counterparts. This is also true in Table 2. Back in Table 3, the $p(ROTS)$

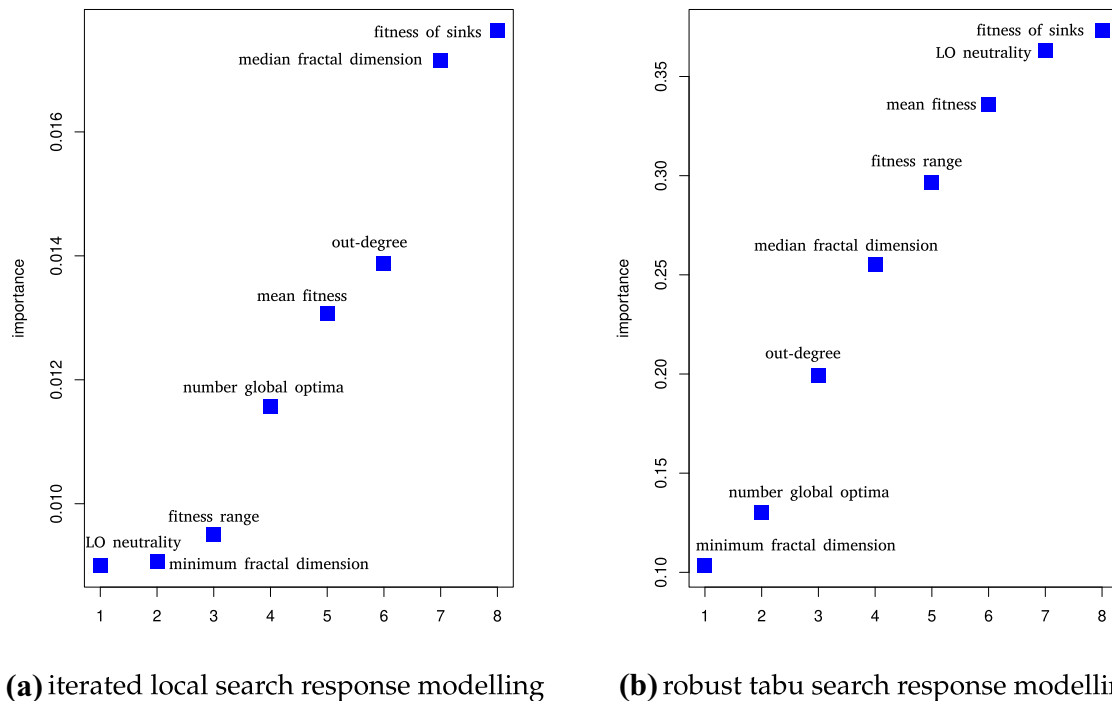


Fig. 9 Variable importance values for random forest models; the models include *deterministic* fractal dimension features as part of the predictor set

Table 3 Summary statistics estimated with 10,000 iterations of random repeated subsampling cross-validation for explaining $p(ILS)$ and $p(ROTS)$. Predictors include probabilistic fractal dimension

statistics, as well as other landscape features such as number of local optima and fitness distribution measures

Type of regression	Response variable	R^2	Mean squared error
Linear	Iterated local search performance	0.293	0.002
Linear	Robust tabu search performance	0.332	0.017
Random forest	Iterated local search performance	0.304	0.000
Random forest	Robust tabu search performance	0.578	0.033

models have higher mean squared errors but the random forest model is definitely the strongest with respect to search algorithm explanation, with around 58% being accounted for by the predictors. Although a smaller portion of variance is explained in the $p(ILS)$ models, the low mean squared errors are encouraging in accuracy terms.

In the predictor rankings, seen in Fig. 10, we draw your attention to the two fractal dimension metrics in the iterated local search plot (Fig. 10a): they are among the most important predictors, ranking third and fourth. Their values are between 0.005 and 0.006; for comparison, the importance value of the most dominant predictor, the number of global optima, is just above 0.007. The fractal dimension predictors are among the least important in the tabu model setup seen in Fig. 10b. Instead, the strongest predictors for $p(ROTS)$ appear to be relating to the local optima level fitness distribution: *mean fitness*, *LO neutrality*, *fitness of sinks*, and *fitness range*. In fact, these four form a distinct

group on the plot, far higher in importance than the remaining four (which include the fractal dimension features). Nevertheless, the lower group are not useless: their values are approximately in the range 0.04–0.07; the more important features have values between around 0.155 and 0.171.

