

PageRank Centrality for Performance Prediction: The Impact of the Local Optima Network Model

Sebastian Herrmann · Gabriela Ochoa ·
Franz Rothlauf

Received: 28 July 2016 / Accepted: 3 May 2017

Abstract A local optima network (LON) compresses relevant features of fitness landscapes in a complex network, where nodes are local optima and edges represent transition probabilities between different basins of attraction. Previous work has found that the PageRank centrality of local optima can be used to predict the success rate and average fitness achieved by local search based metaheuristics. Results are available for LONs where edges describe either basin transition probabilities or escape edges. This paper studies the interplay between the type of LON edges and the ability of the PageRank centrality for the resulting LON to predict the performance of local search based metaheuristics. It finds that LONs are stochastic models of the search heuristic. Thus, to achieve an accurate prediction, the definition of the LON edges must properly reflect the type of diversification steps used in the metaheuristic. LONs with edges representing basin transition probabilities capture well the diversification mechanism of simulated annealing which sometimes also accepts worse solutions that allow the search process to pass between basins. In contrast, LONs with escape edges capture well the diversification step of iterated local search, which escapes from local optima by applying a larger perturbation step.

Keywords Fitness landscape analysis · search difficulty · PageRank centrality · local optima networks · NK landscapes

S. Herrmann
Johannes Gutenberg-Universität, Mainz, Germany
E-mail: s.herrmann@uni-mainz.de

G. Ochoa
University of Stirling, Scotland, UK
E-mail: gabriela.ochoa@cs.stir.ac.uk

F. Rothlauf
Johannes Gutenberg-Universität, Mainz, Germany
E-mail: rothlauf@uni-mainz.de

1 Introduction

Local optima networks (Ochoa et al, 2008) can be used to study the structure of optimization problems. A local optima network (LON) is a compressed representation of a combinatorial fitness landscape. In particular, a LON is a directed graph where the vertices are the search space’s local optima. The edges¹ model the transitions between the local optima and are weighted by the transition probabilities of moving from one basin of attraction to another. In the literature, we find three types of LON models: edges with basin transition probabilities (Ochoa et al, 2010; Herrmann and Rothlauf, 2015), escape edges (Vérel et al, 2012; Herrmann, 2016) and LONs for partition crossover (Ochoa et al, 2015). LONs can be used to predict the performance of search algorithms on particular problem instances (search difficulty, Lu et al, 2014; Malan and Engelbrecht, 2014). This is possible as the network features of LONs like centrality or path length can capture the search difficulty of landscapes (Ochoa et al, 2014; Herrmann and Rothlauf, 2015).

For simulated annealing (SA) and a steady-state genetic algorithm (GA), Chicano et al (2012) studied the correlation between network features of landscapes and the search difficulty of instances for the quadratic assignment problem (QAP). They found that the average length of the shortest paths to the global optimum is a good predictor of search difficulty of SA ($R^2 \approx 0.5$), although the correlation is lower when predicting GA performance. Daolio et al (2012) conducted a similar study with the NK model (Kauffman and Weinberger, 1989), which is a binary combinatorial optimization problem. They tried to predict the difficulty of problem instances for iterated local search, and also found that the average length of the shortest paths to the global optimum had a high accuracy in predicting difficulty ($R^2 \approx 0.5$). Herrmann and Rothlauf (2015) suggested to use the PageRank centrality of the LON for predicting the expected performance of local search based metaheuristics. They found that the performance of first-improvement local search as well as SA can be well predicted using the PageRank centrality of the global optimum for LONs with edges representing basin transition probabilities. For the NK model, the use of the PageRank explained more than 90% of the variance of search performance. Furthermore, PageRank centrality is a better predictor of search performance than traditional approaches such as ruggedness, deceptiveness as well as the length of the shortest path to the optimum. Herrmann (2016) took up the idea of using PageRank centrality and found that the performance of iterated local search (ILS) can be well predicted using the PageRank centrality of a LON where the edges model the number of escape edges. Similar to Herrmann and Rothlauf (2015), the PageRank of the global optimum in the LON with escape edges explains more than 95% of the success rate of ILS.

¹ The edges in a directed graph are usually denoted as arcs. However, to be consistent with the current terminology in the LON literature, we use the term *edge* instead of *arc* throughout the paper.

Although the existing exploratory studies gained valuable insights into the possibilities of LONs and it is now known that PageRank centrality can be an accurate predictor for the expected performance of local search based metaheuristics, our current understanding of how the definition of the LON affects the prediction quality of the PageRank centrality is still rather vague. Furthermore, there are no studies that directly compare the different types of LON models in this context.

Consequently, this paper extends previous work (Herrmann and Rothlauf, 2015; Herrmann, 2016) and studies whether and how the definition of LON edges affects the ability of PageRank centrality to accurately predict the performance of local search based metaheuristics. Furthermore, it directly compares different types of local search based metaheuristics like SA and ILS for different instances of the Kauffman NK model. As expected, we are able to confirm the results of previous work and find that PageRank centrality is a good predictor for search difficulty. The PageRank centrality of the global optimum in LONs accurately predicts the success rate (probability of finding the global optimum) of SA as well as ILS and the average of the fitness values of the local optima weighted by their corresponding PageRanks is a good predictor of the solution quality at the end of a run (Herrmann and Rothlauf, 2015; Herrmann, 2016). As a new contribution, we find that LONs are stochastic models of a search algorithm in a landscape and, to accurately predict search difficulty, a LON must properly reflect the behavior of the local search based metaheuristic. We demonstrate for ILS and SA that the definition of the LON edges must fit to the diversification mechanisms of the metaheuristic. In particular, LONs with edges representing basin transition probabilities capture well the diversification mechanism of SA, whereas LONs with escape edges capture well the search dynamics of ILS. Thus, to ensure high prediction accuracy, the LON type must fit to the search paradigm. For our experiments, we use instances of the Kauffman NK model as scalable test problems. To assess the search difficulty of a landscape, we measure the performance of SA and ILS. As performance measures, we use the success rate, the average fitness achieved, and the average number of function evaluations (running time).

The work at hand indicates that LONs are a stochastic model of heuristic search in a fitness landscape and the search process is a random walk through the LON. Thus, the linkage structure of the LON and the distribution of nodes allow us to predict search difficulty of the landscape, which is just the probability that a random walk through the LON visits the node that represents the global optimum.

Section 2 briefly describes simulated annealing and iterated local search. Section 3 gives a short introduction to fitness landscapes and provides some formal definitions. In section 4, we define LONs with either basin transition probabilities or escape edges. In sections 5 and 6, we describe PageRank centrality, our experimental design, and the search space used (Kauffman NK family of landscapes). Experimental results are presented in section 7. We end with some conclusions, limitations, and future work.

2 Search Algorithms

We summarize the two variants of heuristic optimization methods used as representatives for local search based methods: simulated annealing (SA) and iterated local search (ILS).

