The Multi-Funnel Structure of TSP Fitness Landscapes: A Visual Exploration

Gabriela Ochoa[1], Nadarajen Veerapen[1],
Darrell Whitley[2], and Edmund K. Burke[1]

[1] Computing Science and Mathematics, University of Stirling, Scotland, UK.
[2] Department of Computer Science, Colorado State University, USA.

Abstract. We use the Local Optima Network model to study the structure of symmetric TSP fitness landscapes. The 'big-valley' hypothesis holds that for TSP and other combinatorial problems, local optima are not randomly distributed, instead they tend to be clustered around the global optimum. However, a recent study has observed that, for solutions close in evaluation to the global optimum, this structure breaks down into multiple valleys, forming what has been called 'multiple funnels'. The multiple funnel concept implies that local optima are organised into clusters, so that a particular local optimum largely belongs to a particular funnel. Our study is the first to extract and visualise local optima networks for TSP and is based on a sampling methodology relying on the Chained Lin-Kernighan algorithm. We confirm the existence of multiple funnels on two selected TSP instances, finding additional funnels in a previously studied instance. Our results suggests that transitions among funnels are possible using operators such as 'double-bridge'. However, for consistently escaping sub-optimal funnels, more robust escaping mechanisms are required.

1 Introduction

The structure of combinatorial fitness landscapes is known to impact the performance of heuristic search algorithms. Features such as the number and distribution of local optima and their basins of attraction are among the most studied. The relationship among local optima for the symmetric Traveling Salesman Problem (TSP) under the standard 2-change neighbourhood was first analysed in [4], where a *globally convex* structure was discovered. The global optimum was found to be 'central' to all other local optima conforming a 'big-valley' structure. This is interpreted as a landscape where many local optima exists, but they are easy to escape and the gradient, when viewed at a coarse level, leads to the global optimum (Fig. 1). However, a more recent study has found that the big valley structure breaks down when considering solutions near in evaluation to the global optimum [7]. The big-valley separates into multiple valleys, conforming what has been called 'multiple funnels' in the study of energy surfaces in chemical-physics [19]. The multi-funnel concept implies that local optima are organised into clusters, so that a particular local optimum largely belongs to a
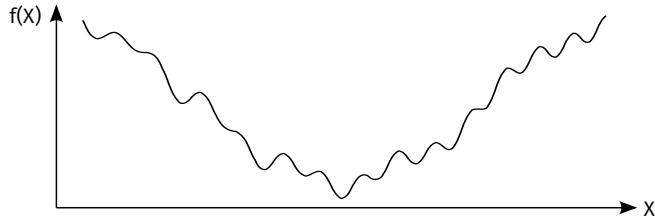
Fig. 1: Depiction of the 'big-valley' structure.

particular funnel. The appearance of multiple funnels explains why certain iterated local search heuristics can quickly find high-quality solutions, but fail to consistently find the global optimum. In a series of studies, Whitley et al. [20, 7, 21] have proposed a crossover operator (Partition Crossover), which has demonstrated the ability to escape funnels at evaluations close to the global optimum. A similar recombination operator [12] is used by Helsgaun [8] in the so called LKH-solver.

This article uses the Local Optima Network (LON) model [15, 16, 14, 18] in order to explore in more detail the structure of TSP landscapes near the global optimum. Local optima networks compress the whole search spaces into a graph having as vertices the local optima, and as edges transitions among them according to a given search operator. This network-based model brings the tools from the *new science of networks* [13] (e.g., metrics and visualisation) to the study of fitness landscapes in combinatorial optimisation.

Our study considers Chained Lin-Kernighan (Chained-LK), one of the best performing heuristic algorithms for TSP [11, 2]. Chained LK is an iterated local search approach combining the variable depth local search of Lin and Kernighan (LK-search) [10] with the *double-bridge* move [11] (a form of 4-change, depicted in Fig. 2b) as the perturbation or 'kick' operator. Therefore, the proposed LON model considers local minima according to LK-search, and transitions among them according to the double-bridge move. Our goal is to gain a deeper understanding of the multi-funnel structure of the TSP under Chained-LK, which will help in selecting and designing stronger escape mechanisms (such as Partition Crossover [20, 21]) to avoid being trapped in a sub-optimal funnel. The main contributions of this article are the following:

