

Multifractality and Dimensional Determinism in Local Optima Networks

Sarah L. Thomson
University of Stirling
Stirling, United Kingdom
s.l.thomson@stir.ac.uk

Sébastien Verel
Université du Littoral Côte d'Opale
Calais, France
verel@uni-littoral.fr

Gabriela Ochoa
University of Stirling
Stirling, United Kingdom
gabriela.ochoa@stir.ac.uk

Nadarajen Veerapen
University of Stirling
Stirling, United Kingdom
nadarajen.veerapen@stir.ac.uk

David Cairns
University of Stirling
Stirling, United Kingdom
david.cairns@stir.ac.uk

ABSTRACT

We conduct a study of local optima networks (LONs) in a search space using fractal dimensions. The fractal dimension (FD) of these networks is a complexity index which assigns a non-integer dimension to an object. We propose a fine-grained approach to obtaining the FD of LONs, using the probabilistic search transitions encoded in LON edge weights. We then apply multi-fractal calculations to LONs for the first time, comparing with mono-fractal analysis. For complex systems such as LONs, the dimensionality may be different between two sub-systems and multi-fractal analysis is needed. Here we focus on the Quadratic Assignment Problem (QAP), conducting fractal analyses on *sampled* LONs of reasonable size for the first time. We also include fully enumerated LONs of smaller size. Our results show that local optima spaces can be multi-fractal and that valuable information regarding probabilistic self-similarity is encoded in the edge weights of local optima networks. Links are drawn between these phenomena and the performance of two competitive metaheuristic algorithms.

CCS CONCEPTS

• **Computing methodologies** → **Search methodologies**;

KEYWORDS

Fitness Landscapes, Quadratic Assignment Problem, Local Optima Networks, Fractal Dimension

ACM Reference Format:

Sarah L. Thomson, Sébastien Verel, Gabriela Ochoa, Nadarajen Veerapen, and David Cairns. 2018. Multifractality and Dimensional Determinism in Local Optima Networks. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan*, Keiki Takadama and Hernan Aguirre (Eds.). ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3205455.3205472>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '18, July 15–19, 2018, Kyoto, Japan

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5618-3/18/07...\$15.00

<https://doi.org/10.1145/3205455.3205472>

1 INTRODUCTION

Interest is growing in the fractal structure of complex networks. In particular, some networks have recently been shown to be *self-similar* [2, 9], in that smaller copies of patterns can be found. The *fractal dimension* has been of particular interest, where a non-integer dimension can be assigned to an object as an index of complexity. This captures the way that a shape fills the geometric space it resides in, in terms of the relationship between scale and observed detail.

Figure 1 shows two patterns with different fractal dimensions. While Fig. 1a has dimension 1.4649, Fig. 1b is 1.7848. We can see the latter is convoluted and complex. The former is more simple in the way that it fills space. We can imagine if these were search spaces, that intuitively the latter would provide more complications during optimisation.

Assigning a single fractal dimension may not be appropriate for a real-world complex network. Doing so has an inherent strong assumption that the self-similarity in the network is roughly homogeneous. Such an assumption can sometimes be misguided, as found in recent studies proposing *multifractal* analysis for complex networks [2, 10]. This type of analysis means that a spectrum of dimensions is produced to describe an object.

Indeed, Benoit Mandelbrot — the pioneer of fractal geometry — argued that a continuous spectrum of dimensions are necessary to properly capture the complicated dynamics of a real-world system [7].

The main contributions of this paper are:

- (1) First application of fractal analysis to *sampled* local optima networks.
- (2) A proposed methodology for probabilistic fractal dimensions when studying LONs.
- (3) First multi-fractal analysis of LONs.
- (4) A comparison of monofractal, multifractal, and probabilistic fractal characterisation for LONs.
- (5) First study to contrast LON fractal attributes with competitive search algorithm performance.




Figure 1: Two fractals, both with dimensions between one and two.

2 BACKGROUND

2.1 Fitness Landscapes in Network Form

A fitness landscape [11] is a triplet (S, N, f) where S is the set of all possible solutions, $N : S \rightarrow 2^S$, a neighbourhood structure, is a function that assigns to every $s \in S$ a set of neighbours $N(s)$, and f is a fitness (objective value) function such that $f : S \rightarrow \mathbb{R}$, where the fitness value is a real number that can be viewed as the *height* of a given solution in the landscape.

Local Optima Network. A *Local Optima Network* (LON) models a subset of the fitness landscape. The local optima comprise the node set, V . The edges, E , represent search connectivity between two local minima. Specifically, a directed *escape* edge is traced if the destination can be reached from the source by doing one perturbation followed by hill-climbing.

