The Effect of Landscape Funnels in QAPLIB Instances

Sarah L. Thomson, Gabriela Ochoa, Fabio Daolio and Nadarajen Veerapen

University of Stirling

Computing Science and Mathematics

Stirling, Scotland, UK

{sto,goc,fda,nve}@cs.stir.ac.uk

ABSTRACT

The effectiveness of common metaheuristics on combinatorial optimisation problems can be limited by certain characteristics of the fitness landscape. We use the *local optima network* model to compress the 'inherent structure' of a problem space into a network whose structure relates to the empirical hardness of the underlying landscape. *Monotonic sequences* are used on the local optima networks of a benchmark set of QAP instances (QAPLIB) to expose landscape *funnels*. The results suggest links between features of these structures and lowered metaheuristic performance.

1 INTRODUCTION

The number of local optima in a fitness landscape has long been linked to search difficulty. Some recent research, however, has suggested that their global distribution and connectivity might be might be more relevant to the performance of certain metaheuristics such as iterated local search (ILS) [7, 11, 12]. We model fitness landscapes as *networks* of local optima - local optima networks (LONs) [16]. The nodes are the local optima configurations, and the edges are possible transitions between them with a given search operator. The study of LONs can reveal valuable information about navigability of a fitness landscape. If there are sub-networks of local optima which are densely connected to one another, and yet with sparse ties to other areas, this could impact optimisation.

In this work, we focus on the nature of landscape *funnels*. The notion of funnels has received attention in the combinatorial optimisation literature [6, 10, 12], largely because if there are multiple funnels, ILS in particular can become trapped in them and never reach the globally optimal solution.

We expose landscape funnels using *monotonic sequences*, which in our context are paths of local minima with only improving or equal (in objective value) moves between them [13]. A LON can be trimmed to retain only local optima links that lead to a better solution — or one with the same fitness — and funnels can subsequently be identified. This is done by first looking for the end of the funnels, which are the local optima with no outgoing links left. These nodes have no links to higher-quality solutions and they are at the end of a monotonically descending sequence of local minima.

If there is only one such node, this would mean the underlying landscape's global structure conforms to the so-called 'big valley' phenomenon, and the sole funnel bottom would be the global optimum. This distribution of local minima means that the coarse-grained trajectory of a local search naturally leads to optimality. If, however, there is more than one, a more complex picture emerges. If even one sub-optimal funnel is present, local search could get trapped there.

In this work, we use the Quadratic Assignment Problem Library (QAPLIB) [2]. QAPLIB is a well-studied benchmark set of QAP instances. The problem classes vary from random uniform to occurrences from real life. Because the prevalence of multiple funnels in combinatorial landscapes is not yet known, the QAPLIB is an interesting test-bed for funnel analysis because of the diversity of problem class, problem size and landscape ruggedness. This variance could potentially mean that findings would transfer to unseen QAP problems. To gain an empirical view of the relationship between the features of landscape funnels in QAPLIB and metaheuristic ability, ILS and Simulated Annealing (SA) are used.

2 FITNESS LANDSCAPES AS NETWORKS

We first define the fundamentals of fitness landscapes. Following this, the local optima network model which we use to examine the underlying landscape connectivity is described.

2.1 Preliminaries

A fitness landscape [15] is a triplet $\{S, V, f\}$ where S is the set of all possible solutions, $V : S \longrightarrow 2^{|S|}$, a neighbourhood structure, is a function that assigns to every $s \in S$ a set of neighbours $V(s)$, and f is a fitness (objective value) function such that $f : S \longrightarrow \mathbb{R}$, where the fitness value is a real number that can be viewed as the *height* of a given solution in the landscape.

In the QAP, a solution is encoded as a permutation of the set $\{1, 2, ..., n\}$. A configuration, S, is considered to be a neighbour of another solution S' if a pairwise exchange in S can result in S'.

2.2 Local Optima Networks

Before defining the local optima network model, we first articulate the components.