1. First study of local optima networks for TSP, including their sampling and analysis.
2. Definition of the DLON model (distance local optima networks) and adaptation of the escape edges model (ELON) to TSP.
3. Network visualisation of the multi-funnel structure of TSP fitness landscapes.
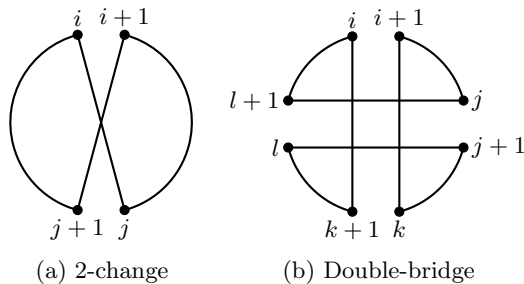
(a) 2-change   (b) Double-bridge

Fig. 2: Illustration of tours obtained after 2-change or double-bridge move.

2   Local Optima Networks for TSP

For a TSP instance with n cities, the search space is the set of permutations of the n cities. The number of tours, which equals the number of permutations, is factorial in n. The fitness function f is given by the length of the tour, which is to be minimised. Before presenting formal definitions in Section 2.1, we briefly describe the following notions relevant to our model.

LK-search: The well-known Lin-Kernighan heuristic is a powerful local search algorithm. It is based on the idea of k-change moves: 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-change exists which decreases its length. Fig. 2a illustrates a 2-change move. LK applies 2, 3 and higher-order k-changes. The order of a change is not predetermined, rather k is increased until a stopping criterion is met. Thus many kinds of k-changes and all 3-changes are included. There are many ways to choose the stopping criteria and the best implementations are rather involved. We use here the implementation available in the Concorde software package [1], which uses *do not look* bits and candidate lists.

Double-bridge move: Proposed by Martin et al. [11] as the 'kick' mechanism in the Chained-LK heuristic, the double-bridge move (drawn in Fig. 2b) is a type of 4-change. It consists of two improper 2-changes, each of which is a 'bridge' (i.e, it takes a legal, connected tour into two disconnected parts). The combination of both bridges, must then be chosen as to produce a legal final tour.

Bond distance: Measures the difference between two tours t_1 and t_2 according to the number of edges or 'bonds' that differ in both tours. Specifically, $b(t_1, t_2)$ is equal to n minus the number of edges that are present in both t_1 and t_2 disregarding edge direction [4].

Our approach requires defining and extracting local optima networks for TSP instances. To construct the networks, we need to define their nodes and edges. Nodes will be local optima according to LK-search, and two types of weighted edges are considered: escape edges and distance edges. The escape edges are

based on the number of double-bridge moves required to escape from a local optimum, while distance edges consider the bond distance between solutions.

Since combinatorial explosion renders the full enumeration of local optima for TSP instances of non-trivial size impossible, we resort to sampling local optima which are close in evaluation to the global optimum. The sampling procedure is further described in Section 2.2.

2.1 Definitions

Definition 1. A *funnel floor* solution is a high quality local optimum that is conjectured to be at the bottom of a funnel. Indeed, they were called funnel bottom solutions in [7], and are generated running Chained-LK for a large enough number of iterations. The set of funnel floor solutions is denoted by F.

Definition 2. A *funnel basin* solution is a local optimum within a funnel. Each funnel basin solution is obtained by first locating a funnel floor, and then escaping from the funnel floor in order to discover a nearby local optimum. In this article, this is done using a random walk with double-bridge followed by improvement using LK-search. The set of local optima defining the funnel basins is denoted by B. Specifically, for some $x \in F$, $y \in B_x \subseteq B$ if it can be obtained from x after a sequence of length d of double-bridge moves followed by LK-search. Since after a double-bridge followed by LK-search the local optimum obtained y can be equal to the starting point x, the length d of the random walk is incremented until $y \neq x$.

The set of local optima, L, is the union of the funnel floors and local optima that define the funnel basins, $L = F \cup B$.

Definition 3. An *escape edge* is a weighted edge from a funnel floor to a local optimum. Specifically, there is an edge $e_{x,y}$ of weight d between the funnel floor point $x \in F$ and the local optima $y \in B$ if y can be obtained from x after a sequence of length d of double-bridge moves followed by LK-search. No self-loops are considered. The set of escape edges is denoted by E_{esc}.

Definition 4. A *distance edge* is a weighted edge, according to the bond distance, between any two local optima. Specifically, there is an edge $e_{x,y}$ of weight d between local optima x and $y \in L$ if the bond distance $b(x,y) = d$. The set of distance edges between any two local optima in L is denoted by E_{dist}.