2.1 Simulated Annealing

Simulated annealing (SA) is a method inspired by statistical mechanics (Kirkpatrick et al, 1983). Other than standard hill climbing procedures, SA has the ability to overcome local optima. It is based on an analogy with cooling down a liquid to a solid substance. Algorithm 1 outlines the basic principle of SA using a fixed cooling schedule for a maximization problem.

Algorithm 1 Simulated Annealing (SA)

Require: Solution space S ,
 Fitness function $f(S)$,
 Neighborhood $N(S)$,
 Initial Temperature T_0 , Cooling Rate c ,
 Stopping Threshold t

- 1: Choose random initial solution $s_0 \in S$
- 2: $i \leftarrow 0$
- 3: **repeat**
- 4: choose random $s_{i+1} \in N(s_i)$
- 5: **if** $(f(s_i) < f(s_{i+1}))$ **then**
- 6: $s_i \leftarrow s_{i+1}$
- 7: **else if** $\frac{e^{-|f(s_i) - f(s_{i+1})|}}{T_i} \geq \text{rand}(0, 1)$ **then**
- 8: $s_i \leftarrow s_{i+1}$
- 9: **end if**
- 10: $T_{i+1} \leftarrow c \times T_i$
- 11: $i \leftarrow i + 1$
- 12: **until** $i \geq t$
- 13: **return** s_{i+1}

Let S be the solution space, i.e. the set of all valid problem solutions. The function $f : S \rightarrow \mathbb{R}_{\geq 0}$ assigns a fitness value to each $s \in S$. SA starts with a randomly selected solution $s_0 \in S$. Then, in iterative steps, SA applies incremental changes to improve the initial solution: it selects a random solution s_{i+1} from the neighborhood of s_i . The neighborhood $N(s)$ is the set of solutions that can be reached by applying an incremental change to s . This procedure is repeated until a fixed limit of iterations has been reached. SA always accepts better solutions; solutions with a lower fitness are accepted with probability $e^{-|f(s_i) - f(s_{i+1})|} / T_i$. The probability to accept worse solutions decreases with a higher fitness difference $|f(s_i) - f(s_{i+1})|$ and a lower temperature T_i . As soon as the temperature approaches zero ($T \rightarrow 0$), SA turns into a standard hill climber performing local search.

The initial temperature T_0 and cooling rate c are parameters, which must be set problem-specific. A rule of thumb for a proper starting temperature is to randomly generate an initial number of solutions and set $T_0 \approx \sigma(f(s)) \dots 2 \times \sigma(f(s))$ (where σ is the standard deviation). A proper setting for the cooling rate is $c \in (0.9, 0.999)$ (Laarhoven and Aarts, 1988). For our experiments, we randomly sampled 1,000 initial solutions s_i before each SA run and set the initial temperature to $T_0 = 1.5 \times \sigma(f(s_i))$. For all runs, we set $c = 0.97$.

2.2 Iterated Local Search

Iterated local search (ILS, Lourenço et al, 2003) has so far been used in a variety of studies on local optima networks (Daolio et al, 2012; Vérel et al, 2012; Ochoa et al, 2016; Herrmann, 2016). The concept of ILS is used in many practically relevant search methods, e.g. the Chained Lin Kernighan heuristic (Applegate et al, 2003; Lin and Kernighan, 1973).

During intensification, metaheuristics focus search on promising areas of the search space, whereas during diversification, new areas are explored (Rothlauf, 2011). ILS intensifies search by performing a sequence of local search steps (standard hill climbing). As pure intensification would not allow ILS to escape local optima, it regularly performs perturbation steps to diversify search and to explore new areas of the search space. Algorithm 2 outlines the functionality of ILS, which starts with a randomly selected solution $s_0 \in S$. Then, the algorithm performs a hill climbing procedure with best improvement as selection rule (Algorithm 3): from the neighborhood $N(s)$, the solution with highest fitness is selected. The neighborhood $N(s)$ is the set of solutions that can be reached by performing an incremental change to s . The effort of this step depends on the size of the neighborhood $|N(s)|$, as it requires a scan of the whole neighborhood of s . This local search step (hill climbing) is repeated until it reaches a local optimum s^* , i.e. no further improvement is possible. Then, ILS performs a diversification step by applying a limited perturbation to the local optimum, resulting in s' . As a next step, hill climbing is started again from s' , until the next local optimum $s^{*'}$ is reached. If the new local optimum $s^{*'}$ is different from the previous s^* and has higher fitness, the algorithm has “escaped” to a new local optimum, and the change is accepted. Otherwise, another perturbation is applied to s^* . This procedure is repeated until a termination condition is met, e.g. a fixed number of escapes without any further improvement.

3 Fitness Landscape Analysis

3.1 Concept

The notion of fitness landscapes originated from evolutionary biology (Wright, 1932). The idea is that each solution has a corresponding fitness value. Combining the distances between the solutions with the fitness values of the solutions

Algorithm 2 Iterated Local Search (ILS)

Require: Solution space S ,
 Fitness function $f(S)$,
 Neighborhood function $N(S)$,
 Stopping Threshold t

- 1: $i \leftarrow 0$
- 2: Choose initial random solution $s_0 \in S$
- 3: $s^* \leftarrow \text{hillClimbBI}(s_0)$
- 4: **repeat**
- 5: $s' \leftarrow \text{perturbation}(s^*)$
- 6: $s^{*'} \leftarrow \text{hillClimbBI}(s')$
- 7: **if** $f(s^{*'}) > f(s^*)$ **then**
- 8: $s^* \leftarrow s^{*'}$
- 9: $i \leftarrow 0$
- 10: **end if**
- 11: $i \leftarrow i + 1$
- 12: **until** $i \geq t$
- 13: **return** s^*

Algorithm 3 Best Improvement Hill Climbing (hillClimbBI)

Require: Solution space S ,
 Fitness function $f(S)$,
 Neighborhood function $N(S)$,
 Initial solution s_0

- 1: $i \leftarrow 0$
- 2: **repeat**
- 3: choose x s.t. $f(x) = \max_{x \in N(s_i)}(f(x))$
- 4: **if** $f(x) > f(s_i)$ **then**
- 5: $s_{i+1} \leftarrow x$
- 6: **else**
- 7: $s_{i+1} \leftarrow s_i$
- 8: **end if**
- 9: $i \leftarrow i + 1$
- 10: **until** s_i is local optimum: $\{s \in N(s_i) : f(s) < f(s_i)\} = \emptyset$
- 11: **return** s_i

creates a landscape where each solution’s fitness is the height. In combinatorial optimization, a motivation to analyze fitness landscapes is to gain a better understanding of algorithm performance on a set of problem instances. Landscape characteristics reflect the difficulty for a variety of heuristics (Lu et al, 2014; Malan and Engelbrecht, 2014), thus problem-specific knowledge can help construct better search methods (Pitzer and Affenzeller, 2012).

