

Additional Dimensions to the Study of Funnels in Combinatorial Landscapes

Gabriela Ochoa and Nadarajen Veerapen
Computing Science and Mathematics
University of Stirling
Stirling, FK9 4LA, Scotland
{goc,nve}@cs.stir.ac.uk

ABSTRACT

The global structure of travelling salesman’s fitness landscapes has recently revealed the presence of multiple ‘funnels’. This implies that local optima are organised into several clusters, so that a particular local optimum largely belongs to a particular funnel. Such a global structure can increase search difficulty, especially, when the global optimum is located in a deep, narrow funnel. Our study brings more precision (and dimensions) to the notion of funnels with a data-driven approach using Local Optima Networks and the Chained Lin-Kernighan heuristic. We start by exploring the funnel ‘floors’, characterising them using the notion of *communities* from complex networks. We then analyse the more complex funnel ‘basins’. Since their depth is relevant to search, we visualise them in 3D. Our study, across a set of TSP instances, reveals a multi-funnel structure in most of them. However, the specific topology varies across instances and relates to search difficulty. Finally, including a stronger perturbation into Chained Lin-Kernighan proved to smooth the funnel structure, reducing the number of funnels and enlarging the valley leading to global optima.

CCS Concepts

- **Computing methodologies** → Discrete space search;
- **Mathematics of computing** → *Randomized local search*;
- **Human-centered computing** → *Visualization techniques*;

Keywords

Local Optima Networks; Travelling Salesman Problem; Hill Climbing; Local Search; Fitness Landscapes

1. INTRODUCTION

The global structure of combinatorial fitness landscapes is little understood, partly due to the lack of tools for exploring its complexity on realistic search spaces. The big-valley hypothesis [3] holds that TSP and other combinatorial problems are highly multi-modal, but local optima

are not randomly distributed, instead they tend to cluster around the global optimum conforming a globally convex structure. Therefore, a decrease in fitness (when minimising) implies that, on average, search is getting closer to the global optimum. Problems that exhibit this characteristic are sometimes referred to as single-funnel landscapes. However, recent studies on TSP landscapes have revealed a more complex picture [8, 22, 21]. The big valley seems to decompose into several sub-valleys or multiple funnels. A similar multi-funnel structure has been observed on some continuous optimisation problems [14, 15, 16], where its impact on search difficulty has been established. In particular, landscapes with more than one funnel, where the global optimum is located in a deep, narrow funnel are significantly harder. However, the literature on characterising the multi-funnel structure of combinatorial landscapes is mostly lacking.

We use the local optima networks model to analyse and visualise the global structure of combinatorial fitness landscapes. Local optima networks compress the whole search space into a graph, where nodes are local optima and edges are transitions among them with a given search operator [20, 5, 10]. The model emphasises the number, distribution and most importantly, the connectivity pattern of local optima. Modelling landscapes as networks introduces a new set of metrics to analyse fitness landscapes and the possibility of visualising them [11, 19]. Most previous studies on local optima networks fully enumerate local optima, and therefore consider small problem instances of up to size 30 for binary spaces and 10 for permutation spaces. A data-driven approach is proposed in [10] where quadratic assignment instances of up to size 32 area analysed.

This contribution considers instead TSP instances of several hundred cities (in the range of 500 to 700). A powerful sampling procedure based on the Chained Lin-Kernighan heuristic [17] is implemented, which allows us to collect tens of thousands of local optima on each of the studied landscapes. This data is analysed and visualised using the local optima network model. Our main goal is to bring more rigour to the characterisation of funnels on realistic combinatorial spaces. We propose new 3D visualisation techniques to assist our understanding. Finally, we explore ways in which such understating can inspire algorithmic design.

2. WHAT IS A FUNNEL?

The notion of ‘funnel’ remains vague. The intuition is captured by Figure 1 where two funnels are depicted as two groups of local optima which are close in configuration space within a group, but well-separated between groups. Con-

ducting a complete survey of the concept of funnels and their characterisation is beyond the scope of this article. Our intention is instead to introduce key concepts and tools from theoretical chemistry and optimisation, respectively.

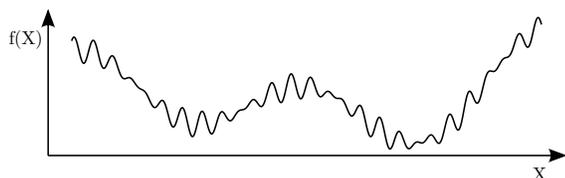


Figure 1: Depiction of two funnels.

2.1 Funnels in Energy Landscapes

The notion of energy landscapes (or energy surfaces) provides a common ground for theories of structure, dynamics and thermodynamics in molecular science [26]. An energy landscape is a mapping of all possible conformations of a molecular entity (clusters, glasses and proteins), to their corresponding energy levels. The structural and dynamic properties of these entities are known to be related to the topographical features of their energy landscapes. An extensive body of work has been devoted to characterising the energy landscapes of molecular clusters, glasses and proteins [6, 26, 7].

The term ‘funnel’ was introduced in the protein folding community to describe “a region of configuration space that can be described in terms of a set of downhill pathways that converge on a single low-energy structure or a set of closely-related low-energy structures.” [6]. It has been suggested that the energy landscape of proteins is characterised by a single deep funnel, a feature that underpins their ability to fold to their native state. In contrast, some shorter polymer chains (polypeptides) that misfold, are expected to have other funnels that can act as traps.