Nodes. A locally optimal solution, *lo*, has superior fitness to all its neighbouring solutions. Formally, $\forall s \in V(lo)$, $f(s) \geq f(lo)$ for a minimisation problem such as the QAP. The nodes in an optima network, *LO*, are the set of local optima. In this work, due to the large search spaces the nodes and edges of the LONs are obtained through the sampling process described in [9].

Escape edges. The edges in the network represent potential 'escapes' from one local optimum to another. These links are characterised by two parameters: a representation-specific measure of distance between two configurations, *d*, and an integer, *D*, which is the permitted number of escape moves between two local optima in order for an edge to be traced between them. Formally, there exists a link e_{ij} between the local optima lo_i and lo_j if $\exists s \in S$ where $d(s, lo_i) \leq D$ and s is a solution which — following hill-climbing — converges to the local optimum lo_j. In the present study, *d* is a pairwise exchange in the permutation and *D* is set at two. We denote the full set of edges in the network as *E*. These links are weighted with the *probability* of that escape move.

Local optima network (LON). We can now define the weighted, directed local optima network (LON) as a graph which has the set of local optima *LO* as nodes, and a set of links, *E*, traced between optima if the weight, w_{ij}, of the edge, e_{ij}, is greater than zero, i.e. if it's possible to 'escape' from the source local minimum to the destination.

2.3 Funnel Detection with Monotonic Sequences

We take advantage of the level of abstraction LONs provide to find *funnels* — which are essentially higher-level basins of attraction — in the configuration space. In this work, we define a funnel as being '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'[4]. In our context, the landscape funnels are sequences of local optima that form downhill paths converging to a local optimum, which itself has no possible transitions to a better solution. Another name for a path of local minima such as this is a *monotonic sequence* [1]. We extract monotonic sequences from the LONs of QAPLIB instances as a case study, to identify trajectory paths that lead to the same end point. Doing this allows us to find landscape 'traps' which could lower search performance. With the sampled local optima networks [9], we prune off any edges which are deteriorating in fitness. Therefore, what is left are monotonic sequences of local minima which can be seen as the 'backbone' of the fitness landscape. Note that neutral links are retained; the QAPLIB instances under study contain varying degrees of neutrality at the LON level. This is dealt with by the compression of sets of local minima who have equal fitness (and are connected) into a single node.

Monotonic Edges. The set of monotonic edges are LON escape edges, and therefore characterized by the two parameters *d* and *D*. These edges are all equal or improving in nature, i.e. they are directed from a worse-fitness optimum to either a better-fitness one or one equal in fitness. Specifically, a monotonic edge e_{ij} exists between lo_i and lo_i if $f(lo_i) \geq f(lo_j)$ and if $\exists s \in S$ where $d(s, lo_i) \leq D$ and s is a solution that — following hill-climbing — converges to the local optimum lo_j.

Monotonic Local Optima Network (M-LON). A directed and weighted local optima network with the set of local optima as the vertices, and the set of links being non-deteriorating fitness moves between them.

M-LON plateau. A group of local optima with equal fitness which are connected to one another in the M-LON. In other words, a neutral network at the LON level instead of at the solution level.

Compressed LON nodes. Nodes which form a M-LON plateau being compressed into a single node. The set of vertices in the compressed LON is the set of compressed nodes, with some nodes only forming a LON plateau of size one.

Compressed Monotonic Local Optima Network (CM-LON). A local optima network where the vertices are the compressed LON plateaus — even if they consist of a single local optimum — and the directed links between them are non-deteriorating in terms of objective function evaluation. The links are *weighted* by the aggregation of links leading out of a compressed plateau node, resulting in a single link. From the *CM-LON*, we can proceed to extract monotonic sequences and subsequently identify funnels in the underlying fitness landscape. A crucial part of a landscape funnel is its eventual end point. These can be seen as the 'destination' of a search algorithm following one of the downhill pathways of local minima. We can find these by looking at our monotonic sequences, and identifying nodes which have no outgoing links, which means an algorithm using local search would not be able to make any improving moves from there. In graph theory, such a node (one with out-degree zero) is called a *sink*. We use this terminology to refer to the end of landscape funnels. Counting the number of sinks tells us how many funnels we are dealing with.