3.2 Neighborhood Structure

In combinatorial optimization, a fitness landscape (Stadler, 2002) is a triplet of the search space S , the fitness function f , and the neighborhood structure $N(S)$. S contains all valid solution candidates. The fitness function $f : S \rightarrow \mathbb{R}_{\geq 0}$ assigns a fitness value to each $s \in S$. The neighborhood function

$N : S \rightarrow \mathcal{P}(S)$ assigns a set of neighbors $N(s)$ to every $s \in S$. The neighborhood structure determines the position of each s in the landscape (Reidys and Stadler, 2002). We define a distance function between all pairs of solutions s_i and s_j as

$$d : (s_i, s_j) \rightarrow \mathbb{N}_0; s_i, s_j \in S. \quad (1)$$

Usually, the distance function depends on the search operator used. For local search, we apply a small search step to solution s_i and obtain a new solution s_j , where the distance d_{s_i, s_j} is small. With d_{max} being the maximal distance between two neighboring solutions, we define the neighborhood function for local search as

$$N : s_i \rightarrow \{s_j \in S : s_j \neq s_i \wedge 0 < d(s_i, s_j) \leq d_{max}\}. \quad (2)$$

For simulated annealing, the distance between an existing s_i and new solution s_{i+1} (we denote this distance as step size) usually remains constant throughout a run. The incremental change applied to a solution results in a neighboring solution with minimal distance, $d_{max} = 1$. Iterated local search uses two types of search operators: for the hill climbing phase (hillClimbBI), the step size is usually minimal and neighboring solutions with minimal distance are created, $d_{max} = 1$. In contrast, the perturbation step (restart of search) performs a larger step through the search space where $d_{max} > 1$. This larger search step allows the search process to escape from the local optimum returned by hillClimbBI. This continuous switch of landscapes during a run makes a static analysis of ILS difficult. Despite that, it is generally accepted to study a landscape defined by a fitness function and one or several neighborhood functions (Pitzer and Affenzeller, 2012).

3.3 Local Optima

A fitness landscape can have one or more local optima. A local optimum is a solution that has no neighbors with higher fitness. For a maximization problem, we define a function

$$N_{sup}(s) = \{n \in N(s) : f(n) > (f(s))\} \quad (3)$$

which returns the neighbors n of a solution $s \in S$ that have higher fitness. As local optima have no neighbors with higher fitness, the set

$$LO = \{lo \in S : N_{sup}(lo) = \emptyset\}. \quad (4)$$

contains all local optima, which also includes one or more global optima.

3.4 Basins of Attraction

A basin of attraction is defined as the set of solution candidates from which a best improvement hill climbing algorithm converges to the same local or global optimum. Extracting the basins of a fitness landscape is required in order to calculate the edges of LONs. The function

$$B : lo \rightarrow \mathcal{P}(S \setminus LO) \quad (5)$$

assigns a subset from the power set \mathcal{P} of solutions over the search space to each local optimum $lo \in LO$, which is the basin around lo . Consequently, we use B as a function which returns the basins.

3.5 Landscape Features

Structural features of fitness landscapes are often used to predict the performance of algorithms. A well-elaborated collection of such features is given by Kallel et al (2001). Two frequently used features are ruggedness (Weinberger, 1990) and deceptiveness (Jones and Forrest, 1995). The idea of ruggedness is that the smoother the landscape is, the easier it is to search the landscape in order to find the global optimum. Ruggedness is a consequence of modality, i.e. the presence of local optima. The higher the number of local optima, the more rugged is the landscape. Ruggedness is usually measured as the correlation ρ_{nn} of fitness values between pairs of neighboring solutions (nearest-neighbor-correlation). A common way to determine an approximation of ρ_{nn} is to perform random walks across the search space and draw samples of the fitness of neighboring solutions.

A landscape is deceptive if the structure of the search space leads a search process that is guided by the fitness of solutions away from the global optimum. A common measurement for deceptiveness is the correlation ρ_{fd} between the fitness of the solutions and their distance to the global optimum (Jones and Forrest, 1995). For the calculation of the fitness-distance correlation ρ_{fd} , the global optimum must be known in advance. Often, ρ_{fd} is approximated by drawing a random sample of solutions determining the correlation between their fitness and their distance to the global optimum. A misleading landscape with a fitness-distance correlation $\rho_{fd} \approx -1$ is often referred to as a trap.

4 Local Optima Networks

LONs have been inspired by the study of energy landscapes in chemical physics (Stillinger, 1995) and used for the analysis of fitness landscapes (Ochoa et al, 2008).

4.1 Concept

A LON is a graph representation of a fitness landscape. A directed graph $G = (V, A)$ is defined by its set of vertices V and its set of directed edges (arcs²) A . In a LON graph, V contains all the local optima of the fitness landscape. The edge set A contains the directed, weighted edges that model transitions between the local optima. A variety of approaches have been proposed to compute A . For our experiments, we used the two most-studied models: LONs with edges modeling basin transition probabilities (LON_{btp}, Ochoa et al, 2010) and LONs with escape edges (LON_{ee}, Vérel et al, 2012). When analyzing fitness landscapes, we compute both LON variants per landscape. Both have the same set of nodes V , but different edge sets A_{btp} and A_{ee} , resp. For a visualization of both LON models, we refer to Figure 2, where the nodes represent local optima and edges represent either basin transition probabilities (left) or escape edges (right). The following two sections give some details on how the two LON models are computed.

4.2 LONs with Edges Representing Basin Transition Probabilities

The initial approach of local optima networks (Ochoa et al, 2008) was defined for edges that represent basin transition probabilities. Given A_{btp} as the set of directed edges, there is a weighted edge from a local optimum lo_x to a local optimum lo_y if local search can move from the basin of attraction $B(lo_x)$ to the basin $B(lo_y)$. For such a transition, both basins must be connected through a solution on the border between the two basins, such that a solution in $B(lo_y)$ with a high fitness is a neighbor to a solution in $B(lo_x)$ with lower fitness. The more such connections between two basins exist, the higher is the transition probability between their corresponding local optima, and the higher is the corresponding edge weight.

We use the following rule set for the calculation of the edge weights (Ochoa et al, 2008): the probability to perform a move from any solution s_x to another neighbor s_y depends on the number of superior neighbors of s_x . For example, simulated annealing selects a random neighbor and replaces the current solution if the neighboring solution has higher fitness. Thus, the probability to move from any solution s_x to a neighbor solution s_y can be calculated as

$$p(s_x, s_y) = \frac{1}{|N_{\text{sup}}(s_x)|}. \quad (6)$$

To determine the probability that the algorithm moves from s_x to any solution in the basin around lo_y , we compute the fraction of superior neighbors of s_x belonging to $B(lo_y)$ as

$$p(s_x, B(lo_y)) = \frac{|N_{\text{sup}}(s_x) \cap B(lo_y)|}{|N_{\text{sup}}(s_x)|}. \quad (7)$$

² We follow the standard LON terminology and denote A as the set of directed edges.