The energy landscapes of so called *atomic clusters* (spatial arrangements of atoms) have been widely studied, as they represent a convenient benchmark of controllable complexity for which excellent putative global minima have been tabulated. In particular, *Lennard-Jones* clusters have become an archetypal test problem for global optimisation algorithms that specialise in finding the lowest energy molecular configuration, represented as the atoms’ spatial positions. The energy landscape of Lennard-Jones clusters is highly multimodal, the complexity of the landscape in terms of the number of local optima can be controlled through the cluster size. Two types of difficulty can be distinguished. First, there is the general increase in the number of optima with cluster size. Second, there are size-specific effects related to the topography of the landscape. For example, the landscape of the 55-atom Lennard-Jones cluster has a single deep funnel leading to the global minimum. On the other hand, the 38-atom Lennard-Jones cluster has a paradigmatic double-funnel topography, where the dominant funnel leads down to a sub-optimal structure and a second narrower, less accessible, funnel leads down to the global optimum. These double-funnel examples are one or two orders of magnitude harder to optimise [6]. In order to characterise energy landscapes, the notions of optima basins and *barriers* between them have been studied. Approaches to elucidate the basin structure have led to the concept of *disconnectivity graphs*

[6, 26], also known as *barrier trees* [7]. To construct these graphs, a database of local optima and *transition states* connecting them is required. Transition states refer to intermediate configurations at which the potential energy has a maximum value, they are also called *saddle points* in the study of fitness landscapes. To visualise the graphs, which are actually trees, local optima are identified with leaves, while the interior nodes represent saddle points separating groups of local optima.

2.2 Funnels in Fitness Landscapes

Energy landscapes in theoretical chemistry and fitness landscapes in optimisation are conceptually related, as has been already observed [25]. This relationship is particularly close for continuous optimisation. Indeed, Müller and Sbalzarini [18] propose including atomic cluster problems as part of black-box benchmark suites. For combinatorial optimisation, a discrete version of the surface, and therefore, different techniques are required.

2.2.1 Continuous optimisation

Locatelli [14] studied the sources of difficulty in continuous optimisation and finds that it is not strictly related to the number of local optima, but to how *chaotic* their positions are. He suggests that problems are structured in multiple levels. At each level, different ‘objects’ are observed, but all levels display a similar structure. The objects are hierarchies of neighbourhoods related to distances among solutions. The number of levels that can be recognised in a problem is a more indicative measure of difficulty than the number of local optima. Lunacek and Whitley [15] propose a metric, *dispersion*, that quantifies the proximity of the best regions in the search space. A high dispersion metric indicates the presence of multiple funnels. In a follow up work [16], the authors studied abstract landscapes with two funnels and find that evolutionary algorithms mostly fail when the global optimum is located in a proportionally smaller funnel.

2.2.2 Combinatorial optimisation

The literature is much more scarce for discrete search spaces, where studies mostly deal with small examples. Disconnectivity graphs (barrier trees) have been applied to discrete optimisation problems where the notions of local optima, basins and saddle points are clearly defined. Flamm et al. [7] extended these definitions so that barrier trees can be constructed for highly degenerate problems (i.e landscapes with neutrality). They present empirical results for binary strings of up to length 10. Hallam and Prügel-Bennett [9] construct barrier trees for MAX-SAT problems with up to 40 variables using branch-and-bound to find only the best local optima in the space. Daolio et al [5] studied the community structure of local optima networks on two classes of instances of the quadratic assignment problem. The two problem classes give rise to different configuration spaces, with the so-called real-like instances revealing a modular structure. The approach is based on a full enumeration of local optima, therefore, instances of size up to 10 were analysed. In a follow up work with a data-driven approach, the modularity of instances up to size 32 is studied [10]. Although not mentioned in these articles, we argue that there is a connection between groupings (communities) of local optima and the notion of funnels.

3. LOCAL OPTIMA NETWORKS FOR TSP

To construct the networks we need to define their nodes and edges. The definition is closely related to the methodology for extracting the network data, which is based on a number of runs of the Chained-LK algorithm described below. Clearly, a full enumeration of the local optima for TSP instances of non-trivial size becomes unmanageable. Therefore, networks are constructed from a sample of high-quality local optima in the search space. We provide below some preliminaries, basic definitions, and details of the sampling algorithm.

3.1 Chained Lin-Kernighan

Lin-Kernighan (LK) [13], is a powerful and well-known heuristic for solving the TSP. For about two decades, it was the best local search method, and nowadays it is a key component of state-of-the-art TSP solvers. LK search is based on the idea of k -exchanges: take the current tour and remove k different links from it, which are then reconnected in a new way to achieve a legal tour. A tour is considered to be ‘ k -opt’ if no k -exchange exists which decreases its length. LK applies 2, 3 and higher-order k -exchanges. The order of a change is not predetermined, rather k is increased until a stopping criterion is met. Thus many kinds of k -exchanges and all 3-exchanges are included. There are many ways to choose the stopping criteria and the best implementations are rather involved. Here, we use the implementation available in the Concorde software package [1], which uses *do not look* bits and candidate lists.