Definition 5. The *Escape Local Optima Network (ELON)* is the graph $ELON = (L, E_{esc})$ where nodes are the local optima L, and edges E_{esc} are the escape edges.

Definition 6. The *Distance Local Optima Network (DLON)* is the graph $DLON = (L, E_{dist})$ where nodes are the local optima L, and edges E_{dist} are the distance edges.

Data: I, a TSP instance
Result: F, the set of tours on the funnel floors
$\qquad\qquad$ B, the set containing the escape tours from sampled funnel floors
$F \leftarrow \emptyset$;
for $i \leftarrow 1$ **to** $10,000$ **do**
$\quad$ $x \leftarrow$ chainedLK$(I, \text{stallcount} = 10,000)$;
$\quad$ **if** $x \notin F$ **then**
$\quad\quad$ $F \leftarrow F \cup \{x\}$;
$\quad$ **end**
end
$S \leftarrow$ mostFrequentSolutionForEachFitnessLevel(F);
$B \leftarrow \emptyset$;
for $v_0 \in S$ **do**
$\quad$ $B_{v_0} \leftarrow \emptyset$;
$\quad$ **for** $j \leftarrow 1$ **to** $1,000$ **do**
$\quad\quad$ $i \leftarrow 0$;
$\quad\quad$ **repeat**
$\quad\quad\quad$ $i \leftarrow i + 1$;
$\quad\quad\quad$ $v_i \leftarrow$ randomDoubleBridgeMove(v_{i-1});
$\quad\quad\quad$ $v' \leftarrow$ LK(v_i);
$\quad\quad$ **until** $v' \neq v_0$;
$\quad\quad$ $B_{v_0} \leftarrow B_{v_0} \cup \{v'\}$;
$\quad$ **end**
$\quad$ $B \leftarrow B \cup B_{v_0}$;
end

Algorithm 1: TSP local optima network sampling procedure

2.2 Sampling Methodology

We apply a sampling strategy similar to that used by Hains et al. [7] where two stages are considered. This process also resembles the one used by Iclan-zan et al. [9] to sample the landscape of Quadratic Assignment Problem instances. In the first stage, local optima of very good quality are identified which define the funnel floors (set F defined in Section 2.1). In the second stage, random walks are generated to escape these local optima in order to determine the funnels' basins (set B defined in Section 2.1). These approaches are detailed below and through pseudocode in Algorithm 1.

The funnel floor solutions are tours obtained when Chained-LK stalls. In practice, we determine stalling to occur when fitness does not improve for 10,000 consecutive iterations of Chained-LK. This procedure is itself repeated 10,000 times from a randomly generated initial tour and the unique tours produced are saved in F, the set of funnel floor solutions. This procedure corresponds to the first loop in Algorithm 1.

To determine a funnel's basin, we identify a start point in its floor, let us call it v_0, and follow a random walk using a sequence of double bridge perturbations. More precisely, at each step i of the random walk, a random move is performed on v_{i-1}, producing a tour v_i. An LK-search is then applied to v_i to produce a

locally optimal tour v'. If v' is different from v_0, then we have escaped from the basin of attraction of v_0. The random walk is stopped and its length i is the escape distance. Tour v' is saved in B_{v_0}, the set of tours having escaped from v_0. This escape procedure is repeated 1,000 times.

When there are many tours on the funnel floors, it is impractical to try to escape from all of them. When Hains et al. [7] computed the funnels floors from 1,000 Chained-LK applications, they found that tours with the same fitness level formed a connected component under 2-change. These could thus be considered to form a plateau and they, therefore, randomly chose one tour to escape from out of each plateau.

In our case, having performed 10,000 Chained-LK applications, we find many more tours on the funnel floors and, furthermore, they are not all on 2-change plateaus. Our approach selects the most frequently occurring solution within each fitness level as a starting solution. Ties are broken at random.

3 Results

Our study considers two 'milestone' TSP instances: *lin318* and *att532* (as named in TSPLIB [17], also listed in Table 1.5 from [3]). They are composed of 318 and 532 cities, and were first solved to optimality in 1980 and 1987, respectively. The lin318 instance is a circuit board drilling example (i.e., it models the routing of a numerically controlled drilling machine efficiently through a set of hole positions), and was presented by Lin and Kernighan in their seminal paper [10]. It remained the largest TSP instance solved to optimality for a span of seven years in the 1980s. The att532 instance is comprised of pseudo-Euclidean coordinates that go through the 532 largest cities of the USA. It is very well known given the difficulty that the distances to the next node are very short at the east coast, whereas in other regions of the USA they are very long.