2.3.1 Modularity. Because we are dealing with networks of optima, we can take advantage of clustering algorithms to assess whether the optima form distinct groups. An interesting LON feature which has been previously studied [7, 9] is called *modularity*. Modularity is a measure for how strong — or significant — a network partition is [5]. In other words, if the network has 'true' community structure, different from a random graph, the modularity should be high. This would mean that links are much more dense within clusters than between them. It's possible that if clusters of optima have this property, local search could encounter difficulty trying to travel between communities.

Figure 1 shows the top 10% highest-quality optima in the LONs of two seperate QAPLIB instances, with the top 30% heaviest-weighted edges (most likely paths through the configuration space) retained between nodes. Figure 1a shows a LON which had been found to exhibit multiple-funnel structure through the extraction of monotonic sequences. Conversely, Figure 1b shows an instance which was discovered to have only one funnel, otherwise known as the 'big valley' global topology. The globally optimal solutions are shown in red.

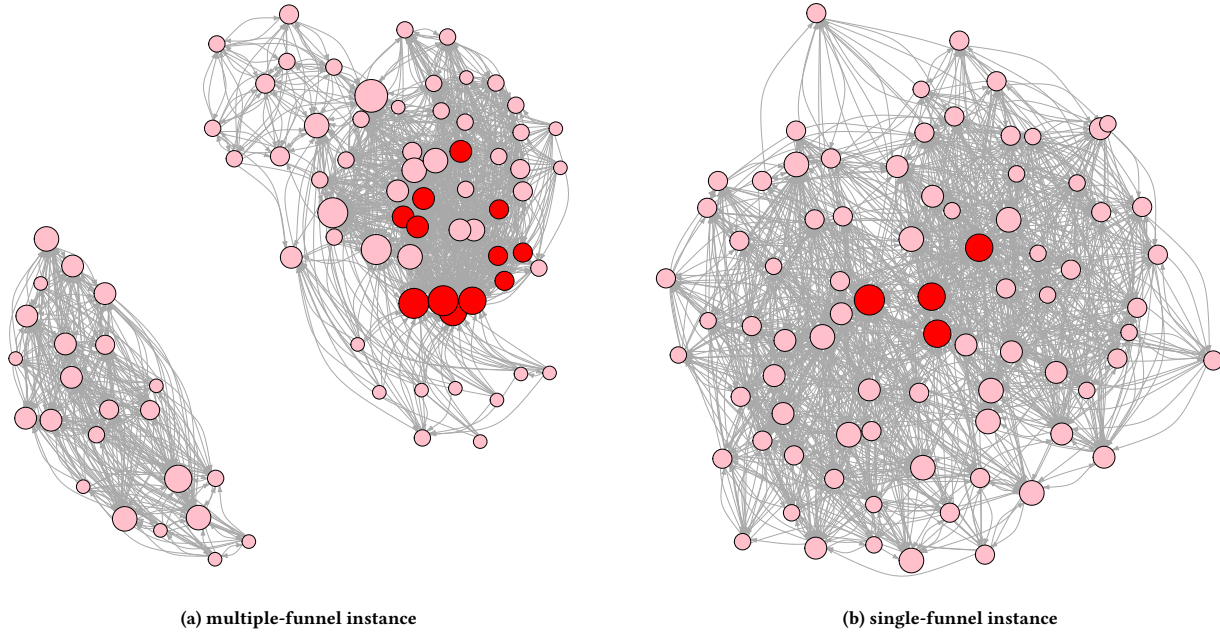(a) multiple-funnel instance　　　　　　　　　　　　　　　　　　(b) single-funnel instance

Figure 1: the top 10% of local optima in a multiple-funnel and single-funnel QAPLIB instance LON, respectively. The global optima are red in colour and node size is proportional to pagerank centrality. The multiple-funnel instance original LON had 867 optima, and the single-funnel had 910.