The overall tour-finding strategy using LK-search was to repeatedly start the basic LK routine from different starting points keeping the best solution found. This practice ended in the 1990s with the seminal work of Martin, Otto and Felten [17], who proposed the alternative of *kicking* (perturbing slightly) the LK tour and reapplying the algorithm. If a better tour is produced, the old LK tour is discarded and the new one kept. Otherwise, the search continues with the old tour and kicks it gain. This simple yet powerful strategy is nowadays best known as *iterated local search*. It was named *Chained Lin-Kernighan* (Chained LK) by Applegate et al. [2], who also provided an improved implementation to solve large TSP instances. The kick or escape operator in Chained-LK is a type of 4-exchange, named *double-bridge* by Martin et al. [17]. It consists of two improper 2-exchanges, each of which is a ‘bridge’ as it takes a legal, connected tour into two disconnected parts. The combination of both bridges, must then be chosen in order to produce a legal final tour.

3.2 Model description

Local Optima. A tour is a *local optimum* if no tour in its neighbourhood is shorter than it. The neighbourhood is imposed by LK-search, which considers variable values of k . The local optimality criterion is, therefore, rather stringent. Only a small proportion of all possible tours are LK-optimum. The set of local optima is denoted by LO .

Escape Edges. Edges are directed and based on the double-bridge operator. There is an *escape* edge from local optimum LO_i to local optimum LO_j , if LO_j can be obtained after applying a double-bridge kick to LO_i followed by LK-Search. The set of escape edges is denoted by E_{esc} .

Local Optima Network (LON). This is the graph $LON = (LO, E_{esc})$ where nodes are the local optima LO , and edges

E_{esc} are the escape edges. Edges are directed and weighted. The weight indicates the probability of transition and is in this work empirically estimated.

3.3 Gathering network data

To extract the network data, we instrumented the Chained-LK implementation of Concorde (see Algorithm 1). We simply store, in LO , every unique local optima obtained after an LK application, and create and store, in E_{esc} , an edge between the starting and end optima after a double-bridge move. We also keep count of the number of times and edge was visited, this number is stored as the edge’s weight.

Data: I , a TSP instance

Result: LO , the set of local optima,

E_{esc} , the set of edges between local optima

$n \leftarrow \text{numberOfCities}(I); LO \leftarrow \{\}; E_{esc} \leftarrow \{\}$

for $i \leftarrow 1$ **to** 1000 **do**

$s \leftarrow \text{initialSolution}()$

$s \leftarrow \text{LK}(s)$

$LO \leftarrow LO \cup \{s\}$

for $k \leftarrow 1$ **to** $10 \times n$ **do**

$s_{start} \leftarrow s$

$s_{end} \leftarrow \text{applyKick}(s)$

$s_{end} \leftarrow \text{LK}(s_{end})$

$LO \leftarrow LO \cup \{s_{end}\}$

$E_{esc} \leftarrow E_{esc} \cup \{(s_{start}, s_{end})\}$

if $\text{fitness}(s_{end}) < \text{fitness}(s_{start})$ **then** $s \leftarrow s_{end}$

end

end

Algorithm 1: Local optima network sampling combining 1000 runs of Chained-LK.

A thousand independent runs of Chained-LK are executed for each TSP instance. We chose to use two different starting mechanisms, one producing “better” solutions, the other “worse” solutions, to have a broader picture of the search space. Half of the runs start from a relatively good solution, built using the Quick-Borůvka method. The latter is the default initialisation for Concorde’s Chained-LK and is based on the minimum-weight spanning tree algorithm of Borůvka [2]. The other half starts from a random solution. Each run performs i iterations or kicks, where i is ten times the size of the tour ($10 \times$ the number of cities). The default kicking procedure in Concorde is used: the edges involved in the double-bridge are selected using random walks along connected vertices.

3.4 TSP instances

Our goal is to visualise and characterise in detail the global structure of TSP fitness landscapes. We consider 8 instances of moderate size (from 500 to 700 cities) and different types. The first four instances are randomly generated using the DIMACS TSP instance generator¹. Sizes of 570 and 670 are chosen to roughly match the size of the selected TSPLIB instances. The instances starting with ‘E’ are composed of uniformly distributed cities, while a ‘C’ indicates that the cities are clustered. The “.0” suffix in the instance name indicates, as per DIMACS convention, a seed of 1000. The second group are well-known instances taken from TSPLIB [23]. A popular way of constructing TSP instances is to choose a set

¹dimacs.rutgers.edu/Challenges/TSP/download.html

of actual cities and to define the cost of travel between any two cities as the distance between them. This is the case of instances att532 (532 US cities) and gr666 (666 World cities). The other two instances: u574 and d657, arise from the task of drilling 574 and 657 holes, respectively, in printed circuit boards. Table 1 summarises the TSP instances studied, including results obtained when running the Concorde solver.

Table 1: TSP Instances: Name with number of cities as suffix, edge type and features resulting from running the Concorde solver.