Results are discussed in the following two subsections. Section 3.1 analyses the sampled local optima and the bond and escape distances among them. Section 3.2 visualises the escape and distance local optima networks.

3.1 Local Optima and Distances

For instance lin318, 4 unique funnel floor solutions were identified, each with a different fitness level (Table 1). The global optimum was found in the over-whelming majority, 96 %, of cases. The other funnel floor solutions' fitness is within 0.32 % of the global optimum.

When considering att532, 47 unique funnel floor solutions were identified, distributed among 8 different fitness levels (Table 2). This is in contrast to the 20 unique solutions and 4 different fitness levels found by Hains et al. [7]. A closer look at the data reveals that these 4 fitness levels amount to the most frequent fitness levels in our data, comprising 99 % of the solutions found. The seldom found solutions are therefore a result of carrying out a greater number of Chained-LK searches to sample solutions close to the global optimum.

Table 1: lin318 summary data

	All Sols	Fitness Levels			
		42029	42143	42155	42163
Unique Solutions	4	1	1	1	1
Fitness Level Freq. (%)		96.02	3.59	0.09	0.30
Colour of funnel in figures		🟥	🟩	🟦	🟪
Symbol in Fig. 4a		○	△	+	×

Table 2: att532 summary data

	All Sols	Fitness Levels							
		27686	27693	27703	27704	27705	27706	27708	27715
Unique Solutions	47	2	1	8	8	13	8	5	2
Fitness Level Freq. (%)		41.78	0.04	33.17	0.65	20.69	3.58	0.07	0.02
Start Point Freq. (%)		21.35	0.04	5.80	0.16	4.64	0.57	0.03	0.01
Colour of funnel in figures		🟥	🟧	🟩	🟩	🟦	🟦	🟪	🟪
Symbol in Fig. 4b		○	△	+	×	◇	▽	⊠	✳

The two globally optimal solutions account for only 42 % of all funnel floor solutions found but all the fitnesses are within 0.10 % of the global optimum. As previously mentioned, for att532, the starting points we try to escape from are the most frequent funnel floor solution within each fitness level. These make up 33 % of the solutions found.

The pairwise bond distances between the starting points for both instances are given in Fig. 3. In most cases, the pairwise distance between any two solutions is non-trivial. For example, the bond distance between the first two best solutions for lin318 is 37.

For att532, the smallest bond distance between start points is only 16. This seems to be a bridgeable distance with a small number of double-bridge moves. The starting point with fitness 27,693 only represents 0.04 % of funnel floor solutions. It is at distance 16 from the start point with fitness 27,686 that constitutes 21 % of solutions found. These numbers suggest that there is a reasonable way to move between these funnels, which explains why so few solutions with fitness 27,693 are found. This is corroborated by the local optima networks visualised in Section 3.2.

To analyse the fitness distribution of local optima within funnels, let us consider Figure 4. Dot plots of fitness versus bond distance to the global optimum are presented for both instances. In addition, kernel density estimation distributions of points are provided.

Here our results match those of Hains et al. [7]. Firstly, local optima within a funnel are correlated in fitness and distance to their own respective starting

(a) lin318 – 4 start points
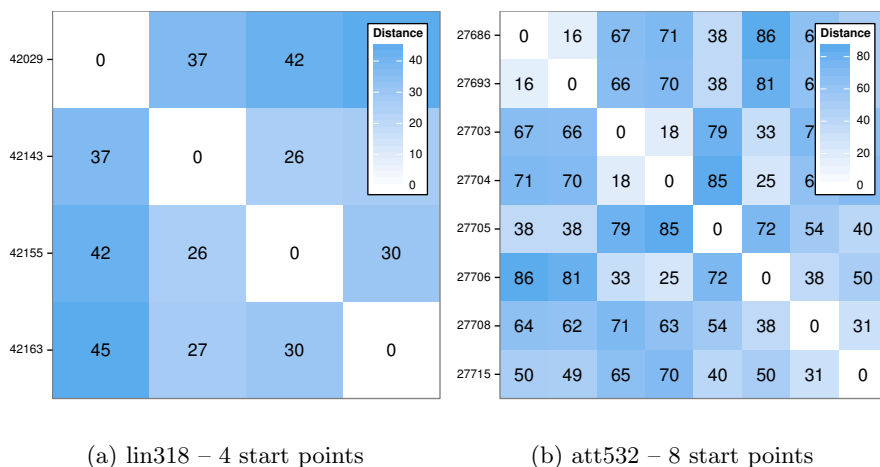

(b) att532 – 8 start points

Fig. 3: Pairwise distances between funnel floor solutions for instances lin318 and att532. Fitness levels are indicated on the left of each plot. In (a), instance lin318 has a single solution per fitness level. In (b), the most frequent solution is selected for each fitness level of att532.