2.2 A Fractional Dimension

Generally, dimensions are measured as integer numbers: a two-dimensional square, for example, or a three-dimensional cube. In reality, though, many patterns in nature exhibit a *fractal* or *fractional* dimension [7]. This means that the shape doesn't fill space in a way that conforms to the available integer dimensions. Increased interest in complex networks has naturally lead to the question of whether these might contain some self-similarity. For this, a decision must be made on how to calculate the fractal geometry of a network. In general, to obtain the fractal dimension of an object, two measures are needed: the *scale* at which the shape is measured, s , and the detail observed at that scale, d . A popular algorithm for networks found in the literature is termed *box-counting* [9]. The general principle is boxing together nodes which can be considered 'neighbours'. Nodes are iteratively boxed together if they are within s links of each other. An elementary example of this is shown in Figure 2. The parameter s provides the scaling or coarseness factor needed for fractal dimension calculation. The extent of detail observed is simply the number of boxes necessary to fully cover the network at that scale.

The box-counting algorithm is agnostic of the semantics of the network, and does not take into account node attributes. In a local optima network, we are modelling a subset of the fitness landscape, meaning node fitness is of great importance.

A previous study modified box-counting [14] for local optima networks such that fitness distance was considered as well as link distance: two nodes can be boxed together iff $d(lo_i, lo_j) < s$ and

$|f(lo_i) - f(lo_j)| < \epsilon$. We use and extend this approach in our fractal analysis algorithms for this study.

2.3 The Multifractal Spectrum

Real world systems often do not conform to the assumed homogeneity traditional (mono)fractal dimension takes [6]. In this case, the extent of fractal geometry cannot be characterised by a single value, but rather a spectrum of numbers.

Indeed, pioneering authors in the fractal community have previously stated that the multifractal approach is necessary for many real world shapes [4, 7].

In the past few years, studies have surfaced where the concept of *multifractality* has been applied to complex networks [2, 8, 10]. In all cases, some of the networks studied required a dimension spectrum to characterise the self-similar properties.

The competitive approach to calculating a set of generalised fractal dimensions for a complex network is called a *Sandbox Algorithm*, which is a variant of the *box-counting* process introduced in Section 2.2.

Algorithm 1 Sandbox Algorithm for Multifractal Analysis of a Local Optima Network

Initialisation:

- ▷ V : nodes in network
- ▷ CV, NCV : center nodes, non-center nodes
- ▷ R : set of sandbox radius values
- ▷ E : set of fitness difference values
- ▷ Q : set of values for q , i.e. set of numbers for the spectrum of dimension

$CV = []$, $NCV = V$

for q in Q **do**

for r in R **do**

for e in E **do**

$V = \text{shuffle}(V)$

$CV = \text{rand}(V, 100)$

$\text{sizes} = \{ \}$

 ▷ sizes : set of sandbox sizes

for c in CV **do**

$\text{num.covered} = 0$

for v in V **do**

$d = \text{dist}(c, v)$

$j = \text{diff}(f(c), f(v))$

if $j < e$ and $d < r$ **then**:

$\text{num.covered} = \text{num.covered} + 1$

$\text{sizes} = \text{sizes} \cup \{[\text{num.covered}]^{q-1}\}$

$\text{mean.sizes}[q][e][r] = \text{mean}(\text{sizes})$

A set of nodes are randomly selected to be *sandbox centers*. A sandbox surrounds a center. These are allocated a maximum radius, r . For each sandbox, s , nodes which are at most r links away from the central node are added to s . The average of the box sizes is taken. This process is repeated for various values for r . The basis for the spectrum of dimensions is provided by using an arbitrary set of numbers, $q \in Q$. The value for q is used in the equation to obtain fractal dimension:




Figure 2: Three iterations of a box-counting algorithm with $s = 2$ (s is the maximum allowed edge distance separating nodes which can be boxed together). Node colour indicates membership to a box. The start point is four boxes. A node is selected as seed (hub) for first box, and so on for remaining nodes. Boxing two nodes is amalgamating them and considering them as a single node. The links retained are those belonging to the amalgamated node, or hub. For this illustrative network, scale = 2 and detail = $\frac{1}{8}$, i.e. $2^x = 0.125$, which after solving for x gives a calculated dimension of 3.00.

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

We use the range of values suggested in the literature [10], starting at -10 and ending at +10, in step sizes of one.

We then extend this process, to allow for the special case of a local optima network. Because these networks are a compression of the fitness landscape, node fitness should also be taken into account. We can easily add this to the Sandbox algorithm, in the form of an additional condition for a particular node's inclusion or exclusion from a sandbox. Specifically, a node n can be included in a central node c 's sandbox iff $\text{distance}(n, c) \leq r$ and $|(f(n) - f(c))| < \epsilon$, where r is the radius of the sandbox and ϵ is the maximum allowable fitness difference between the two nodes (local optima).