6 Conclusion

We conducted multifractal analysis on the local optima networks (LONs) associated with a benchmark combinatorial optimisation problem library, QAPLIB. The QAPLIB instance set was more than three times the size of the set used in a prior study (Thomson et al. 2018a) and raised the considered problem sizes from $N \leq 28$ to $N \leq 50$. A recent and refined LON construction algorithm (Ochoa and Herrmann 2018) was used to build the LONs.

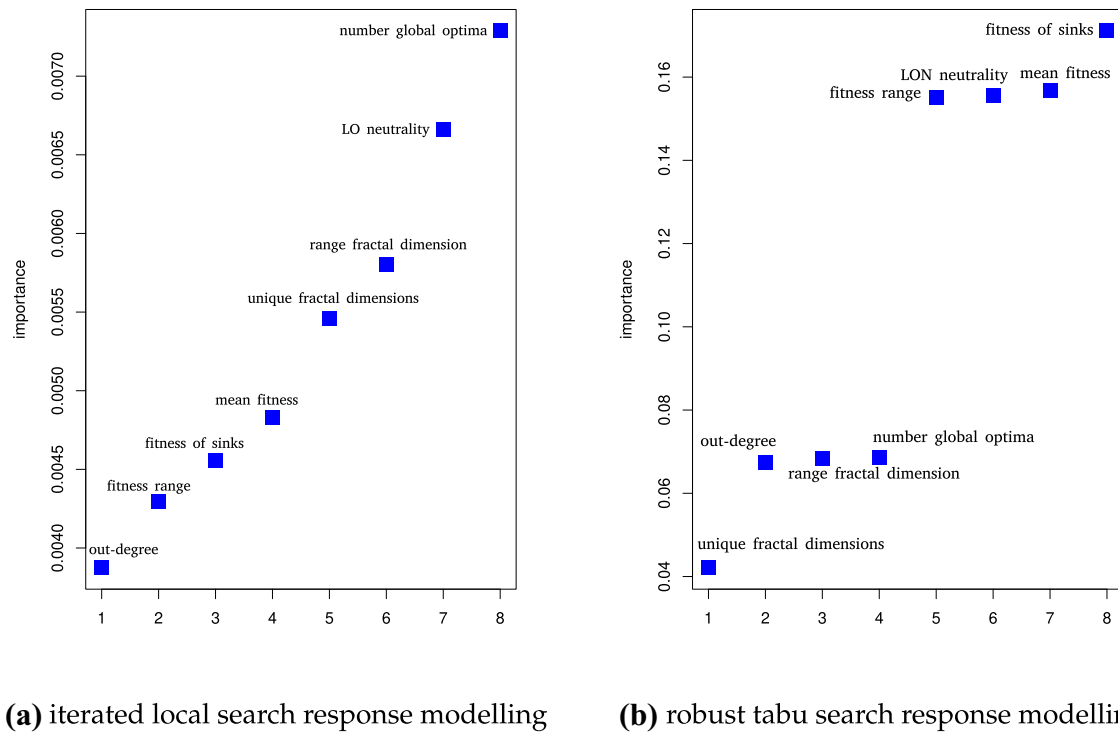


Fig. 10 Variable importance values for random forest models; the models include *probabilistic* fractal dimension features as part of the predictor set

Relationships between fractal dimension features of LONs and algorithm performance by iterated local search (ILS) were established using correlation analysis, visual analysis tools, and linear and random forest regression with random repeated subsampling cross-validation. The results showed that the *extent* of multifractality and the highness of values in the dimension spectrum can contribute towards partially predicting or explaining ILS algorithm performance. Features of the fractal dimension distribution for the LONs also displayed individual pairwise correlations to ILS algorithm performance. Fractal dimension features in LONs were less important for predicting tabu search but could still contribute some information. Sampled fitness levels in the LON were more important in these models.

A limitation to our approach is that the features are computed from sampled LONs, whose characteristics can alter markedly with a different sampling effort (Božejko et al. 2018). Nevertheless, the LON features can contribute towards explaining algorithm performance within regression models—it follows that they are useful, even if the sample illustrates a certain version of the fitness landscape. In addition, there is no alternative to sampling when analysing QAP LONs of moderate size (greater than $N = 11$ according to Daolio et al. (2011)). Another consideration is the random selection of box centres in the sandbox algorithm for multifractal analysis. We argue that the number of algorithm iterations—each of which contains a random

selection of centres and produces its own fractal dimension—should mitigate the variation induced by random selection. In addition, the resulting fractal dimension features help to explain metaheuristic performance in statistical analysis over 85 observations. This implies that the randomness inherent to the approach does not affect the empirical usefulness of the computed fractal dimensions.