point. Secondly, there is little correlation between fitness of local optima near the global optimum and their distance to it. However, for att532, the great majority of the local optima observed by Hains et al. when using double-bridge were below the 27,750 fitness level and a plot similar to ours was only obtained when using 2-change instead of double-bridge. They therefore concluded that double-bridge exacerbates the multi-funnel structure. We found instead that, when comparing the two escape operators, it is 2-change that exacerbates the multi-funnel structure. In other words, it is harder to escape funnels using 2-change as compared to double-bridge.

Figure 5 gives the escape and pairwise bond distance distributions for both instances. With a mean and mode of 1 for the escape distance, we can see that the double-bridge move is highly effective in escaping from the starting points.

For bond distances, the distribution for all edges differs from the distribution considering only edges between a start point and the solutions it escaped to. For lin318, when considering all start points, the distribution roughly resembles a step function with 2 steps which then quickly tapers off. The same distribution can be observed when considering each start point separately (not shown here). For att532, the bond distance distributions when considering a single start point to the local optima within the funnel appear to be bimodal (not shown here) or similar to the distribution when considering all start points. We intend to look more closely at distributions within individual funnels in future work.

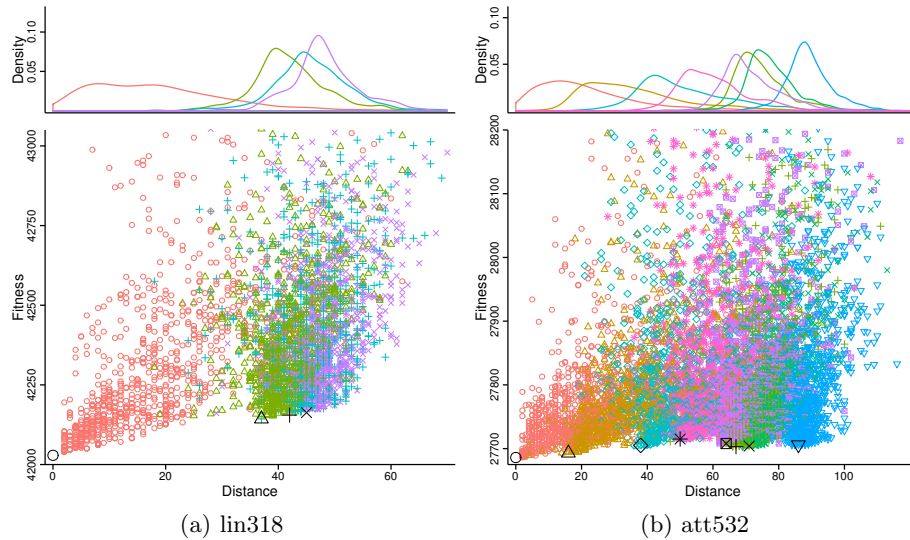(a) lin318   (b) att532

Fig. 4: Dot plots and corresponding density distribution plots of the local optima generated when escaping from funnel floors. Bond distance is computed w.r.t. to the global optimum, or the most frequent of the two global optima in the case of att532. The range of fitness values displayed is chosen to encompass at least 95 % of points. Start points are indicated by a black symbol.

3.2 Local Optima Networks

The two local optima networks models, using escape and bond distance edges, were extracted and visualised for the two selected TSP instances. Both models clearly suggest a multi-cluster (multi-funnel) structure (see Figure 6 explained below). The escape edges give a network view of the search process by Chained-LK, while the bond distance model is more general and illustrates the distributions of local optima which are close in distance.

At the heart of network visualisation is the graph layout. We use here the Fruchterman and Reingold's method [6] provided by the igraph package [5] for the R statistical language. The method is based on exploiting analogies between the relational structure in graphs and the forces among elements in physical systems. Specifically, considering attractive and repulsive forces by associating vertices with balls and edges with springs. The heuristic is concerned with drawing graphs according to some generally accepted aesthetic criteria such as a) distribute the vertices evenly in the frame, b) minimise edge crossings, c) make edge lengths uniform, and d) reflect inherent symmetry [6].