Instance	Edge Type	Concorde solver		
		Optimum	Run time	B&B nodes
E570.0	EUC-2D	17 832 234	55.4	49.7
E670.0	EUC-2D	19 316 839	15.3	10.3
C570.0	EUC-2D	8 900 015	2.2	1.0
C670.0	EUC-2D	8 908 620	13.2	6.1
att532	ATT	27 686	8.9	5.2
gr666	GEO	294 358	6.5	3.2
u574	EUC-2D	36 905	3.8	1.7
d657	EUC-2D	48 913 ²	23.1	14.8

4. NETWORK ANALYSIS

The data collection process combines a thousand Chained-LK runs as described in Section 3.3. This section analyses and visualises the networks constructed from the data. The number of unique local optima gathered ranges from 1.3 to 2.7 million nodes on the studied instances. A close inspection revealed that over 98% of the local optima produced are failed escape attempts. That is, solutions generated with a double-bridge move, which were not accepted by the search process as their cost was higher than the incumbent solution. We do not consider those nodes when constructing the networks as they do not impact the search dynamic.

Table 2: Local Optima Networks. Properties: number of different global optima go ; percentage of runs reaching a global optimum % $succ$. Network metrics: number of vertices n ; number of edges m ; mean incoming vertex strength \hat{s} , average path length to a global optimum \hat{l} .

	go	% $succ$	n	m	\hat{s}	\hat{l}
E570.0	1	38.7	18 052	37 083	127.83	7.53
E670.0	1	43.5	24 203	46 487	120.81	9.84
C570.0	1	95.3	18 790	34 904	122.38	3.94
C670.0	1	86.7	28 108	51 524	83.27	4.91
att532	2	37.7	24 079	52 413	88.07	6.26
gr666	2	12.8	32 945	61 168	70.66	9.61
u574	4	41.1	27 959	55 210	83.84	6.58
d657	8	20.3	32 669	73 210	69.42	9.74

²The best fitness reported for d657 in TSPLIB is 48 912, however we find 48 913 when using Concorde on both our machine and NEOS Server (www.neos-server.org). The TSPLIB FAQ notes that “due to round-off errors distance evaluations may slightly differ on different machines.”

Table 2 summarises basic statistics and network metrics from the studied instances. There are clear differences between the random and the structured instances. All the random instances have a single global optimum, whereas the structured instances have more than one, which indicates that these landscapes have neutrality. The percentage of runs that find a global optimum % $success$, is lower the larger the number of cities for all instance types except the uniform random ‘E’ instances, where the smaller instance is harder to solve (as is also revealed in Table 1 for the Concorde solver), despite having the smallest number of local optima. The highest success percentage is by far that of the clustered random ‘C’ instances. The lower search difficulty of these instances is also revealed by their lower average path length to hit a global optima \hat{l} when such path exists. The number of edges is approximately twice the number of nodes in all cases, indicating that the sampled networks are relatively sparse. The random instances have a smaller number of local optima and a higher incoming vertex strength (sum of the edge weights of the incoming adjacent edges of a vertex) as compared to the structured TSPLIB instances. Once the data is filtered and the networks constructed, we proceed to characterise and analyse the funnel *floors*, which are easier to define and clearly organise into clusters, and then consider the more complex funnel *basins*.

4.1 Funnel floors

A *funnel floor* solution is a high quality local optimum that is conjectured to be at the bottom of a funnel. Funnel floors are generated by running Chained-LK for a large enough number of iterations. They are recognised by the lack of downward progress in a large number of double-bridge move attempts. As indicated by the first two columns in Table 3, the numbers of unique fitness evaluations u_{fit} and unique solutions u_{sol} increase with the problem size in all instance types. However, for the uniform random instances the increase is small. Notice that for all random instances, the number of unique evaluations matches the number of unique solutions. This is not the case for the structured instances, where several different solutions have the same fitness value.

Table 3: Funnel Floors. Properties: number of unique fitness evaluations u_{fit} , and number of unique solutions u_{sol} . Network metrics: number of vertices n ; number of edges m ; number of weak connected components cp ; number of communities cm and modularity metric mod .

	u_{fit}	u_{sol}	n	m	cp	cm	mod
E570.0	24	24	5	9	3	3	0.50
E670.0	27	27	3	3	3	3	0.65
C570.0	6	6	1	1	1	1	0
C670.0	12	12	1	1	1	1	0
att532	10	73	42	299	1	8	0.78
gr666	64	114	8	21	1	4	0.74
u574	11	44	24	138	1	5	0.78
d657	56	325	44	330	5	6	0.81

To analyse the connectivity pattern of the funnel floors, we induced a sub-graph of the local optima network where the nodes are those in the funnel floors. We only included