For LONs with basin transition probabilities (LON_{btp}), the edge weight between nodes lo_x and lo_y is the overall probability to move from one basin $B(lo_x)$ to another basin $B(lo_y)$. We calculate it by summing up the probabilities to move from any $s_i \in B(lo_x)$ to $B(lo_y)$:

$$p(B(lo_x), B(lo_y)) = \frac{\sum_{s_i \in B(lo_x)} p(s_i, B(lo_y))}{|B(lo_x)|}. \quad (8)$$

We calculate the transition probabilities between all local optima basins, and obtain the weights for the edges $e_{x,y} \in A_{\text{btp}}$ ($x, y \in LO$)

$$|e_{x,y}| = p(B(x), B(y)); \quad x, y \in LO. \quad (9)$$

4.3 LONs with Escape Edges

An alternative model to analyze a fitness landscape using LONs was introduced by Vérel et al (2012): LONs with escape edges. Like in the model with basin transition probabilities, V contains all the local optima of the fitness landscape. The escape edges are defined according to the distance function d of the fitness landscape (minimal number of moves between two solutions). To compute the escape edges, we need to define an integer $D > 0$, which is the step size (distance) of the perturbation steps that are, for example, used in ILS to escape from a local optimum. Consequently, there is a directed edge $E_{x,y} > 0$ between local optimum lo_x and lo_y if there exists a solution s such that

$$d(lo_x, s) \leq D \wedge s \in B(lo_y). \quad (10)$$

The weight of an escape edge should reflect the probability that an algorithm escapes from lo_x to lo_y . It is determined by the number of solutions that can be reached by one escape step (perturbation step) and belong to the basin around lo_y . Thus, the edge weight between lo_x and lo_y is calculated as the number of solutions in the basin around lo_y that can be reached by a perturbation step. This number is normalized by the total number of solutions in the basin around lo_x that can be reached. As a result, the edge weights of the edges $e_{x,y} \in A_{\text{ee}}$, where $x, y \in LO$ are calculated as

$$|e_{x,y}| = \frac{|\{s \in S \mid d(s, lo_x) \leq D \wedge s \in B(lo_y)\}|}{|\{s \in S \mid d(s, lo_x) \leq D\}|}; \quad x, y \in LO. \quad (11)$$

5 PageRank Centrality

In network analysis, the concept of centrality describes how important or influential a node is. The importance of a node depends on the network's topology and the position of the node in the network (Borgatti, 2005). For instance, Google uses the concept of PageRank centrality to assess the importance of websites (Brin and Page, 1998; Franceschet, 2011). The PageRank centrality,

which is a variant of the Eigenvector centrality, is based on the model of a user who surfs the web by randomly clicking links. The PageRank value of a website reflects the probability that the surfer currently is on this web page. There are three factors determining the PageRank of a website: the number of links it receives, the number of outgoing links of the linking web pages, and the PageRank of the linking web pages. Thus, PageRank is a recursively defined concept. For detailed information on the notion and application of PageRank, we refer to Franceschet (2011).

To calculate the PageRank of the LON nodes, we need a transition matrix Π_{btp} or Π_{ee} (here denoted as Π , as the procedure is the same for both LON models). The size of Π is $|LO| \times |LO|$. The entries $\pi_{x,y}$ ($x, y \in LO$) of Π are the edge weights $|e_{x,y}|$ of the edges $e_{x,y} \in A_{btp}$ or $e_{x,y} \in A_{ee}$, resp. As the matrix Π is stochastic, all rows and columns are normalized to sum up to 1.

For the calculation of the PageRanks, we have to consider a damping factor α , which reflects the fact that a random surfer may—instead of following links—visit a totally random page at some point. This probability is calculated as $1-\alpha$. A typical value is $\alpha = 0.85$, which says that a surfer instead of following a page link chooses a random page after about six page visits. The local search algorithms used in our experiments do not perform any random jumps in the search space, thus we set $\alpha = 1.0$ for analyzing both the LONs with basin transition probabilities and LONs with escape edges. Although ILS uses a perturbation operator, which can be viewed as a surfer that visits a random page, the escape edges in the corresponding LON model already reflect this behavior, leading to $\alpha = 1.0$.

The PageRank vector, which contains the centralities of all nodes of a LON (local optima), is the vector P , which is the Eigenvector of Π :

$$P = \Pi \times P. \quad (12)$$

If Π is a strongly connected graph, there exists a solution for P (Frobenius, 1912; Perron, 1907). These conditions are fulfilled in our case, since no negative probabilities occur by definition of the LON transition matrix. In addition, there can be no disconnected components in our LONs (i.e., all LONs were connected graphs). As P contains the PageRank centralities of all local optima, we define $P_{opt,btp}$ as the PageRank value of the global optimum in the LON with basin transition probabilities, and $P_{opt,ee}$ as the PageRank value of the global optimum in the LON with escape edges.

6 Experiments

6.1 Search Space: Kauffman NK Model

We used the well-known NK model (Kauffman and Weinberger, 1989), which is a family of combinatorial optimization problems from the class of pseudo-Boolean functions. Each instance of the model can be generated by the two parameters N and K . Each solution $s \in S$ consists of N binary decision

variables, forming a search space of $|S| = 2^N$ possible states. The fitness function

$$f_{NK} : [0, 1]^N \rightarrow [0, 1] \quad (13)$$

assigns a score to every combination of bits. It consists of N sub-functions, which assign a fitness for each bit i , depending on the state of bit i and the states of K other bits

$$f_i : [0, 1]^{K+1} \rightarrow [0, 1]. \quad (14)$$

The total fitness $f_{NK}(s)$ is the average of the values of the N sub-functions. All function values are normalized between 0 and 1, with 1.0 as the fitness of the global optimum. The parameter K determines the number of co-variables per decision variable and thus the complexity of an instance (epistasis). A value of $K = 0$ results in a decomposable and easy problem solvable with polynomial effort. $K = N - 1$ leads to a difficult problem where each decision variable can only be set to the optimal value if all other $N - 1$ co-variables are considered. Even though it is commonly accepted that a higher level of epistasis leads to higher search difficulty of landscapes, it is only a rough measure for difficulty (and an exogenous parameter in our case). Landscapes with an identical level of epistasis (same K) can have a significant variety of search difficulty. Our results on the performance of SA and ILS in Section 7.1 underpin this assumption.

The distance between two binary solutions $x, y \in S$ is calculated by the Hamming distance

$$d(x, y) = \sum_{i=0}^n |x_i - y_i| \quad (15)$$

which is the number of bits that are set to different values when comparing two solutions. For our implementation of SA and for the hill climbing procedure in ILS, we assumed that two solutions x, y are neighbors if their Hamming distance equals one ($d_{\max} = 1$). Thus, a local search step flips exactly one bit of the current solution. As perturbation operator in ILS, we flip two bits in one step.

6.2 Experimental Setting

We generated 300 instances in total of the NK model with $N = 20$ bits. To test different levels of epistasis, we used 100 instances each for $K \in \{3, 7, 10\}$. Each search space contained $2^{20} \approx 1.4M$ solutions, which is a rather small size, but allowed us to fully determine the LONs with reasonable effort. Using larger problem sizes would exponentially increase the effort for determining the LONs.