Both r and ϵ should be varied to assess the nature of the scaling in the object. Therefore, we have a set of values for q , r , and ϵ . Pseudocode for the sandbox algorithm, modified for the local optima network case, is seen in Algorithm 1.

To obtain the set of generalised fractal dimensions, $D(q)$, we perform a linear regression of the observed detail, $\log(\text{mean}(\text{detail})^{q-1})$, and the scale, $(q-1)\log(\text{scale})$, where detail is the mean size of a sandbox as a proportion of the size of the network, and where scale is the radius used as a proportion of the diameter of the network.

2.4 Probabilistic Dimensions

In a previous study, fractal dimensions were calculated on the LONs of a set of NK Landscape instances [14]. The box-counting algorithm used considered edge distance between nodes as a scaling factor for dimension calculation.

The problem with this methodology is that if there is a small network diameter, there can be a coarse-levelled reduction in the detail observed at a particular scale. To arrive at a finer-grained view of the relationship between scale and observed detail in networks, we can use the edge weightings in the network.

Edge weights in local optima networks represent the probability of a search path between two optima being followed. We can modify the traditional box-counting algorithm to consider this instead: two nodes can only be boxed together if it is deemed *likely* they will be linked, in that the edge between them is of a certain weight.

To normalise the edge weights in the networks, we simply subtract them from one, i.e. the standardised weight sw is defined as

$1-w$ and we have the opposite probability. Mathematically, nodes x and y are neighbours iff $sw(E_{x,y}) < \beta$. Deciding a suitable value for β is important. Here we use a set of values for the probability parameter, $\beta \in 0.90, 0.96$, in step sizes of 0.02. The lower the value for β , the stricter the probability condition. We use this range based on preliminary runs and the observed distribution of network edge weights. In this way, dimensions can be calculated such that the stochastic nature of the information in a local optima network is respected.

3 EXPERIMENTAL SETUP

3.1 Benchmark Problem

We focus on a benchmark combinatorial domain here: the Quadratic Assignment Problem (QAP). In the QAP a set of facilities with given flows have to be assigned to a set of locations with given distances in such a way that the sum of the product of flows and distances is minimised.

The cost associated with a permutation π is given by:

$$C(\pi) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\pi_i \pi_j}$$

where n denotes the number of facilities/locations and $A = \{a_{ij}\}$ and $B = \{b_{ij}\}$ are referred to as the distance and flow matrices, respectively. The structure of these two matrices characterises the class of instances of the QAP problem.

Our experiments are based on both benchmark and synthetic instances. Some are from the well studied QAP Library (QAPLIB). One of the contributions of this paper is the fractal analysis of *sampled* local optima networks. A previous study [14], the first of its kind, considered only fully enumerated networks. The instances from the QAPLIB are larger (between twelve and 28 locations, detailed in Table 1). For these, a full enumeration is not feasible. We use sampled local optima networks for these, provided by the authors of [3]. To obtain the local optima during the sampling, they used one-swap best-improvement hill-climbing. For defining the edges, they used random kick moves of three swaps. While the sampling introduces an inevitable bias, the bias is towards the search space regions likely to be encountered by heuristic search.

In addition to those, we use structured instances produced using the generator proposed in [5]. The generator produces flow entries

that are non-uniform random values. Clusters of points are placed in compact circular areas, and all of these clusters are enclosed in a large circle. These instances have the so-called “real-like” structure since they resemble the structure of QAP problems found in practical applications.

The synthetic problem instances are fully enumerated, and as such are constrained to a small size of 11. We use 30 local optima networks of this type, as made available by the authors of [1]. To obtain the local optima, the authors conducted best-improvement pairwise exchange local search. Two nodes are connected by an edge if the destination node can be reached from the source using two pairwise exchanges (perturbation) followed by single-exchange hill-climbing (local search).

3.2 Normalisation of Node Fitness Difference

We use *logarithmic returns* to normalise fitness values here, to allow the setting of the fitness parameter ϵ to be independent of any particular fitness distribution. To calculate the logarithmic return of two values, x_1 and x_2 — in our case, the objective values of two local optima which are connected by a directed LON edge — we do $\log(x_2/x_1)$.

3.3 Algorithms

The nature of the experiments can be split largely in two: those relating to *probabilistic* dimension and the calculations of the *multifractal* spectrum. For both algorithms, we begin with a box-counting algorithm for (mono)fractal analysis from the literature [9], which is written in C.

Probabilistic Dimension Analysis. A separate variant of the box-counting algorithm is used to calculate what we term here *probabilistic dimensions*, as described in full in Section 2.4. The network edge weights represent search probabilities, so calculating a dimension using them reflects probabilistic complexity. We also apply the original box-counting algorithm to the LONs for comparison.