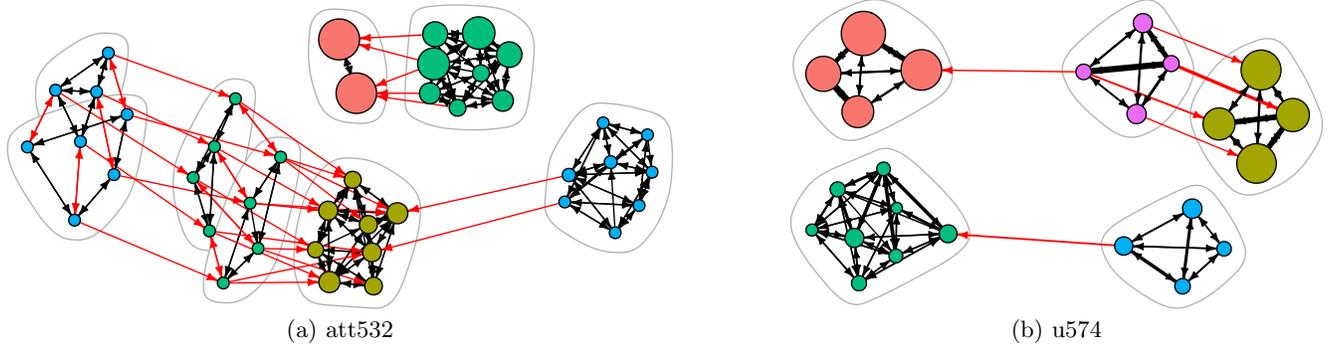


Figure 2: Funnel floor networks for two selected instances. Nodes are local optima in the funnel floors and edges represent escapes with double-bridge moves. Node colours identify the top 5 fitness levels as indicated in Table 3. Node sizes are proportional to their strength (weighted incoming degree), and edges width to their weight. The groups outlined identify communities. We argue communities correspond to funnels. Black edges indicate connections within communities, and red across communities.

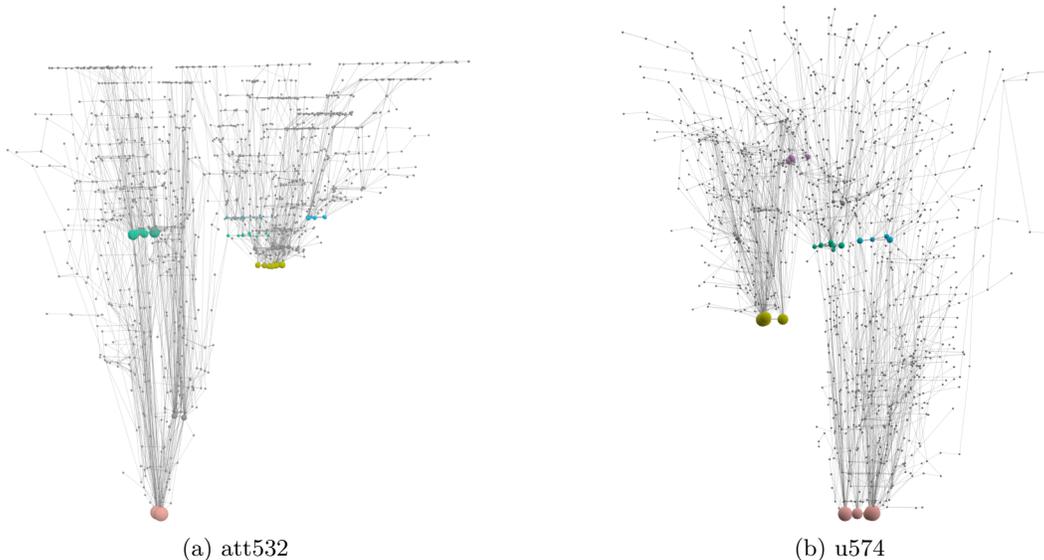


Figure 3: Local optima networks for two selected instances. Node colours and sizes are as indicated in Figure 2. Edges widths are constant. The 3D network models are oriented to roughly display the x - z plane in 2D, where z is the fitness.

fitness levels obtained in above 5% of runs as these correspond to the areas of the search space more frequently visited. Columns 3 to 5 in Table 3 show some basic metrics for these networks. The random instances have very small funnel floor networks, with the clustered cities problems trivially conformed by a single node (the global optimum). In contrast, the TSPLIB instances revealed a complex funnel floor structure with several tens of nodes and a number of groupings or communities. Table 4 gives further details on the structure of the floors of two selected TSPLIB instances, namely att532 and u574. The number of unique solutions and the percentage of runs attaining the top 5 fitness levels are reported. Notice that the number of unique solutions on each level is generally greater than 2 and can be as high as 16 for att532. This is evidence of symmetries in TSP instances producing plateaus of solutions. Once a solution in the fitness level is reached, the search process traverses the

plateau finding different solutions with the same fitness.

Figure 2 visualises the funnel floor local optima networks for instances att532 and u574. The graph visualisation uses the so-called force-directed layout, as implemented in the `igraph` package [4]. In order to bring more rigour to the global structure visualised, we resort to the notion of *communities* in complex networks. A community is a group of nodes that have a higher likelihood of connecting to each other than to nodes from other communities. More formally, a community is a locally dense connected sub-graph in a network. Community detection is related to *graph partitioning* in computer science. There is, however, an important difference between two: graph partitioning divides a network into a predefined number of smaller sub-graphs. In contrast, community detection aims to uncover (in an unsupervised fashion) the inherent community structure of a network. Many community detection algorithms have been

Table 4: Funnel floors for two selected TSPLIB instances. Top 5 fitness levels obtained at the end of 1000 Chained-LK runs. Fitness levels obtained in less than 5% of the runs are not included.