Comparing the two networks, which have a similar number of local optima, we can see that the multiple-funnel LON is split cleanly into two isolated components. Conversely, the nodes in the single-funnel LON are one large connected component. Recalling that these are the highest-quality optima in the instances — and that only the most likely paths between them are traced — we can imagine the implications for local search. Intuitively, the underlying landscape of Figure 1a should be more difficult to navigate. Indeed, the success rates for ILS and SA on this QAPLIB instance were 0.639 and 0.441, respectively. On the other hand, the landscape corresponding to the LON in Figure 1b should be less challenging. Fittingly, the success rate for ILS was 0.729 and SA had a proportional success of 0.7.

3　EMPIRICAL SETTING

3.1　Test Problem

The Quadratic Assignment Problem (QAP) Library (QAPLIB). The Quadratic Assignment Problem (QAP) is one of the most widely-known combinatorial optimisation problems. Solutions are encoded permutations assigning a series of *n* facilities to *n* locations. Each facility and location pair have a distance, d_{ij}, and a flow, f_{ij} between them. The objective function is the cost of the product of the flows and the distances, making it a minimization problem [3]. The QAPLIB[2] is a benchmark set of QAP instances of varying nature arising from the structural differences in their distance and flow matrices. Some are random, some are grid-based, and some are from real-life QAP problems. For a detailed introduction, the reader is referred to [2].

Table 1: The three sets of problem classes. All are from the QAPLIB.

class	instance name	description
bur	{26{a-h}}	real life|8
chr	{12{a-c},15{a-c},18{a-b},20{a-c},22{a-b},25a}	tree/complete|14
nug	{12,14,15,16{a-b},17,18,20,21,22,25,27,28,30}	grid-based|13

In this work, we focus on three full sets of well-known instance classes in the QAPLIB, for a total of 37 instances. This gives rise to varying instance size; these can be seen in the numerical values in the *instance name* column in Table 1, which are stated along with the name of the set (*class*), the names of the instances (*instance name*), and how many instances are in each group (*description*). We extract monotonic sequences from the local optima networks of these instances. Two network clustering algorithms are also applied to the LONs: *Markov Clustering* (MCL) [17] and *Infomap* [14]. For each of 37 instances, the following metrics were computed:

- *optima*: the number of local optima
- *subsinks*: the number of sub-optimal sinks (and therefore landscape funnels).
- *subsinkweight*: the aggregated incoming link weight to the sub-optimal sinks(s) (normalized by the strength of incoming weight to all sinks). This should relate to the probability a local search will be drawn to sub-optimality.
- *meanFun*: the mean (disjoint) funnel size, normalized by the number of local optima in the corresponding LON;

that is, the proportion of local minima lying on monotonic sequences that are unique to a given funnel

- *LONneut*: the neutral degree of the network: for a local minimum, the average proportion of neighbouring local minima with equal fitness value
- *modM*/*modI*: the *modularity* of LON network partitions as found by the MCL and Infomap clustering algorithms, respectively

3.2 Metaheuristics

Our aim was to determine which landscape qualities — with particular interest in funnels - in the QAPLIB lower performance of heuristic search. With this in mind, two trajectory methods — ILS and SA — were used on the instances. The ILS implementation used best-improvement local search with the pairwise exchange (swap) operator on the improvement stage, while the perturbation step considers three swaps. For the SA, a hyper-parameter tuner [8] was used to optimize temperature annealing schedule, which follows an exponential decay with a maximum number of trials allowed between cooling steps, with the aim of minimizing the ratio between the best-found solution evaluation and the fitness of the best-known solution. The optimised parameters are shown in Table 2. *Init.T* is the initial temperature, *Final.T* is the final temperature, and *maxT* is the maximum number of iterations at a given temperature.