Multifractal Dimension Analysis. To conduct multifractal analysis on the local optima networks, we use the algorithm from [9] as a starting point. We then proceed to implement the sandbox algorithm proposed in [6] in C. This process is outlined in Algorithm 1 and described in Section 2.3.

Metaheuristics. For all fractal or landscape analysis conducted, the obtained features of the local optima space must be contrasted with search difficulty to arrive at some useful conclusion. Accordingly, we collect heuristic search data on the underlying problem instances. We select two competitive metaheuristics for the QAP to this end: Robust Tabu Search (ROTS) [13] and Stutzle’s improved Iterated Local Search (ILS) [12]. For the latter, we use the first-improvement version, with a pairwise exchange as the local search and $3n/4$ exchanges for the perturbation.

4 RESULTS

We calculate the various fractal dimensions and associated metrics of a set of local optima networks extracted from QAP instances. For each problem, we run mono-fractal analysis (also referred to as *deterministic* dimension analysis), multi-fractal analysis, and probabilistic dimension analysis. These three are compared with

one another, to assess which is the best fit for local optima networks. In all cases, the fractal results are contrasted with observed search effort by the two metaheuristics.

4.1 LONs: Monofractal or Multifractal?

The question of whether a local optima network is better suited to a single fractal characterisation or a spectrum of dimensions is now addressed.

Figures 3a, 3b, and 3c show boxplots of the fractal dimensions of the local optima networks. Figure 3a represents the “real-like” instances, and 3b and 3c are the *Nug* and *Chr* instances from QAPLIB, respectively. Each bar represents the spread of dimensions when using a particular fractal algorithm variant. The upper two bars represent results obtained from using the traditional box-counting algorithm, which gives a *monofractal* dimension, while the lower three are *multifractal* dimensions — for the same networks — taken from different arbitrary points in the full spectrum. These are annotated with *mf1*, *mf2* and *mf3*.

Surveying Figure 3a, it can be seen that the ranges of the multifractal dimensions (*mf1*, *mf2*, and *mf3*) deviate quite significantly from the monofractal ranges for the same instances (*det1* and *det2*). Generally, the multifractal dimensions are smaller. In addition, the three dimension ranges taken from different points on a single multifractal spectrum are markedly different to one another. This implies a lack of homogeneity or agreement in terms of complexity in different parts of these local optima networks, indicating *multifractality*. A similar phenomenon is seen in Figure 3b (*Nug* instances), although here the dimensions are higher.

The *Chr* benchmark QAP problems display quite different results: observing Figure 3c, we can see that the spread of the multi dimensions here (*mf1*–*mf3*) are much more parallel to the monofractal dimensions in the upper two bars (*det1* and *det2*). The means are very similar; furthermore, when comparing the three lower distributions with one another, they are very close indeed. This means that even when taking dimensions from different points on the spectrum, they are almost the same. This implies that these are *monofractal* networks.

4.2 Dimensional Determinism

We now conduct a comparison between the *deterministic* fractal dimensions we calculated against the *probabilistic* ones. The respective variants of the box-counting algorithm are run on each problem instance considered.

The difference in the algorithms lies in the definition of the scaling factor. In the former, the mere *presence* of links between nodes is used for the boxing criterion. The other box-counting variant uses the *probability* that the search path between the two optima will be traversed; this probability is encoded in the link weights in the networks. In this way, the existence of link between two optima is not enough — it must pass a probability condition.

Figures 4a and 4b show the distribution of the deterministic and probabilistic dimensions. The former reflects the “real-like” instances, while the latter is for the QAPLIB instances. Similarly to Figures 3a–3c, the upper two bars are the values obtained by the traditional deterministic box-counting algorithm for complex networks (*det1* and *det2*). The lower three show results from using the

Table 1: QAPLIB instances used.

class	instance name (number is problem size)	description	number of instances
chr	{12{a-c},15{a-c},18{a-b},20{a-c},22{a-b}}	tree/complete	13
nug	{12,14,15,16{a-b},17,18,20,21,22,25,27,28}	grid-based	12

algorithm variant which uses the probabilistic LON edge weights — instead of simply the existence of a link — as the scaling factor. These are labelled with the parameter setting used in that variant for β , which is the maximum opposite probability (i.e., minimum probability) that two nodes (local optima) will be connected during search. Outliers are shown as red squares.