Figure 6 visualises the two network models (escape and distance edges) on the two studied instances. In order the make the picture manageable in size, sub-graphs of the whole sampled networks were selected for visualisation. The sub-graphs include all the funnel floor solutions (drawn as squares), and all the
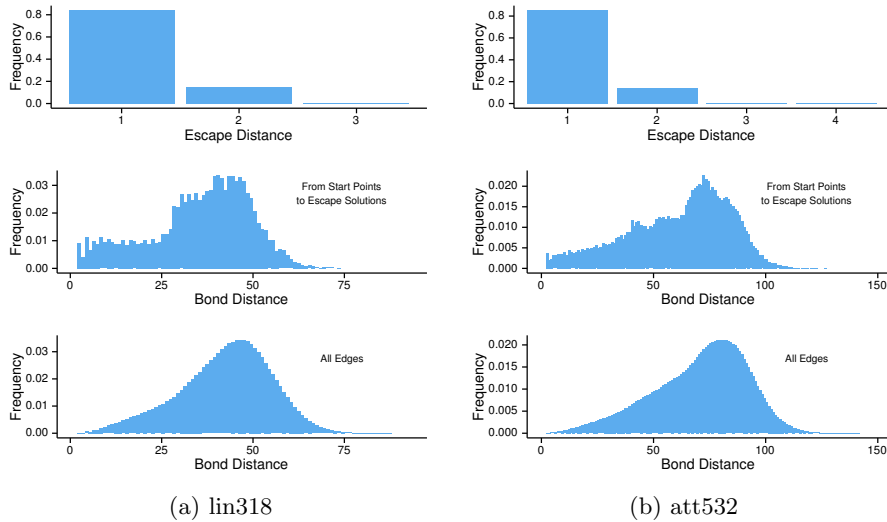
(a) lin318   (b) att532

Fig. 5: Escape and bond distance distributions. The most frequent escape distance is 1. The maximum escape distance is 4 on att532, but occurred only once.

solutions that we call *frontier* nodes (drawn in black). These frontier nodes are those that can be attained from more than one funnel start point by the escaping mechanisms (i.e., a sequence of double bridge moves followed by LK-Search). The colour of the remaining nodes indicates the funnel (fitness level) membership (as indicated in Tables 1 and 2 for lin318 and att532, respectively) with the red colour identifying the funnel of the global optimum. For the lin318 instance, 10 % of the funnel basin points were selected for visualisation. This percentage was 5 % for the larger att532 instance. All the escape edges are visualised, with darker grey indicating edges with escape distance 1. Visualising all bond distance edges is not feasible, so we set a threshold of 1/10 of the maximum distance to the global optimum in the sampled points (i.e., there is an edge if the distance between nodes is below the given threshold). This threshold was a distance of 9 for lin318 and 14 att532. Again the darker grey identifies edges with the minimum distance.

The multi-funnel structure can be visualised in the network plots in Figure 6, which separate in clearly defined clusters of solutions. The lin318 instance features 4 clusters, while att532 has 8 clusters. The clusters are more clearly defined for the escape edges, but interestingly, the same overall structure appears for the distance edges. It is interesting to observe that some points (drawn in black) 'belong' to more than one funnel. That is, they can be reached from more than one funnel floor by double-bridge moves followed by LK-search. Therefore, it is possible for Chained-LK to escape some funnels, but it seems difficult for it to consistently escape from all funnels.