	Global Opt.		Fitness Levels		
att532	27 686	27 703	27 705	27 706	na
Unique	2	8	16	16	-
%	37.7	26.5	22.7	6.8	-
u574	36 905	36 935	36 944	36 945	36 962
Unique	4	4	8	4	4
%	41.1	35.2	9.6	7.5	5.0
Colour	■	■	■	■	■

proposed, in most of them the number and the size of the communities is not predefined, but needs to be discovered by inspecting the network connectivity pattern [12, 11]. Following a recent comparative analysis [12], we selected the *Infomap* method by Rosvall and Bergstrom [24] as it was the best performing on the wide set of benchmarks considered. It is suitable for our study as it considers directed and weighted networks. Moreover, it is a *flow-based* method, which considers the dynamics on the network rather than its topological structure, which makes it relevant to the study of search heuristics. The last column in Table 3 reports the modularity metric of the funnel floor networks for all instances. This metric is designed to measure the strength of division of a network into modules (also called groups, clusters or communities). The random instances reveal lower modularity than the structured instances, with the trivial networks with a single node having zero modularity.

The plots in Figure 2 reveal a clear organisation into groups of local optima. Node colours identify the different fitness levels as indicated in Table 4. Node sizes reflect their incoming strength, which account for the sum of weights of the incoming edges. The groups outlined identify communities according to the *Infomap* algorithm. Edges width reflect the weight of the links, which account for transition probabilities. Black edges indicate connections within communities, and red across communities. It is interesting to note that having the same fitness value does not automatically mean that nodes belong to the same community (putative funnel floor). As Figure 2 (a) for att532 illustrates, for both fitness levels 27 705 in green and 27 706 in blue, there is a clearly distinguished larger community and two additional smaller communities nearby the yellow community. For the u574 (Fig. 2 (b)), the communities seem to correspond to the fitness levels. We argue that the communities detected correspond to the funnel floors in the landscape.

4.2 Funnel basins

So far, visualisations of local optima in the literature have used 2D coordinate systems, sometimes with colour or size providing an additional dimension. We propose a 3D visualisation where the x and y coordinates are, as usual, determined by a graph layout algorithm; the innovation is to use fitness as the z coordinate. This provides a clearer representation of the funnel and basin concepts as shown in Figure 3, depicting the 3D LONs for att532 and u574. The 3D representation brings an almost tangible aspect to the widely used search landscape metaphor.

In order to obtain an image compatible with the con-

straints of printing and page limits, networks are pruned by removing nodes whose tour cost is below a threshold value, selected as sufficiently high to include all the funnel floors encountered. In practice, we used a threshold of 10 units above the fitness of the highest floor. Nodes with strength below the median were also removed. This pruning results in some nodes that appear to stick out of the network but in the original network there are low strength nodes that lead to them.

5. SMOOTHING FUNNELS

Double-bridge is the kick operator in Chained-LK, however its perturbative power is not always the most appropriate. The kick operator defines the connectivity pattern of local optima networks and thus the global structure of the induced landscape. We modified the Concorde Chained-LK implementation to perform p consecutive double-bridges instead of a single one after l non-improving iterations. All the pairwise combinations of $p \in \{2, 3, 4\}$ and l from 0 to 300 with increments of 50, were tested across 1000 runs. In all instances, with the exception of the random instances with clustered cities ‘C’, $p > 1$ produced a better success percentage than $p = 1$. Surprisingly, the value $l = 0$ was always the best performing.

Table 5: Funnel Floors with stronger perturbations. Properties: number of double-bridge moves producing the best success p , percentage of runs reaching a global optimum % *succ*, number of unique fitness evaluations u_{fit} , number of unique solutions u_{sol} . Network metrics: number of vertices n ; number of edges m ; number of weak connected components cp ; number of communities cm and modularity metric mod . Underlined values indicate that the new distribution of fitness is significantly different from the original (p -value < 0.001) using the Wilcoxon-Mann-Whitney test and Bonferonni-Holm correction.

	p	% <i>succ</i>	u_{fit}	u_{sol}	n	m	cp	cm	mod
E570.0	2	40.3	15	15	5	7	3	3	0.52
E670.0	<u>3</u>	73.7	28	28	3	3	3	3	0.59
att532	<u>2</u>	45.5	11	86	37	265	2	7	0.75
gr666	2	14.1	58	100	8	19	2	4	0.73
u574	<u>4</u>	54.9	17	47	20	111	2	4	0.72
d657	<u>3</u>	31.2	57	316	16	100	2	3	0.59

Table 5 outlines relevant metrics on the instances where increasing p improved success %. The effect of increasing p seems to be to *smooth* the landscape, that is, to reduce the number of funnels, increasing the size of the funnel basin leading to global optima. This is evidenced by the reduced number of communities cm , and larger success % with increased p in the structured instances (Table 5). This is further supported by Table 6 and Figure 4, which compare the multi-funnel structure for a selected instance: d657 and two values of p . Increasing the perturbation strength p from 1 to 3 reduces the number of funnels from 6 to 3, enlarging the funnel basin leading to global optima. The 3D visualisations in Figure 5 highlight the structural differences between the two sampled networks. In particular, the single double-bridge network shows a clear separation between the