Figures 4a and 4b indicate that using a deterministic fractal analysis on a local optima network might miss important information relating to *likely* search trajectories, encoded in the edge weightings: we can see that for the QAPLIB instances (4b), the calculated probabilistic dimensions are usually much higher (see *0.90-0.96*), implying a more multi-faceted and complex space-filling behaviour from the networks than the deterministic dimensions hint at (*det1-2*). Notice also that the calculated dimensions are lower for the stricter boxing constraint of $\beta = 0.90$ than for the more lenient $\beta = 0.93$. This will be due to nodes rarely passing the criteria, and therefore not a lot of boxing takes place: the observed dimensions are then constrained to the linear relation between β (0.90) and the proportional number of boxes needed to cover the network, which here must be 0.2 (number of boxes = $\frac{n}{5}$) because the fractal dimensions are ~ 15.28 , obtained from solving for x the equation $0.90^x = 0.2$.

4.3 Search Effort on Fractal Landscapes

To quantify performance of the two competitive heuristics on the instances, we use two metrics for each: the number of iterations to reach the global optimum or optima (*ILS.t* and *ROTS.t* for ILS and ROTs, respectively), and the percentage above the global minimum fitness within a fixed budget (*ILS.p* and *ROTS.p*). The two algorithms are deployed 1000 times on each problem instance.

We now examine the correlations between pairs of variables. Figure 5a shows a correlation matrix and is focused on multifractality in the local optima networks, while Figure 5b looks at probabilistic dimensionalities.

Pairwise Spearman correlation coefficients are shown in the upper triangle of the panel. The overall correlation is shown in black text, with a split into *class* of instance shown through use of colour. Red is the QAPLIB instances, while green is the “real-like” ones. Density plots populate the middle diagonal row, with scatterplots forming the lower triangles.

From left to right by column, Figure 5a includes two dimensions taken from different arbitrary points in the multifractal spectrum (*mf1* and *mf4*) as features; ILS performance metrics, denoted as *ILS.t* (iterations to the global optimum or optima) and *ILS.p* (percentage above the global best fitness within a given budget); ROTs metrics, shown as *ROTS.t* (iterations to the optimum) and *ROTS.p* (percentage above the global minimum within a fixed budget); the number of local optima in the fitness landscape (*optima*), and the monofractal dimensions (*det1* and *det2*).

Figure 5b contains the same variables, with the exception of fractal dimension type. Instead of dimensions taken from different points on a spectrum, this matrix considers probabilistic dimensions with different values for β : *probabilistic1* (0.90) *probabilistic2* (0.93) and *probabilistic3* (0.96).

Let us look at each of the correlation matrices in turn. Figure 5a displays a few important points for this study. First, one of the multifractal dimensions (*mf1*) has moderately strong positive correlations with three out of the four search performance metrics we consider here. This can be seen by following along the row labelled *mf1* and checking the intersections with the search measures. For all three metaheuristic performance features (*ILS.t*, *ILS.p*, *ROTS.t*), the positive association with the dimensionality suggests that the latter hinders the efficiency of the algorithms. What we notice, however, is the relationships with search in the case of the other dimension point from the spectrum are much weaker. This has two implications: a single fractal dimension is insufficient to characterise the dynamics written in a local optima network (if it was, each of these would exhibit similar behaviour), and not all fractal geometries in the space of local optima are significantly linked to search challenges.

A further remark could be made when considering the dimension-dimension correlations. While they are correlated, with p-values indicating statistical significance, there is a marked difference when comparing those of the benchmark instances with the synthetic ‘real-like’ instances. We can see this clearly from noting the contrast between the correlations in red text (library instances) with those in green (synthetic, structured instances).

The extreme resemblance of the dimensions calculated on the benchmark instances hints that these might actually be monofractal networks, with somewhat uniform self-similarity. Conversely, the multifractal dimensions of the generated instances have far weaker correlations with one another.

Turning our attention to Figure 5b, we can observe that the connections between fractal geometry and search difficulty is more pronounced in the probabilistic dimensions than in the deterministic ones. This can be noted by checking the intersections of *probabilistic1*, for example, against *ILS.t* and *ROTS.t*, and then comparing with the deterministic dimension rows (*det1*, *det2*).

The probabilistic dimensions all exhibit moderate-to-strong correlations with the runtime of the metaheuristics, with $p < 0.001$ in all cases. This suggests that intricate paths or patterns in the local minima space — and specifically, probabilistically *likely* search trajectories — are linked to slower performance by search algorithms on the underlying problem.

This is more prominent in the case of the QAPLIB instances, shown in red, than the generated “real-like” ones in green.