(a) lin318 − Escape Edges (b) lin318 − Distance Edges

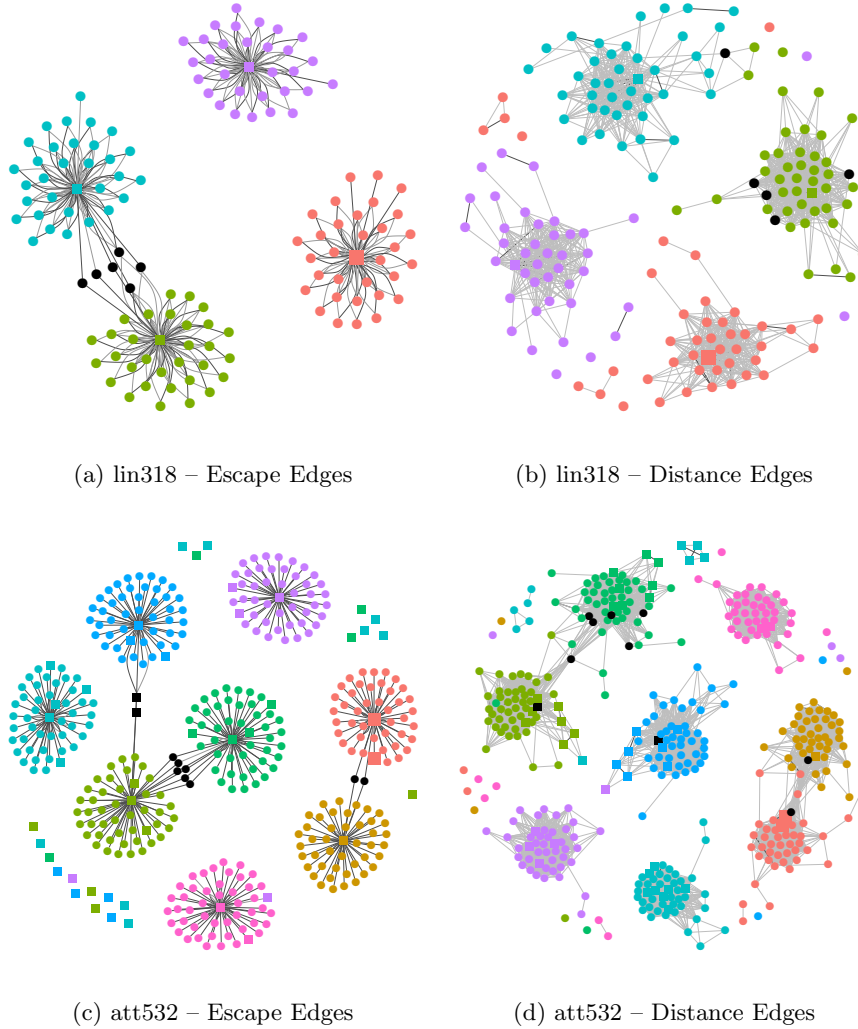(c) att532 − Escape Edges (d) att532 − Distance Edges

Fig. 6: Visualisation of Local Optima Networks for lin318 (top) and att532 (bottom). Both networks models, using escape and distance edges, are visualised. Nodes are local optima and edges represent escape or distance edges (with a set threshold), respectively. Square nodes represent solutions that belong to the funnel floors, while circle nodes to funnel basins. The larger square nodes (in red) are the global optima. Colours identify the different funnels (or fitness levels) as indicated in Tables 1 and 2. The black nodes are 'frontier' points, i.e., points that can be reached from more than one funnel.

An interpretation of the effectiveness of Chained-LK may be obtained when considering the local optima networks together with the fitness levels of the start points of each funnel, their frequency when sampling the funnel floors and the pairwise bond distance between start points.

For lin318, the two connected funnels are the ones whose start points have fitness 42,143 and 42,155 and were sampled 3.59 % and 0.09 % of the time respectively. They are also the two closest start points for lin318, with a distance of 26. For att532, as was observed in Section 3.1, start points with fitness 27,686 and 27,693 are at a distance of 16 and constitute 21.35 % and 0.04 % of sampled funnel floors. They are at a bond distance of 16 to each other and their corresponding funnels are linked in the local optima networks. The start point with fitness 27,703 (5.80 %) is connected to the start point with fitness 27,704 (0.16 % and distance 18). While these three observations are not sufficient to draw broad conclusions, an initial interpretation is that 'close enough' start points exhibit funnels that are linked to each other. Furthermore, when two funnels are connected, it is highly probable that the search will end up in the funnel with the funnel floor with better fitness.

The start point with fitness 27,703 is also connected to the one with fitness 27,706 (0.57 % and distance 33), but through two other floor solutions (indicated by black squares in the figure) that were not used as start points. These two floor solutions are also of fitness 27,703 and 27,706 and are only at a bond distance of two from the start point with the same fitness.

4   Conclusions

We have implemented a sampling procedure to extract local optima networks for TSP instances. In particular, we studied the search space structure close to the global optimum and confirmed the existence of multiple funnels. Our study is the first to analyse local optima networks for TSP and provide a clear visualisation of its multi-funnel structure. The proposed distance local optima network model is a contribution of this article, which may find easy application in other combinatorial optimisation problems.