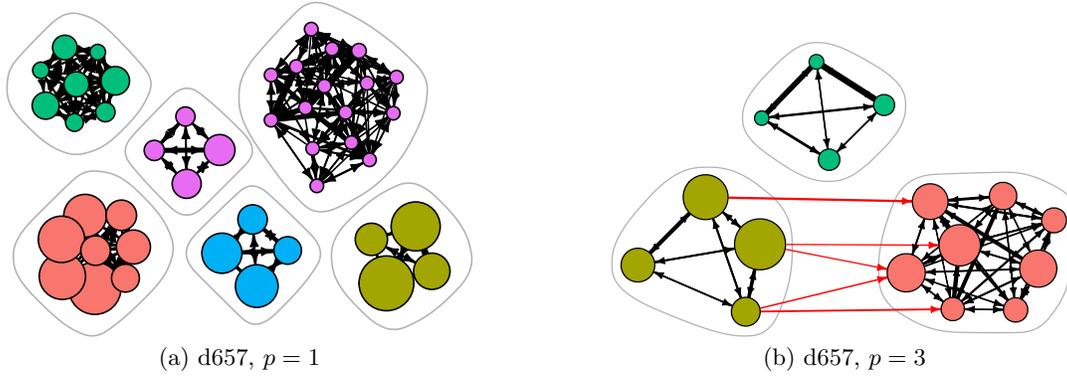


Figure 4: Funnel floor networks for for d657 and two perturbation strengths p . Node colours identify the top 5 fitness levels as indicated in Table 5. Node sizes are proportional to their incoming strength (weighted incoming degree), and edges width to their weight. The groups outlined identify communities. The number of communities (funnels) decreases with increasing p .

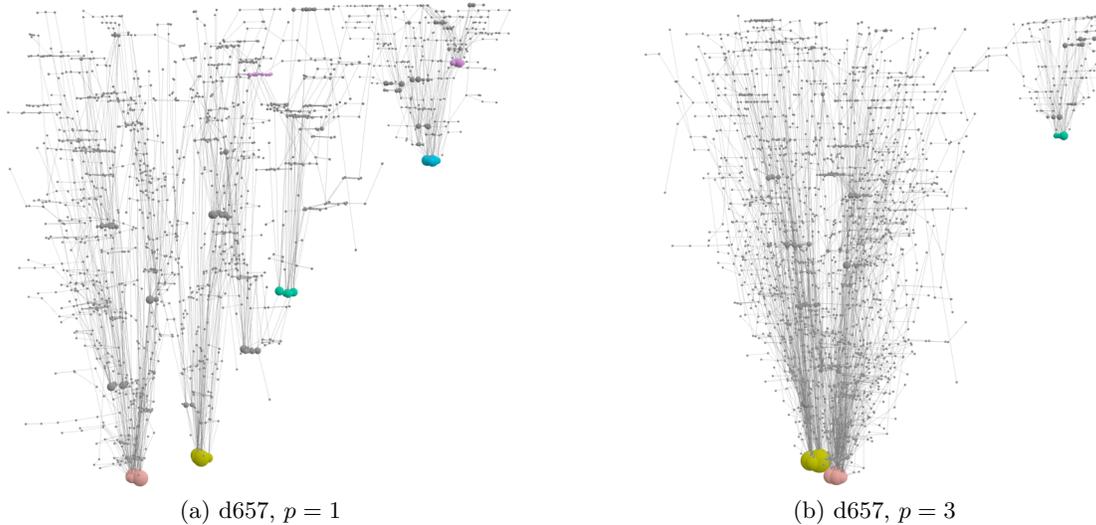


Figure 5: Local optima networks for d657 and two perturbation strengths p . Node colours and sizes are as indicated in Figure 4, edges are visualised with a constant width.

two deepest funnels. On the other hand, the triple double-bridge network shows less funnels, connections between the two deepest funnels and a wider basin leading to the global optima.

6. CONCLUSION

The global structure of combinatorial landscapes has received little attention in the literature. This is due in part to the lack of tools to study their complexity. Local optima networks help to fill this gap. We found evidence of multiple funnels in TSP landscapes of moderate size (500 to 700 cities). Our data-driven approach incorporates the notion of community structure together with a novel 3D local optima network visualisation, bringing new insights into landscapes' global structure. In particular, the 3D plots provide a concrete and intuitive depiction of the fitness landscape metaphor. We found significant differences among the studied instance classes. Randomly generated instances with

uniformly distributed cities are hard to solve due to their multi-modality. There is little clustering of good local optima in this case. In sharp contrast are the random instances with clustered cities, where good local optima are gathered in a single large funnel guiding the search process straight into the unique global optimum. These are the easiest instances to solve with Chained-LK. The structured TSPLIB instances lie somewhere in between with a rich multi-funnel structure. Good local optima decompose into multiple valleys of different depths, each channelling the search process to a separate low cost solution or group of solutions conforming a plateau. Search difficulty in these cases is related to the size and reachability of the funnel basin containing the global optima.

Experiments with a stronger perturbation proved to help in *smoothing* the funnel structure, that is, reducing the number of funnels and making the global optima more reachable. Future work will consider larger instances and additional problems. We will also explore alternative mechanisms such

Table 6: Funnel floors for d657 and two perturbation strengths $p = \{1, 3\}$. Top 5 fitness levels obtained at the end of 1000 Chained-LK runs. Fitness levels obtained in less than 5% of the runs are not included.