Table 2: SA parameters, optimized by hyper-parameter tuner

Parameter	value
Alpha	0.8198
Init.t	893 649
Final.t	0.482
Max.t	1889

4 RESULTS

4.1 Correlations

Figure 2 shows a pairwise correlation matrix of the landscape features with ILS and SA performance. The blue text in the upper right triangle represents the Spearman correlation with indication of p-value significance, as described in the caption. The diagonal panels report density plots, indicating the distributions of each variable studied. The lower left triangle shows scatter plots. The variable names are described in Section 3.1. The most important rows are *SAs*, *SAt*, *ILSs* and *ILSt*, which represent the success rate and runtime of SA and for ILS, respectively.

The success rate of ILS is lowered in a significant way not only by the *number* of lower-than-optimal sinks in the landscape, but also by the incoming weight to these local optima, which can be seen as the probability a stochastic local search will end up there. We can see these associations by looking at the *ILSs* column and checking against the *subsinks* and *subsinkweight* rows. The ILS seems to have raised efficiency and effectiveness when the funnels are larger.

Table 3: Linear model fit using the sub-optimal sink weight (*subsinkweight*), modularity (strength) of local optima clusters (*mod*), and the neutral degree (*LONneut*) as predictors and success rate of ILS as the response. $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

coef	value	error
subsinkweight	-0.46068 ***	0.11870
mod	-0.58469***	0.14982
LONneut	-0.47332 ***	0.11678
adj R^2	0.6321	
p-val	6.283e-08	

We can see the relationship in Figure 2 by checking where the *meanFun* row meets the *ILSs* and *ILSt* columns. The neutral degree of the networks contributes to a slower search, as we can see from the point where the *LONneut* row and *ILSt* column intersect. Our results suggest that modularity slows down search for ILS. A similar effect was observed in [9] for a Tabu search implementation.

To see the landscape features contrasted with the SA perfomance, we can look along the rows labelled *SAs* and *SAt*, which are the success rate and function evaluations, respectively. Checking these against the columns *subsinkweight* and *subsinks* in particular, we can see that more sub-optimal funnels — and the flow towards them — seems to lower the effectiveness and efficiency of SA on QAPLIB instances considerably. The number of local optima — and therefore ruggedness — of the landscape affects both measures of performance in a significant way, and so do strong (modular) clusters of local optima (see *modI*).

Looking at the row corresponding to the number of local optima (*optima*), notice that in the case of both algorithms, the ruggedness does lower their respective success rates; however, the sub-optimal sinks and their associated incoming flow have a stronger association (*subsinks* and *subsinkweight*). In terms of function evaluations (*ILSt* and *SAt*), the amount of ruggedness seems to have a slightly stronger effect on ILS than the funnel metrics do; for SA, the opposite is true.

4.2 Linear Model

Table 3 shows the coefficients, standard error, adjusted R^2 and p-value for a linear model fit to the success rate of ILS. The dependent variable is the proportional success rate and the predictors are the neutral degree of the LON (*LONneut*), the modularity of the local optima clusters found by the Infomap algorithm (*mod*), and the incoming link weight to sinks that were not optimal (*subsinkweight*). The sink weight is normalised by the total incoming weight to all sinks in that instance.

These variables can explain just over 63% of the variance seen in the success rate of ILS. We can see this looking at the value for $adjR^2$ in Table 3. The p-value is sufficiently small that it would be reasonable to reject the null hypothesis.