Our analysis considered the well-known Chained-LK heuristic as implemented in the Concorde software package. Chained-LK is an iterated local search approach combining LK-search with double-bridge as the perturbation or escape operator. On two selected TSP instances, we found that while some funnels are directly connected to other funnels via double-bridge escape moves, most of them are not. This gives a visual insight of why Chained-LK produces suboptimal solutions in some runs, and justify the multiple restarts used in the default Concorde implementation. We hypothesise that when Chained-LK produces sub-optimal solutions, it is because it gets trapped in a sub-optimal funnel and the double-bridge escape mechanism, while generally efficient to escape local optima, is not strong enough to escape some funnels. Future work will explore alternative funnel-escape mechanisms such as the recently proposed Partition Crossover [20, 21], and will study Tunneling Crossover Networks for TSP [14].

Data Access. All data generated during this research are openly available from the Zenodo repository (`http://doi.org/10.5281/zenodo.20732`).

References

1. Applegate, D., Bixby, R., Chvátal, V., Cook, W.: Concorde TSP solver (2003), `http://www.math.uwaterloo.ca/tsp/concorde.html`
2. Applegate, D., Cook, W., Rohe, A.: Chained Lin-Kernighan for Large Traveling Salesman Problems. INFORMS Journal on Computing 15, 82–92 (2003)
3. Applegate, D.L., Bixby, R.E., Chvátal, V., Cook, W.J.: The Traveling Salesman Problem: A Computational Study. Princeton University Press (2007)
4. Boese, K.D., Kahng, A.B., Muddu, S.: A new adaptive multi-start technique for combinatorial global optimizations. Operations Research Letters 16, 101–113 (1994)
5. Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal Complex Systems, 1695 (2006)
6. Fruchterman, T.M.J., Reingold, E.M.: Graph drawing by force-directed placement. Software Practice Exper. 21(11), 1129–1164 (Nov 1991)
7. Hains, D.R., Whitley, L.D., Howe, a.E.: Revisiting the big valley search space structure in the TSP. Journal of the Operational Research Society 62(2), 305–312 (2011)
8. Helsgaun, K.: An effective implementation of the LinKernighan traveling salesman heuristic. European Journal of Operational Research 126(1), 106–130 (Oct 2000)
9. Iclanzan, D., Daolio, F., Tomassini, M.: Data-driven Local Optima Network Characterization of QAPLIB Instances. In: Proceedings of the 2014 Conference on Genetic and Evolutionary Computation. pp. 453–460. GECCO '14, ACM, New York, NY, USA (2014)
10. Lin, S., Kernighan, B.W.: An Effective Heuristic Algorithm for the Traveling-Salesman Problem. Operations Research 21, 498–516 (1973)
11. Martin, O., Otto, S.W., Felten, E.W.: Large-Step Markov Chains for the Traveling Salesman Problem. Complex System 5, 299—-326 (1991)
12. Möbius, A., Freisleben, B., Merz, P., Schreiber, M.: Combinatorial optimization by iterative partial transcription. Physical Review E 59(4), 4667–4674 (Apr 1999)
13. Newman, M.E.J.: Networks: An Introduction. Oxford University Press, Oxford, UK (2010)
14. Ochoa, G., Chicano, F., Tinos, R., Whitley, D.: Tunnelling crossover networks. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). pp. 449–456. ACM (2015)
15. Ochoa, G., Tomassini, M., Verel, S., Darabos, C.: A study of NK landscapes' basins and local optima networks. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO). pp. 555–562. ACM (2008)
16. Ochoa, G., Verel, S., Daolio, F., Tomassini, M.: Local optima networks: A new model of combinatorial fitness landscapes. In: Richter, H., Engelbrecht, A. (eds.) Recent Advances in the Theory and Application of Fitness Landscapes, Emergence, Complexity and Computation, vol. 6, pp. 233–262. Springer Berlin Heidelberg (2014)

17. Reinelt, G.: Tsplib—a traveling salesman problem library. ORSA Journal on Computing 3(4), 376–384 (1991), `http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/`
18. Verel, S., Ochoa, G., Tomassini, M.: Local optima networks of NK landscapes with neutrality. IEEE Transactions on Evolutionary Computation 15(6), 783–797 (2011)
19. Wales, D.J., Miller, M.a., Walsh, T.R.: Archetypal energy landscapes. Nature 394(August), 758–760 (1998)
20. Whitley, D., Hains, D., Howe, A.: Tunneling Between Optima: Partition Crossover for the Traveling Salesman Problem. In: Proceedings Genetic and Evolutionary Computation Conference. pp. 915–922. GECCO '09, ACM, New York, NY, USA (2009)
21. Whitley, D., Hains, D., Howe, A.: A hybrid genetic algorithm for the traveling salesman problem using generalized partition crossover. In: Parallel Problem Solving from NaturePPSN XI. pp. 566–575 (2010)