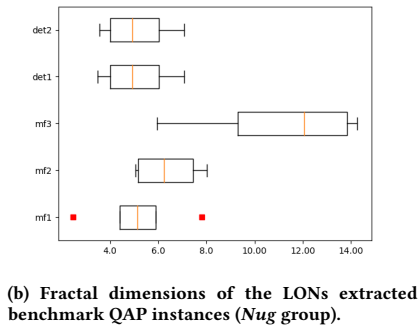



Figure 3: Quantile ranges of some of the multifractal dimensions (taken from arbitrary positions on the spectrum) for all LONs considered in this study (the three lower bars: *mf1*, *mf2* and *mf3*). The three are different only in their value for q . Deterministic fractal dimension ranges for the same networks are shown with the two right-most bars: *det1* is the dimensionalities when considering a fitness-discrepancy condition for boxing nodes, while *det2* reflects the results arising from not using this condition. The class which the instances belong to are indicated in the captions.

Noteable also is the apparent discrepancies between deterministic and probabilistic dimensions. With each other, these have only a weak-to-moderate correlation.

This says that the two fractal dimension definitions provide different information about repeating patterns in the space of the





Figure 4: Quantile ranges of the probabilistic fractal dimensions for all LONs considered in this study (the lower three bars: 0.90, 0.93, and 0.96). These are dimensions calculated with β set at 0.90, 0.93 and 0.96, respectively. Deterministic fractal dimension ranges for the same networks are shown with the two right-most bars: *det1* is the dimensionalities when considering a fitness-discrepancy condition for boxing nodes, while *det2* reflects the results arising from not using this condition.

local optima; indeed, the density plots (of *det1* and *probabilistic2*, for example) show very different distributions.

Table 2: Linear models using fractal dimensions of LONs to predict runtime of metaheuristics.

	IIS	ROTS
Predictor	Estimate	Estimate
Probabilistic Dim. ¹	0.1454 (0.212)	-0.031 (0.111)
Probabilistic Dim. ²	0.768(0.288)*	0.260(0.151)
Deterministic Dim. ¹	0.857 (1.090)	0.034 (0.570)
Deterministic Dim. ²	-0.751(1.068)	0.064 (0.559)
Optima	1.270(0.271)***	0.453(0.142)**
R^2	0.640	0.422

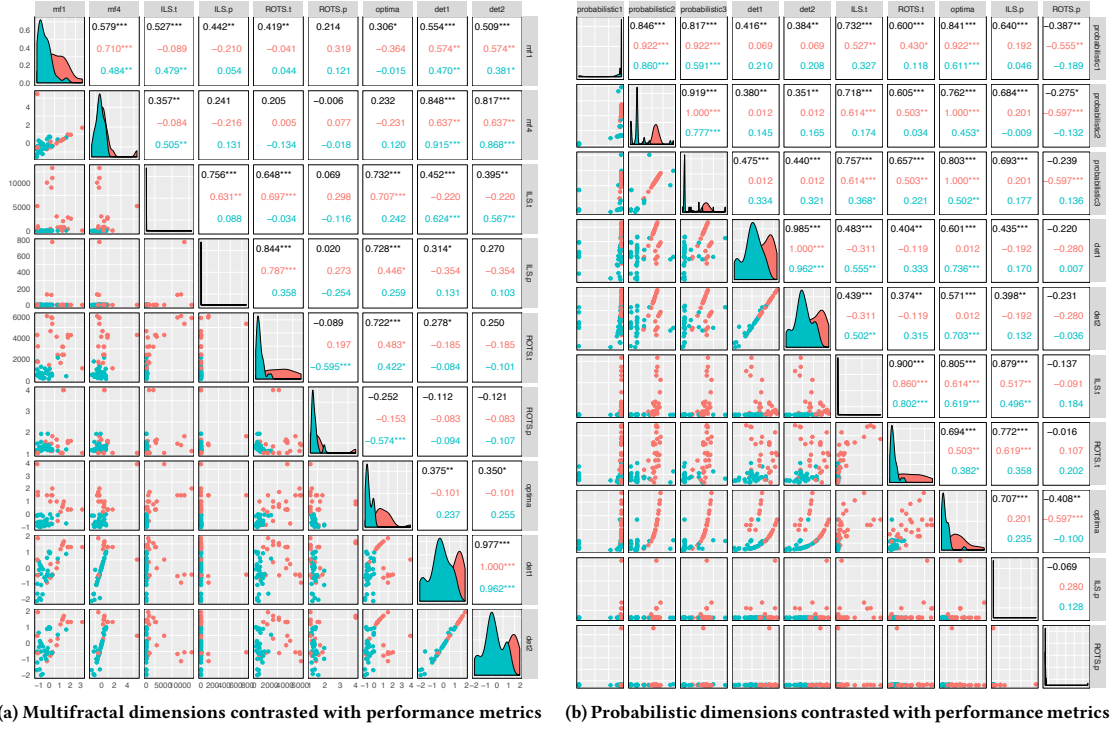


Figure 5: Correlation matrices of performance metrics and landscape features (see facet titles). Lower triangle: pairwise scatter plots. Diagonal: density plots. Upper triangle: pairwise Spearman's rank correlation, *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Table 3: Linear models using fractal dimensions of LONs to predict success of metaheuristics.