The objective of our experiments is to predict the search difficulty of landscapes for SA and ILS and to study how the choice of the LON edges affects the accuracy of performance prediction of local search based metaheuristics using PageRank centrality. In order to determine the search difficulty, we measure

Performance Variable	Description
hit_{sa} and hit_{ils}	the success rate, defined as the fraction of the 1,000 independent runs that hit the global optimum
fit_{sa} and fit_{ils}	the average fitness (solution quality) achieved
$cost_{sa}$ and $cost_{ils}$	the average number of fitness function evaluations performed by the algorithm in those runs in which the global optimum was found

Table 1: Performance measures for SA and ILS

Landscape Feature	Description
ρ_{nn}	the ruggedness of the fitness landscape, measured by the Pearson correlation between the fitness of nearest neighbors (Weinberger, 1990)
ρ_{fd}	the deceptiveness of the landscape, measured by the Pearson correlation between fitness and distance to the global optimum (Jones and Forrest, 1995)
$P_{opt,btp}$, $P_{opt,ee}$	the PageRank value of the global optimum for the LON _{btp} and the LON _{ee} model
$\overline{F_{lo}} = avg(F_{lo})$	the average fitness of all local optima
$\overline{F_{lo,btp}}$ and $\overline{F_{lo,ee}}$	the expected PageRankFitness (solution quality) as average fitness of the local optima weighted by their corresponding PageRank

Table 2: Variables to predict the performance of SA and ILS

the empirical performance of both metaheuristics. For each problem instance, we performed 1,000 independent runs of SA and ILS. The initial solutions were randomly selected. As a stopping threshold, we limited the running time by a reasonable number of function evaluations, which was 1/5th of the search space size (Daolio et al, 2012). For each NK fitness landscape, we determined three measures of performance for SA and ILS, respectively: the success rate hit , the average solution quality fit at the end of a run, and the average number $cost$ of fitness functions evaluations (cf. Table 1).

Furthermore, for each instance of the NK model, we computed the fitness landscape, the LON with basin transition probabilities, the LON with escape edges and calculated the following vectors for further analysis of the LONs:

1. P_{btp} and P_{ee} : the PageRank vectors for both LON models and
2. F_{lo} : the vector containing the fitness values of all the local optima.

To study the relationship between fitness landscape features and the search difficulty for SA and ILS, we used the statistical measures described in Table 2. We examined two relations: first, we studied whether the PageRank values of the global optimum $P_{opt,btp}$ and $P_{opt,ee}$, respectively, predict the success rates hit_{sa} and hit_{ils} in the different problem instances. Furthermore, we compared the PageRank values to the average number of function evaluations $cost_{sa}$ and $cost_{ils}$ that were used in those runs in which the global optimum was found.

Second, we aimed to predict the solution quality of SA and ILS (the average fitness). To achieve this, we calculated the fitness vectors F_{btp} and F_{ee} of the local optima as well as the PageRank vectors P_{btp} and P_{ee} for each search space, resp. We assume that the use of the PageRank vector allows us to assess the centrality of all local optima in the search space. We use the normalized PageRank centrality values ($\sum P = 1$) to calculate a weighted average of the fitness of the local optima as either

$$\overline{F_{lo,btp}} = P_{btp} \times F_{lo} \quad \text{or} \quad (16)$$

$$\overline{F_{lo,ee}} = P_{ee} \times F_{lo} \quad (17)$$

for both the LON_{btp} and LON_{ee} model, respectively. The result is a scalar value which we call the expected PageRankFitness $\overline{F_{lo,btp}}$ and $\overline{F_{lo,ee}}$, respectively. We calculate the expected PageRankFitness for all problem instances.

In our data set, each NK fitness landscape is a single observation. We assessed the predictive power of the different predictor metrics by computing one univariate, linear regression model for each combination of predictor and performance measure. This gives us a determination coefficient R^2 for each relationship between predictor and performance. The R^2 indicates the proportion of the variance in the dependent variable (in our case: performance) that is predictable from the independent variable (here: landscape feature). The R^2 value is equivalent to the *squared* Pearson correlation coefficient between two variables. As a benchmark, we have also calculated the regression models between the performance measures and the classical metrics from fitness landscapes analysis (ruggedness and deceptiveness).

We implemented our generator for NK landscapes, the extraction procedure for LONs, and the ILS algorithm in Java. For computation, we utilized 150 nodes of a HPC cluster with each 64 cores and 256 GB of RAM. We reserved one node for each instance. The computation time was approx. 1h per instance, depending on the number of local optima. To calculate the PageRank values of the nodes, we used the NetworkX Library (Hagberg et al, 2008). Our statistical analysis was conducted using the R framework (R Core Team, 2014).

7 Results

7.1 Performance of SA and ILS

As a pre-test of our experiments, we compared the performance of SA and ILS. Figure 1 plots the success rate (left), average fitness (middle), and the number of fitness function evaluations to find the global optimum (right) for ILS over SA. Each dot represents one fitness landscape, and the color (green, blue, red) indicates the level of epistasis (K). We plot results for three different levels of epistasis $K \in \{3, 7, 10\}$. The gray crossing line is a 45 degree bisector. A dot on the gray line would indicate that the performance of SA and ILS is




Fig. 1: The three performance measures of SA vs. ILS in 300 instances of NK-landscapes. The color indicates the level of epistasis (green: $K = 3$, blue: $K = 7$, red: $K = 10$): success rate (left), average fitness (middle) and running time measured as the number of fitness evaluations (right).

identical (averaged over 1,000 runs each) for this particular instance. Since the dots are all above the gray line, ILS outperforms SA in the majority of cases, which holds for success rate and average fitness. In terms of running time, ILS requires more function evaluations to locate the global optimum. Regarding the exogenous parameter for epistasis K , we can see that both algorithms have high success rates and we can expect to solve the problem with high probability if K is low (green dots). The average running time of ILS is low in such landscapes, as well. Higher values of K (blue and red dots) generally lead to lower performance.

All of these observations are as expected: higher epistasis leads to a higher search difficulty, and thus lead to a lower success rate, a lower average fitness and longer running times. The distribution of the dots also indicates that there is a high variance of all the performance measures within the different classes of K , indicating that epistasis has only limited explanatory power for search difficulty.

7.2 Visual Inspection

We exemplarily study the local optima networks of two instances of the NK model with fixed epistasis ($K = 7$). The first instance has a low difficulty for both SA and ILS, and the second has a higher difficulty. Figure 2 plots the structure of the resulting LON_{ee} (left) and LON_{btp} (right). Comparing both LON models, we can see that LON_{btp} has more links than LON_{ee} , which is in conjunction with previous studies, e.g. Vérel et al (2012).

Another observation is that in the LONs of the easy instance, there are only a few central nodes, whereas in the case of the hard instance, there are multiple central nodes. We assume that if there are many other local optima in the landscape with high fitness and central location, it is harder to find the global optimum among them. The next section performs a more systematic study of this observation.