The present study could serve as a foundation for further work within this research avenue which remains untapped. In particular, we would like to expand the maximum size of the problems studied, as well as venturing to other domains and to constrained problems. Finally, we conclude with a remark concerning our interest in studying the relationship between perturbation *strength* used to generate the LONs, and the calculated fractal dimensions of that LON.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bożejko W, Gnatowski A, Niżyński T, Affenzeller M, Beham A (2018) Local optima networks in solving algorithm selection problem for tsp. In: International conference on dependability and complex systems, Springer, pp 83–93
- Burkard RE, Karisch SE, Rendl F (1997) Qaplib—a quadratic assignment problem library. *J Global Optim* 10(4):391–403
- Caldwell CB, Stapleton SJ, Holdsworth DW, Jong RA, Weiser WJ, Cooke G, Yaffe MJ (1990) Characterisation of mammographic parenchymal pattern by fractal dimension. *Phys Med Biol* 35(2):235
- Chicano F, Daolio F, Ochoa G, Vérel S, Tomassini M, Alba E (2012) Local optima networks, landscape autocorrelation and heuristic search performance. In: International conference on parallel problem solving from nature, Springer, pp 337–347
- Daolio F, Verel S, Ochoa G, Tomassini M (2010) Local optima networks of the quadratic assignment problem. In: IEEE congress on evolutionary computation, IEEE, pp 1–8
- Daolio F, Tomassini M, Vérel S, Ochoa G (2011) Communities of minima in local optima networks of combinatorial spaces. *Phys A* 390(9):1684–1694
- Dicke M, Burrough PA (1988) Using fractal dimensions for characterizing tortuosity of animal trails. *Physiol Entomol* 13(4):393–398
- Esgiar AN, Naguib RN, Sharif BS, Bennett MK, Murray A (2002) Fractal analysis in the detection of colonic cancer images. *IEEE Trans Inf Technol Biomed* 6(1):54–58
- Furuya S, Yakubo K (2011) Multifractality of complex networks. *Phys Rev E* 84(3):036–118
- Hadjileontiadis L, Douka E (2007) Crack detection in plates using fractal dimension. *Eng Struct* 29(7):1612–1625
- Herrmann S, Ochoa G, Rothlauf F (2016) Communities of local optima as funnels in fitness landscapes. In: Proceedings of the 2016 on genetic and evolutionary computation conference, ACM, pp 325–331
- Hoos HH, Smyth K, Stützle T (2004) Search space features underlying the performance of stochastic local search algorithms for max-sat. In: International conference on parallel problem solving from nature, Springer, pp 51–60
- Hutter F, Hoos HH, Stützle T (2007) Automatic algorithm configuration based on local search. In: Proceedings of the 22nd national conference on artificial intelligence, vol 2, AAAI Press, AAAI'07, p 1152–1157
- Lawler EL (1963) The quadratic assignment problem. *Manag Sci* 9(4):586–599
- Liu JL, Yu ZG, Anh V (2015) Determination of multifractal dimensions of complex networks by means of the sandbox algorithm. *Chaos Interdiscip J Nonlinear Sci* 25(2):023–103
- Locatelli M (2005) On the multilevel structure of global optimization problems. *Comput Optim Appl* 30(1):5–22
- Mandelbrot BB (1972) Possible refinement of the lognormal hypothesis concerning the distribution of energy dissipation in intermittent turbulence. In: Statistical models and turbulence, Springer, pp 333–351
- Mandelbrot BB (1975) Stochastic models for the earth's relief, the shape and the fractal dimension of the coastlines, and the number-area rule for islands. *Proc Nat Acad Sci* 72(10):3825–3828
- Mandelbrot BB, Fisher AJ, Calvet LE (1997) A multifractal model of asset returns. Cowles Foundation discussion paper
- Mashiah A, Wolach O, Sandbank J, Uziel O, Raanani P, Lahav M (2008) Lymphoma and leukemia cells possess fractal dimensions that correlate with their biological features. *Acta Haematol* 119(3):142–150
- McMenemy P, Veerapen N, Ochoa G, (2018) How perturbation strength shapes the global structure of tsp fitness landscapes. In: Liefvooghe A, López-Ibáñez M (eds) Evolutionary computation in combinatorial optimization. *EvoCOP*, (2018) Lecture Notes in Computer Science, vol 10782. Springer, Cham