All predictors appear to have statistical significance, with associated p-values of <0.001 (shown by ***). The attributes of this linear
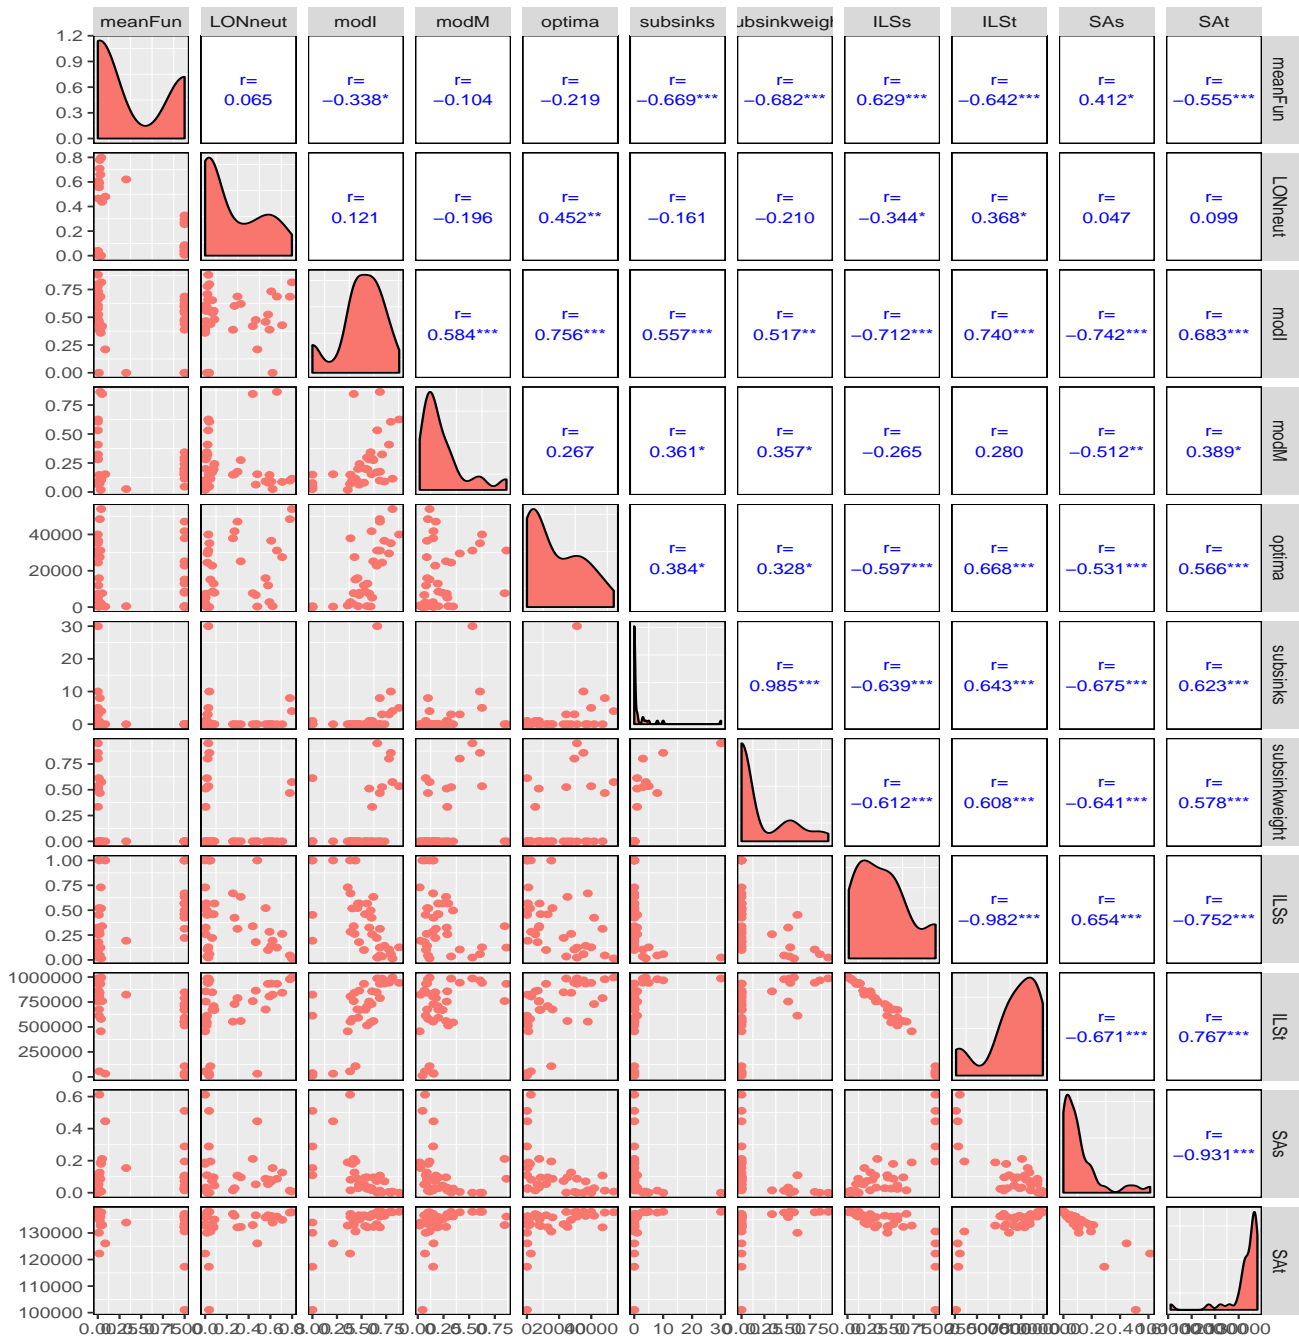
Figure 2: Correlation matrix for all cluster features studied and the performance of ILS and SA, in timesteps and success rate. Lower triangle: relationship scatterplots. Diagonal: density plots showing the distributions. Upper triangle: Spearman correlation coefficients. P-values shown by asterix: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