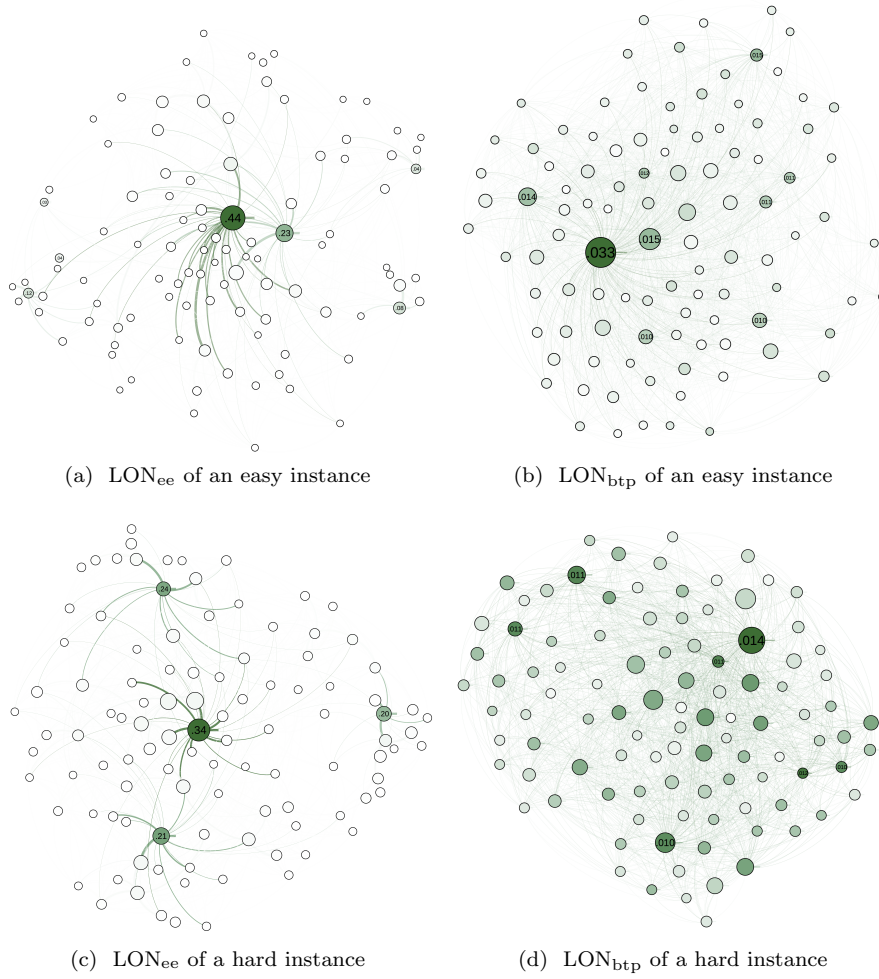


Fig. 2: Visualization of LONs with basin transition probabilities and LONs with escape edges. Each node is a local optimum. Node size indicates fitness, the color saturation ranging from white to green represents the PageRank of the local optimum (dark green indicates high PageRank). For selected nodes with high PageRank, we plot the PageRank within the circle. Figures (a) and (b) show both LONs for an NK landscape in which ILS and SA achieved high success rates (0.5 and 0.06, respectively). Figures (c) and (d) show the LONs for an NK landscape in which ILS and SA achieved significantly lower success rates (0.29 and 0.01, respectively). Both instances were chosen such that the number of local optima is similar (approx. 1000). The plots show only the best 10% of nodes in terms of fitness.

Performance of SA	Predictor	$\forall K$	K = 3	K = 7	K = 10
Success Rate: hit_{sa}	NN Correl.: ρ_{nn}	0.581	0.228	0.135	0.073
	FD Correl.: ρ_{fd}	0.52	0.162	0.061	0.034
	PageRank: $P_{opt,btp}$	0.937	0.863	0.872	0.841
	PageRank: $P_{opt,ee}$	0.504	0.144	0.307	0.095
Avg. Fitness: fit_{sa}	NN Correl.: ρ_{nn}	0.546	0.035	0.001	0.001
	FD Correl.: ρ_{fd}	0.78	0.111	0.009	0.004
	Avg. Fitn.: F_{btp}	0.901	0.513	0.85	0.964
	Exp. PRFitn.: $F_{lo,btp}$	0.988	0.914	0.978	0.994
	Exp. PRFitn.: $F_{lo,ee}$	0.741	0.196	0.518	0.802
Running time: $cost_{sa}$	NN Correl.: ρ_{nn}	0.353	0.075	0.057	0.021
	FD Correl.: ρ_{fd}	0.406	0.042	0.013	0.023
	PageRank: $P_{opt,btp}$	0.394	0.325	0.445	0.359
	PageRank: $P_{opt,ee}$	0.335	0.04	0.111	0.01
Performance of ILS					
Success Rate: hit_{ils}	NN Correl.: ρ_{nn}	0.648	0.265	0.217	0.056
	FD Correl.: ρ_{fd}	0.613	0.168	0.033	0.003
	PageRank: $P_{opt,ee}$	0.996	0.986	0.998	0.998
	PageRank: $P_{opt,btp}$	0.493	0.126	0.216	0.048
Avg. Fitness: fit_{ils}	NN Correl.: ρ_{nn}	0.438	0.241	0.026	0.005
	FD Correl.: ρ_{fd}	0.529	0.205	0.013	0.002
	Avg. Fitn.: F_{ee}	0.608	0.026	0.328	0.748
	Exp. PRFitn.: $F_{lo,ee}$	0.998	0.976	0.998	0.99
	Exp. PRFitn.: $F_{lo,btp}$	0.710	0.099	0.443	0.781
Running time: $cost_{ils}$	NN Correl.: ρ_{nn}	0.066	0.004	0.02	0.001
	FD Correl.: ρ_{fd}	0.106	0.007	0.012	0.003
	PageRank: $P_{opt,ee}$	0.0	0.103	0.276	0.323
	PageRank: $P_{opt,btp}$	0.104	0.064	0.002	0.053

Table 3: R^2 values for all linear regressions between the different performance measures (first column from left) and predictor metrics (second column from left). For example, the PageRank of the global optimum in the LON with basin transition probabilities explains 93.7% of the variance of the success rate of SA (cell is in bold and italic).

7.3 Prediction Accuracy of PageRank and Influence of LON Edges

To confirm the predictive quality of PageRank, we calculated all coefficients of determination (R^2) for each combination of predictor metric, algorithm (SA and ILS) and performance measure. Following previous work (Ochoa et al, 2014; Herrmann and Rothlauf, 2015; Herrmann, 2016), we assumed linear relationships between the two variables. We have also made separate calculations for different levels of epistasis. Table 3 presents results (R^2 values) for all combinations.

To be able to assess the prediction quality of PageRank, we compare it to the standard metrics ruggedness and deceptiveness. Over all K , ruggedness and deceptiveness explains between 52% and 62% of the variance of success rate of ILS and SA. For the average fitness, the R^2 values range from 55% to 78%. This indicates a medium to strong statistical correlation and explains

why these measures are frequently used in the literature. However, these correlations are significantly weaker with a maximum of 0.26 when the level of epistasis (K) is fixed. Moreover, this effect becomes stronger with higher levels of epistasis. If K is large ($K = 10$), the traditional metrics fail to explain the variance in the performance of SA and ILS (over all metrics, the maximum $R^2 = 0.073$). An explanation for this could be that the landscapes with high epistasis have a low variance in their ruggedness ρ_{nn} and deceptiveness ρ_{fd} . A low variance in the regressor variables then results in a low R^2 .