	ILS	ROTS
Predictor	Estimate	Estimate
Multifractal Dim. ¹	0.763 (0.210)***	0.390(0.118)**
Multifractal Dim. ²	-0.519 (0.219)	-0.213 (0.123)
Deterministic Dim. ¹	1.557 (0.818)	-0.034 (0.458)
Deterministic Dim. ²	-1.412 (0.789)	0.000 (0.442)
Optima	2.004(0.186)***	0.744 (0.104)***
R^2	0.765	0.578

4.4 In Pursuit of Explaining Search Variance

We also compute and compare linear mixed models, to assess how much fluctuation in search performance is attributable to fractal complexity in the space of local optima. To this end, we use the probabilistic dimensions of the LONs, alongside dimensions calculated by traditional box-counting, to help explain variance seen in the number of iterations the search algorithms require to find the global optimum. A summary of the models for this purpose is found in Table 2: coefficient estimates are given, alongside indication of p-value and the adjusted R^2 for the model. Predictors used are as follows: the probabilistic dimension with β set at 0.90 (*Probabilistic*

*Dim*¹) and 0.96 (*Probabilistic Dim*²), respectively; the deterministic dimension with a strict value for ϵ (*Deterministic Dim*¹) and with a lenient value (*Deterministic Dim*²), and the number of local optima (*Optima*).

The same format is used in Table 3, except the probabilistic dimensions have been substituted for fractal dimensions taken from the multifractal spectra of the LONs. *Multifractal Dim*¹ and *Multifractal Dim*² are dimensions taken from arbitrary points on the fractal spectrum obtained by the Sandbox algorithm.

Table 2 shows that the predictors in the models are able to explain ~64% and ~42% of variance seen in the runtime of the ILS and ROTS, respectively. This is reflected in the R^2 values.

Surveying and comparing the coefficient estimates, we can consider the emboldened text first. In both models, the strongest fractal characteristic of the four is the feature named as *Probabilistic Dim*², with is the fractal box-counting with β set at 0.96, i.e. two nodes boxed together must have a link weight greater than 0.04 in the local optima network. The p-values indicate that this is only of statistical significance in the ILS model, and not the ROTS.

The coefficient for this probabilistic fractal dimension predictor is positive, implying an increase in the response variable (iterations to the global optimum) with increased dimension of this type.

An interesting observation is available when comparing the coefficient estimates and standard errors for the two probabilistic dimension predictor variables. Both are obtained using edge weight as the scaling factor, but are differentiated by their setting of the

allowed threshold β . While the *Probabilistic Dim*² feature had $\beta = 0.96$, *Probabilistic Dim*¹ had $\beta = 0.90$. If we recall that the LON edge weightings were taken as their opposite ($w_{ij} = 1 - w_{ji}$), i.e. the probability of the search transition *not* happening, it follows that a lower value for β is a stricter boxing condition.

The two probabilistic fractal features have very different coefficient estimates and associated standard error. The dimensions obtained using the stricter criterion appear to have little-to-no effect on the response variables. The error rates are very large.

From a comparative perspective, it seems clear that *Probabilistic Dim*² is more important than the deterministic dimensions as a predictor for metaheuristic runtime. This can be seen comparing the standard errors and p-value indications of the probabilistic and deterministic dimensions in the model summaries.

Table 3 reflects mixed models with multifractal dimension points as predictors, as well as the more traditional monofractal features. Comparing the R^2 values of Table 2 and Table 3, we can see that the models using multifractal dimensions are stronger than the probabilistic dimensions, with $\sim 77\%$ and $\sim 58\%$ of variance explained in the multifractal models.

Immediately on viewing the table, we can see that *Multifractal Dim*¹ is, for both metaheuristic algorithms, the variable with the strongest effect of the four fractal characteristics studied. The corresponding p-values are low, indicating a significant relationship in both cases. The coefficients are positive: a larger dimension connects to a larger runtime for the search algorithms. When viewing these results, it is crucial to remember the multifractal dimensions are taken from arbitrary points on a larger spectrum.

All other predictors in the model (except the landscape ruggedness) are not effective. They exhibit large error rates and weak coefficient estimates; this is an interesting phenomenon, because it implies that the fractal dimension is different depending on which part of the network you are looking at. If the networks displayed the same (or similar) scaling behaviour across the board, the expectation would be that the contributions and error rates for the dimensional predictors would be more uniform. This also implies that dimensionality within local optima connectivity is not always important to metaheuristic search: it depends on the position in the fitness landscape.

5 CONCLUSION

We have conducted a thorough study on the fractal geometry of the local optima space in fitness landscapes. We considered *sampled* LONs for the first time (extracted from benchmark QAPLIB problems), and also some *fully enumerated* LONs for smaller problems.