model seem to suggest that ILS effectiveness on QAP instances is closely related to these three landscape features in particular.

5 DISCUSSION

The findings presented in Section 4.1 suggested a significant link between the number of sub-optimal sinks — with their associated incoming flow — and lowered performance from both ILS and SA. This suggests that in QAP fitness landscapes, funnels are important

for the operation of common trajectory-based metaheuristics. There is also an association between the mean size of the funnels and search success for ILS. The relationship was positive, meaning the larger the disjoint funnels, the more successful the search was. Recalling that the sizes represent a *proportion* of the set of local optima, a large funnel size would either mean a big valley structure (which is naturally the optimal funnel), or a few very wide funnels. This would mean that at least some of the time ILS would begin its search in the correct one, and if there is little overlap in funnel membership the trajectory would lead cleanly to the global best.

Another interesting result in Section 4.1 was that both algorithms had their success rate lowered by the sub-optimal sinks – and their incoming flow — in a more significant way than the ruggedness of the landscape. It could be that the landscape structure at a higher level of abstraction — an overall global picture — might possibly be more important.

Both ILS and SA had their performance lowered by the strength (modularity) of groups of local optima; this has been noted before with regards to Tabu Search on QAPLIB instances in [9]. It is interesting that these findings transfer to other algorithms. The presence of this behaviour across the performance of three metaheuristics hints at the importance of sub-networks of local optima in the QAP.

Section 4.2 shows that over 63% of variance in ILS success could be explained by sub-optimal sink weight, the neutral degree of the LON, and the modularity of local optima clusters. Each of the predictors had a statistically sound p-value, suggesting these three landscape features are important in navigating through a QAP configuration space.