Next, we study the prediction quality of PageRank centrality. In analogy with Herrmann and Rothlauf (2015) and Herrmann (2016), we expect a high correlation between the PageRank of the global optimum and success rate, as well as between the expected PageRankFitness (average of fitness of local optima weighted by PageRank) and the average fitness achieved by SA and ILS. We observe the following results:

- First, we study the prediction of problem difficulty by PageRank centrality. The PageRank of the global optimum in the LON with basin transition probabilities $P_{opt,btp}$ explains approx. 94% of the variance in the success rate of SA. Obviously, the PageRank obtained from the LON_{btp} model is an accurate indicator of the search difficulty for SA. The same holds for the combination of the PageRank in the LON with escape edges $P_{opt,ee}$ and the success rate of ILS: over 99% of the variance is explained.
- The expected PageRankFitness values $\overline{F_{lo,btp}}$ and $\overline{F_{lo,ee}}$ explain more than 98% of the average fitness achieved by SA and ILS ($avg(f)$). This prediction is significantly more accurate than using the average fitness $\overline{F_{lo}}$ of the local optima. Thus, weighting the fitness values by their PageRank makes the prediction more precise since it considers the stationary probability of the algorithm visiting the local optima.
- These results for the average fitness values are robust for different levels of epistasis K . For high values of K , the R^2 is slightly reduced for both predictors, we still observe a very high correlation significantly better than traditional landscape metrics.
- Over all values of K , the number of fitness evaluations needed to locate the global optimum ($cost_{sa}$ and $cost_{ils}$) is only weakly correlated to all of the predictor metrics. Surprisingly, in the case of medium and high epistasis, the PageRank seems to predict 30-40% of the variance of running time by ILS. An explanation for this could be that in cases of high epistasis, the basins are very small. In landscapes with small basins, the running time of ILS is mainly dominated by perturbation steps, and there is nearly no hill climbing. Since the escape edges in the LON_{ee} consider such perturbations, the LON_{ee} well matches the stochastic process of ILS in such cases. Then, the LON is more likely to reflect the running time than in cases where ILS needs to spend many function evaluations for the hill climbing procedure.

Next, we examine the interplay between the type of LON edges and the ability of the PageRank centrality for the resulting LON to predict the performance of local search based metaheuristics. Here, we observe that to achieve an accurate

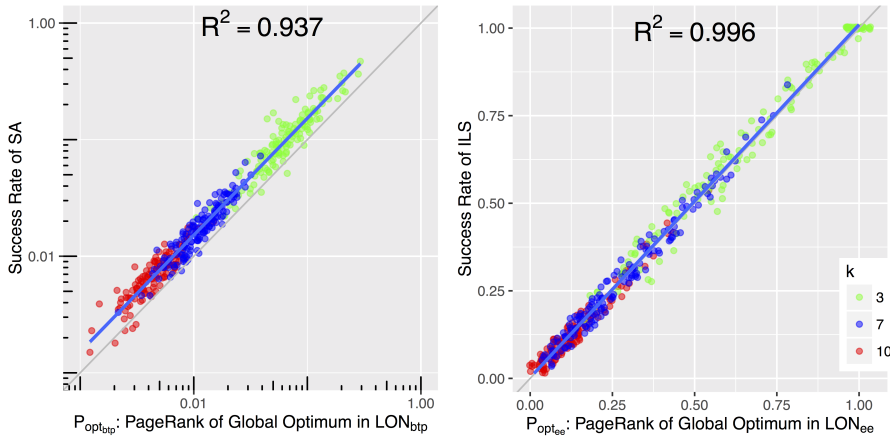


Fig. 3: The success rates of SA (left) and ILS (right) over the PageRank of the optimum. The LON model corresponds to the search operator. Each of the dots represents a single problem instance.

prediction, the definition of the LON edges must properly reflect the type of diversification steps used in the metaheuristic. This can be seen when studying the R^2 values for the combination of the success rate of SA and $P_{opt,ee}$, as well as ILS and $P_{opt,btp}$, which are significantly lower (approx. 50%) than the results on PageRank with corresponding LON models. The explanation is that LONs with basin transition probabilities better reflect the behavior of simulated annealing than escape edges. This is also true the other way round as escape edges better capture the dynamics of ILS in comparison to edges with basin transition probabilities.

We study this observation in more detail. Figure 3 plots the empirical success rates of SA and ILS over the PageRank of the optimum. For SA, we use a LON with edges modeling basin transition probabilities; for ILS, we use a LON with edges representing escape edges. Thus, the definition of the LON fits well to the corresponding local search based metaheuristic. Each dot in the plot represents one problem instance (search space). As expected, both are strongly correlated. In contrast, Figure 4 plots the empirical success rates of SA and ILS over the PageRank of the optimum, where the PageRank values are calculated for the non-corresponding LON model. This means that the success rate of SA is predicted for a LON with escape edges and the success rate of ILS is predicted using a LON with edges modeling basin transition probabilities. The results show that when using the non-corresponding LON model, the PageRank is no accurate predictor of expected search performance any more. With $R^2 \approx 0.5$ the prediction quality is similar to traditional measures like ruggedness or deceptiveness and quite lower in comparison to the case where the definition of the LON edges fits well to the search behavior of the local search based metaheuristic.

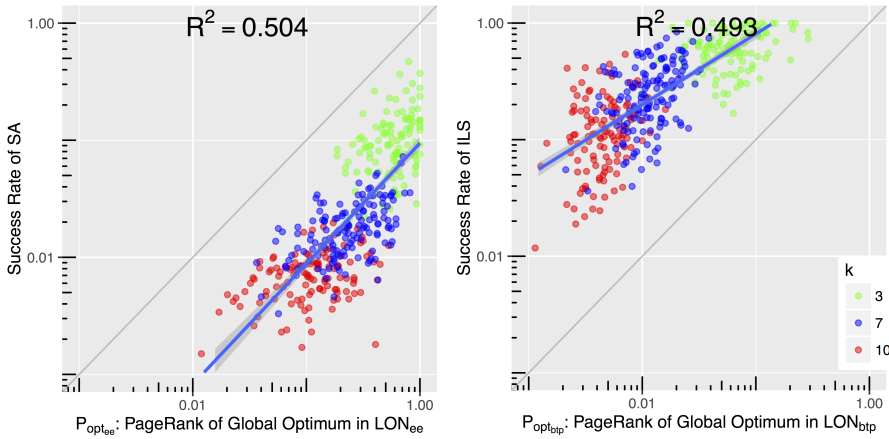


Fig. 4: The success rates of SA (left) and ILS (right) over the PageRank of the optimum. In contrast to Figure 3, the PageRank values were obtained from the non-corresponding LON models. Each of the dots represents a single problem instance.