- Merz P (2004) Advanced fitness landscape analysis and the performance of memetic algorithms. *Evol Comput* 12(3):303–325
- Merz P, Freisleben B (2000) Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. *IEEE Trans Evol Comput* 4(4):337–352
- Ochoa G, Herrmann S (2018) Perturbation strength and the global structure of QAP fitness landscapes. In: International conference on parallel problem solving from nature, Springer, pp 245–256
- Ochoa G, Veerapen N (2016) Additional dimensions to the study of funnels in combinatorial landscapes. *Proc Genetic Evol Comput Conf 2016*:373–380
- Ochoa G, Veerapen N (2018) Mapping the global structure of tsp fitness landscapes. *J Heuristics* 24(3):265–294
- Ochoa G, Tomassini M, Vérel S, Darabos C (2008) A study of nk landscapes' basins and local optima networks. In: Proceedings of the 10th annual conference on Genetic and evolutionary computation, ACM, pp 555–562
- Ochoa G, Veerapen N, Daolio F, Tomassini M, (2017) Understanding phase transitions with local optima networks: Number partitioning as a case study. In: Hu B, López-Ibáñez M (eds) Evolutionary computation in combinatorial optimization. *EvoCOP 2017*, Lecture notes in computer science, vol 10197. Springer, Cham
- Pitzer E, Affenzeller M (2012) A comprehensive survey on fitness landscape analysis. In: Recent advances in intelligent engineering systems, Springer, pp 161–191
- Richter H (2018) Scale-invariance of ruggedness measures in fractal fitness landscapes. *Int J Parallel Emergent Distrib Syst* 33(5):460–473
- Saeedi P, Sorensen S (2009) An algorithmic approach to generate after-disaster test fields for search and rescue agents. *Proc World Congress Eng* 1:93–98
- Song C, Havlin S, Makse HA (2005) Self-similarity of complex networks. *Nature* 433(7024):392–395
- Song C, Havlin S, Makse HA (2006) Origins of fractality in the growth of complex networks. *Nat Phys* 2(4):275
- Song YQ, Liu JL, Yu ZG, Li BG (2015) Multifractal analysis of weighted networks by a modified sandbox algorithm. *Sci Rep* 5:17628
- Stadler PF (2002) Fitness landscapes. Biological evolution and statistical physics lecture notes in physics 585:183–204
- Stützle T (2006) Iterated local search for the quadratic assignment problem. *Eur J Oper Res* 174(3):1519–1539
- Taillard É (1991) Robust taboo search for the quadratic assignment problem. *Parallel Comput* 17(4–5):443–455
- Thomson SL, Verel S, Ochoa G, Veerapen N, Cairns D (2018a) Multifractality and dimensional determinism in local optima networks. In: Proceedings of the genetic and evolutionary computation conference, pp 371–378
- Thomson SL, Verel S, Ochoa G, Veerapen N, McMenemy P, (2018b) On the fractal nature of local optima networks. In: Liefvooghe A, López-Ibáñez M (eds) Evolutionary computation in combinatorial optimization. *EvoCOP 2018*, Lecture Notes in Computer Science, vol 10782. Springer, Cham
- Uahabi LK, Atounti M (2017) New approach to the calculation of fractal dimension of the lungs. *Ann Univ Craiova-Math Comput Sci Ser* 44(1):78–86
- Verel S, Daolio F, Ochoa G, Tomassini M (2011) Local optima networks with escape edges. In: International conference on artificial evolution (*Evolution Artificielle*). Springer, pp 49–60
- Verel S, Daolio F, Ochoa G, Tomassini M (2018) Sampling local optima networks of large combinatorial search spaces: The qap

- case. In: International conference on parallel problem solving from nature. Springer, pp 257–268
- Weinberger ED, Stadler PF (1993) Why some fitness landscapes are fractal. *J Theor Biol* 163(2):255–275
- Werner DH, Ganguly S (2003) An overview of fractal antenna engineering research. *IEEE Antennas Propag Mag* 45(1):38–57
- Zelinka I, Zmeskal O, Saloun P (2014) Fractal analysis of fitness landscapes. In: Recent advances in the theory and application of fitness landscapes, Springer, pp 427–456

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.