6 CONCLUSIONS

We have conducted a case study on a set of benchmark QAP instances, to analyse how the nature of landscape funnels in varying classes of QAP problems affect common metaheuristics. Funnels were extracted using monotonic sequences, and the association between features and the performance of ILS and SA was examined. The results suggest that the number of sub-optimal landscape funnels in QAPLIB instances — along with their incoming flow — negatively impact ILS and SA. Specifically, both algorithms seemed to execute a slower search with lower success in part because of the existence of sub-optimal funnels. While not a funnel feature, the strength of *communities* of local optima found by cluster algorithms were shown to have contributed to lowered success and raised search time, corroborating previous findings [9]. This result has spanned three trajectory-based algorithms — ILS, SA and Tabu Search — suggesting that it could generalise to others. We argue that the ultimate success of trajectory-based metaheuristics in QAP instances is more dependent on the global topology and in-depth analysis of the *nature* of funnels in fitness landscapes — and how commonly they occur — should be fruitful. Future work will include other problem domains, to gain a view of how common multiple-funnel landscapes are.

REFERENCES

[1] R Stephen Berry and Ralph Breitengraser-Kunz. 1995. Topography and dynamics of multidimensional interatomic potential surfaces. *Physical review letters* 74, 20 (1995), 3951.

[2] Rainer E Burkard, Stefan E Karisch, and Franz Rendl. 1997. QAPLIB–a quadratic assignment problem library. *Journal of Global optimization* 10, 4 (1997), 391–403.

[3] Fabio Daolio, Sébastien Verel, Gabriela Ochoa, and Marco Tomassini. 2010. Local optima networks of the quadratic assignment problem. In *Evolutionary Computation (CEC), 2010 IEEE Congress on.* IEEE, 1–8.

[4] J P K Doye, M A Miller, and D J Wales. 1999. The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *Journal of Chemical Physics* 110, 14 (1999), 6896–6906.

[5] Santo Fortunato. 2010. Community detection in graphs. *Physics reports* 486, 3 (2010), 75–174.

[6] DR Hains, L Darrell Whitley, and Adele E Howe. 2011. Revisiting the big valley search space structure in the TSP. *Journal of the Operational Research Society* 62, 2 (2011), 305–312.

[7] Sebastian Herrmann, Gabriela Ochoa, and Franz Rothlauf. 2016. Communities of Local Optima as Funnels in Fitness Landscapes. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Concausing synonymference.* 325–331.

[8] F. Hutter, H. H. Hoos, K. Leyton-Brown, and K. P. Murphy. 2010. Time-Bounded Sequential Parameter Optimization. In *Proc. of LION-4.* To appear.

[9] 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.

[10] Monte Lunacek and Darrell Whitley. 2006. The dispersion metric and the CMA evolution strategy. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation.* ACM, 477–484.

[11] Gabriela Ochoa and Nadarajen Veerapen. 2016. Additional Dimensions to the Study of Funnels in Combinatorial Landscapes. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference (GECCO '16).* ACM, New York, NY, USA, 373–380. DOI:http://dx.doi.org/10.1145/2908812.2908820

[12] Gabriela Ochoa and Nadarajen Veerapen. 2016. Deconstructing the big valley search space hypothesis. In *European Conference on Evolutionary Computation in Combinatorial Optimization.* Springer, 58–73.

[13] Gabriela Ochoa, Nadarajen Veerapen, Fabio Daolio, and Marco Tomassini. 2017. Understanding Phase Transitions with Local Optima Networks: Number Partitioning as a Case Study (Forthcoming). In *17th European Conference on Evolutionary Computation in Combinatorial Optimisation (EvoCOP).* Springer.

[14] Martin Rosvall, Daniel Axelsson, and Carl T Bergstrom. 2009. The map equation. *The European Physical Journal Special Topics* 178, 1 (2009), 13–23.

[15] Peter F Stadler. 2002. Fitness landscapes. In *Biological evolution and statistical physics.* Springer, 183–204.

[16] Marco Tomassini, Sebastien Verel, and Gabriela Ochoa. 2008. Complex-network analysis of combinatorial spaces: The NK landscape case. *Physical Review E* 78, 6 (2008), 066114.

[17] Stijn Van Dongen. 2000. A cluster algorithm for graphs. *Report-Information systems* 10 (2000), 1–40.