Two new approaches for fractal analysis of LONs were proposed: a fine-grained approach for calculating the probabilistic FD of a LON, and multifractal analysis. The fractal characteristics obtained through these were linked to the search performance of two competitive metaheuristics for the QAP (ILS and ROTS).

The fractal characteristics were contrasted with the performance of two competitive metaheuristic algorithms on the problem instances.

The *probabilistic* fractal dimensions, and some of the dimensions taken from the *multifractal* spectra, were shown to be connected to slower search times.

We saw that the the *Chr* QAPLIB LONs studied appeared to be monofractal (able to be characterised by a single fractal dimension), but that “real-like” LONs and *Nug* QAPLIB instances exhibited multifractality. Therefore, it seems that for some local optima networks a single (mono)fractal dimension is not sufficient to capture the heterogeneous search dynamics encoded in them. Instead, a spectrum of dimensions, as we have calculated here, gives more information.

Finally, we have shown that probabilistic dimensions are more significantly correlated with search than deterministic ones, and therefore provide a more accurate picture of the complexity in local optima connectivity patterns when taking into account search probability.

ACKNOWLEDGMENTS

The networks used during this research are from [1] (fully enumerated) and [3] (sampled), respectively. This work was supported by the UK’s Engineering and Physical Sciences Research Council [grant number EP/J017515/1].

REFERENCES

- [1] Fabio Daolio, Marco Tomassini, Sébastien Vétel, and Gabriela Ochoa. 2011. Communities of minima in local optima networks of combinatorial spaces. *Physica A: Statistical Mechanics and its Applications* 390, 9 (2011), 1684–1694.
- [2] Shuhei Furuya and Kousuke Yakubo. 2011. Multifractality of complex networks. *Physical Review E* 84, 3 (2011), 036118.
- [3] David Iclanzan, Fabio Daolio, and Marco Tomassini. 2014. Data-driven local optima network characterization of QAPLIB instances. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 453–460.
- [4] Witold Kinsner. 2005. A unified approach to fractal dimensions. In *Cognitive Informatics, 2005 (ICCI 2005). Fourth IEEE Conference on*. IEEE, 58–72.
- [5] Joshua Knowles and David Corne. 2003. Instance Generators and Test Suites for the Multiobjective Quadratic Assignment Problem. In *Evolutionary Multi-Criterion Optimization, Second International Conference, EMO 2003, Faro, Portugal, April 2003, Proceedings (LNCS)*, Carlos Fonseca, Peter Fleming, Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele (Eds.). Springer, 295–310.
- [6] Jin-Long Liu, Zu-Guo Yu, and Vo Anh. 2015. Determination of multifractal dimensions of complex networks by means of the sandbox algorithm. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 25, 2 (2015), 023103.
- [7] Benoit B Mandelbrot. 1972. Possible refinement of the lognormal hypothesis concerning the distribution of energy dissipation in intermittent turbulence. In *Statistical models and turbulence*. Springer, 333–351.
- [8] Stephanie Rendón de la Torre, Jaan Kalda, Robert Kitt, and Jüri Engelbrecht. 2017. Fractal and multifractal analysis of complex networks: Estonian network of payments. *The European Physical Journal B* 90, 12 (2017), 234. <https://doi.org/10.1140/epjb/e2017-80214-5>
- [9] Chaoming Song, Shlomo Havlin, and Hernán A. Makse. 2005. Self-similarity of complex networks. *Nature* 433, 7024 (2005), 392–395. <https://doi.org/10.1038/nature03248>
- [10] Yu-Qin Song, Jin-Long Liu, Zu-Guo Yu, and Bao-Gen Li. 2015. Multifractal analysis of weighted networks by a modified sandbox algorithm. *Scientific reports* 5 (2015).
- [11] Peter F Stadler. 2002. Fitness landscapes. In *Biological evolution and statistical physics*. Springer, 183–204.
- [12] Thomas Stützle. 2006. Iterated local search for the quadratic assignment problem. *European Journal of Operational Research* 174, 3 (2006), 1519–1539.
- [13] Éric Taillard. 1991. Robust taboo search for the quadratic assignment problem. *Parallel computing* 17, 4-5 (1991), 443–455.
- [14] Sarah L. Thomson, Sébastien Verel, Gabriela Ochoa, Nadarajan Veerapen, and Paul McMenemy. 2018. On the Fractal Nature of Local Optima Networks. In *Evolutionary Computation in Combinatorial Optimization*, Arnaud Liefvooghe and Manuel López-Ibáñez (Eds.). Springer International Publishing, Cham, 18–33. https://doi.org/10.1007/978-3-319-77449-7_2