	Global Opt.		Fitness Levels		
$k = 1$	48 913	48 915	48 948	48 962	48 980
Unique	8	4	8	4	20
%	20.3	13.9	5.5	12.7	8.5
$k = 3$	48 913	48 915	48 962	na	na
Unique	8	4	4	-	-
%	31.2	22.3	8.5	-	-
Colour					

as crossover to smooth the funnel structure. Enhancing visualisation by adding dynamism and interactivity is part of our goal, as is to make our tools and network data-sets available to the research community.

Acknowledgements

This work was supported by the Leverhulme Trust [award number RPG-2015-395]. All data generated during this research are openly available from the Stirling Online Repository for Research Data (<http://hdl.handle.net/11667/73>). Results were obtained using the EPSRC funded ARCHIE-WeSt High Performance Computer (www.archie-west.ac.uk, EPSRC grant EP/K000586/1).

7. REFERENCES

- [1] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. Concorde TSP solver, 2003.
- [2] D. Applegate, W. Cook, and A. Rohe. Chained Lin-Kernighan for Large Traveling Salesman Problems. *INFORMS Journal on Computing*, 15:82–92, 2003.
- [3] K. D. Boese, A. B. Kahng, and S. Muddu. A new adaptive multi-start technique for combinatorial global optimizations. *Operations Research Letters*, 16:101–113, 1994.
- [4] G. Csardi and T. Nepusz. The igraph software package for complex network research. *InterJournal, Complex Systems*:1695, 2006.
- [5] F. Daolio, M. Tomassini, S. Vérel, and G. Ochoa. Communities of minima in local optima networks of combinatorial spaces. *Physica A: Statistical Mechanics and its Applications*, 390(9):1684–1694, 2011.
- [6] J. P. K. Doye, M. A. Miller, and D. J. Wales. The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *Journal of Chemical Physics*, 110(14):6896–6906, 1999.
- [7] C. Flamm, I. L. Hofacker, P. F. Stadler, and M. T. Wolfinger. Barrier Trees of Degenerate Landscapes. *Phys. Chem.*, 216:155–173, 2002.
- [8] D. R. Hains, L. D. Whitley, and A. E. Howe. Revisiting the big valley search space structure in the TSP. *Journal of the Operational Research Society*, 62(2):305–312, 2011.
- [9] J. Hallam and A. Prugel-Bennett. Large barrier trees for studying search. *IEEE Transactions on Evolutionary Computation*, 9(4):385–397, 2005.
- [10] D. Iclanzan, F. Daolio, and M. Tomassini. Data-driven local optima network characterization of QAPLIB instances. In *Genetic and Evolutionary Computation Conference, GECCO 2014*, pages 453–460, 2014.
- [11] E. D. Kolaczyk. *Statistical Analysis of Network Data: Methods and Models*. Springer Series in Statistics. 2010.
- [12] A. Lancichinetti and S. Fortunato. Community detection algorithms: a comparative analysis. *Physical Review E*, 80(5):56117, 2009.
- [13] S. Lin and B. W. Kernighan. An Effective Heuristic Algorithm for the Traveling-Salesman Problem. *Operations Research*, 21:498–516, 1973.
- [14] M. Locatelli. On the Multilevel Structure of Global optimization problems. *Computational Optimization and Applications*, pages 5–22, 2005.
- [15] M. Lunacek and D. Whitley. The dispersion metric and the CMA evolution strategy. In *Genetic and Evolutionary Computation Conference, GECCO 2006*, pages 477–484, 2006.
- [16] M. Lunacek, D. Whitley, and A. M. Sutton. The impact of global structure on search. In *Parallel Problem Solving from Nature - PPSN X*, volume 5199 of *LNCS*, pages 498–507. Springer, 2008.
- [17] O. Martin, S. W. Otto, and E. W. Felten. Large-step markov chains for the TSP incorporating local search heuristics. *Operations Research Letters*, 11:219–224, 1992.
- [18] C. L. Müller and I. F. Sbalzarini. Energy landscapes of atomic clusters as black box optimization benchmarks. *Evolutionary Computation*, 20(4):543–573, 2012.
- [19] M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, Oxford, UK, 2010.
- [20] G. Ochoa, M. Tomassini, S. Verel, and C. Darabos. A study of NK landscapes’ basins and local optima networks. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2008*, pages 555–562. ACM, 2008.
- [21] G. Ochoa and N. Veerapen. Deconstructing the Big Valley Search Space Hypothesis. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2016*, volume 9595 of *Lecture Notes in Computer Science*, pages 58–73. Springer, 2016.
- [22] G. Ochoa, N. Veerapen, D. Whitley, and E. K. Burke. The Multi-Funnel Structure of TSP Fitness Landscapes: A Visual Exploration. In *Artificial Evolution: 12th International Conference, Evolution Artificielle, EA 2015*, volume 9554 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2016.
- [23] G. Reinelt. TSPLIB—a traveling salesman problem library. *ORSA Journal on Computing*, 3(4):376–384, 1991.
- [24] M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4):1118–23, 2008.
- [25] P. F. Stadler. Fitness landscapes. *Appl. Math. and Comput*, 117:187–207, 2002.
- [26] D. J. Wales. Energy landscapes and properties of biomolecules. *Physical Biology*, 2(4):S86–S93, 2005.