In summary, we find that the PageRank of the global optimum in LONs very well predicts the search difficulty of landscapes, i.e. the empirical success rate. Moreover, we found that the stationary distribution of the PageRank vector over all local optima is useful to make better predictions about the expected solution quality obtained when running SA or ILS. Both predictions work for all levels of epistasis (minimal $R^2 = 0.841$), which is a clear advantage to the concepts of ruggedness and deceptiveness. However, it is important to choose the appropriate model for the edges in the LON. Since the edges represent probabilistic transitions between local optima basins, they must properly reflect the behavior of the search algorithm.

8 Conclusions, Limitations, and Future Work

This work focused on the PageRank centrality of local optima and studied how the definition of the LON edges affects the accuracy of the PageRank centrality to predict the success rate and average fitness achieved by local search based metaheuristics like iterated local search and simulated annealing. Furthermore, the study summarized and extended previous work on PageRank centrality for LONs (Ochoa et al, 2014; Herrmann and Rothlauf, 2015) and directly compared the prediction quality of PageRank for LONs where edges describe either basin transition probabilities or escape edges.

We find that in order to achieve accurate predictions, the LON edges must properly reflect the diversification mechanisms of the search algorithm. In particular, LONs with edges for basin transition probabilities capture well

the diversification mechanism of simulated annealing. SA mainly relies on local search by always accepting improvements and only sometimes accepting worse solutions. This is well captured by the concept of basin transition probabilities, which measure the probability that local search passes from one basin of attraction to a different one. On the other hand, LONs with escape edges well capture the search dynamics of iterated local search. Like SA, iterated local search mainly relies on local search, however, the mechanism of escaping from local optima is different from SA. After finding a local optimum, ILS performs some perturbation step. This behavior of ILS to escape from local optima is well captured by LONs with escape edges since the probability to escape from a local optimum is determined by the number of possible escape edges.

The findings confirm that the PageRank centrality of local optima can be an accurate predictor of the success rate and average fitness achieved by metaheuristics. This is because LONs are an approximation of the fitness landscape's Markov chain, and the PageRank vector reflects the stationary distribution of the states in this chain. However, as LONs are stochastic models of an algorithm in a landscape, the LON edges must reflect the behavior of the algorithm. If the LON edges properly capture the diversification mechanisms of the local search based metaheuristic, PageRank centrality can be a high-quality predictor of search difficulty; contrary, if there is a mismatch between the definition of LON edges and a metaheuristic's diversification mechanism, the prediction accuracy of PageRank (like all other measures based on network analysis) will be quite low.

An obvious limitation of our study is the size of the problem instances used. Even though we are convinced that our results also hold for larger instances, it would be interesting to perform further examinations on this, e.g. by sampling the local optima instead of evolving the whole search space.

Another limitation is of more fundamental nature: our cross-check clearly indicates that it is not possible to make general statements on the performance of an algorithm with an arbitrary LON model. Instead, the LON model must closely match the dynamics of the search method. For instance, in the case of ILS, the definition of the escape edges must properly consider the distance of the perturbation step. Our finding that the definition of LON edges is metaheuristic-specific can also be used for a more thorough studying and understanding of the search dynamics of other or new local search operators. If the PageRank is a good predictor of the expected search performance when applying metaheuristics to combinatorial optimization problems, then the metaheuristic's diversification mechanism is well captured by the definition of the LON edges. Contrary, if the prediction quality of the PageRank centrality is low for some given definition of LON edges, then the assumptions made for the definition of LON edges do not fit to the diversification mechanisms of the metaheuristic. Thus, we encourage researchers to come up with more and different types of LON types and to study to which types of metaheuristic they fit.

Acknowledgements G. Ochoa acknowledges support from the Leverhulme Trust [award number RPG-2015-395].

References

- Applegate D, Cook W, Rohe A (2003) Chained Lin-Kernighan for Large Traveling Salesman Problems. *INFORMS Journal on Computing* 15(1):82–92, DOI 10.1287/ijoc.15.1.82.15157
- Borgatti SP (2005) Centrality and network flow. *Social Networks* 27(1):55–71, DOI 10.1016/j.socnet.2004.11.008
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems* 30(1-7):107–117, DOI 10.1016/S0169-7552(98)00110-X
- Chicano F, Daolio F, Ochoa G, Vérel S, Tomassini M, Alba E (2012) Local Optima Networks, Landscape Autocorrelation and Heuristic Search Performance. In: Coello CAC, Cutello V, Deb K, Forrest S, Nicosia G, Pavone M (eds) *Parallel Problem Solving from Nature - PPSN XII: 12th International Conference*, Springer, Berlin and Heidelberg, vol 7492 LNCS, pp 337–347, DOI 10.1007/978-3-642-32964-7_34
- Daolio F, Verel S, Ochoa G, Tomassini M (2012) Local optima networks and the performance of iterated local search. In: Soule T (ed) *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference - GECCO '12*, ACM Press, Philadelphia, Pennsylvania, USA, p 369, DOI 10.1145/2330163.2330217
- Franceschet M (2011) PageRank: standing on the shoulders of giants. *Communications of the ACM* 54(6):92–101, DOI 10.1145/1953122.1953146
- Frobenius FG (1912) Ueber Matrizen aus nicht negativen Elementen. *Sitzungsberichte Preussische Akademie der Wissenschaft*, Berlin pp 456–477
- Hagberg AA, Schult DA, Swart PJ (2008) Exploring network structure, dynamics, and function using NetworkX. *Proceedings of the 7th Python in Science Conference (SciPy 2008)* pp 11–15
- Herrmann S (2016) Determining the difficulty of landscapes by pagerank centrality in local optima networks. In: Chicano F, Hu B, García-Sánchez P (eds) *Evolutionary Computation in Combinatorial Optimization: 16th European Conference, EvoCOP 2016*, Porto, Portugal, March 30 – April 1, 2016, *Proceedings*, Springer International Publishing, Cham, pp 74–87, DOI 10.1007/978-3-319-30698-8_6, URL http://dx.doi.org/10.1007/978-3-319-30698-8_6
- Herrmann S, Rothlauf F (2015) Predicting Heuristic Search Performance with PageRank Centrality in Local Optima Networks. In: Silva S (ed) *Proceedings of the 2015 Genetic and Evolutionary Computation Conference - GECCO '15*, ACM Press, Madrid, Spain, pp 401–408, DOI 10.1145/2739480.2754691
- Jones T, Forrest S (1995) Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms. In: *Proceedings of the Sixth Inter-*

- national Conference on Genetic Algorithms, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp 184–192
- Kallel L, Naudts B, Reeves CR (2001) Properties of Fitness Functions and Search Landscapes. In: Kallel L, Naudts B, Rogers A (eds) *Theoretical Aspects of Evolutionary Computing*, Natural Computing Series, Springer, Berlin and Heidelberg, pp 175–206, DOI 10.1007/978-3-662-04448-3_8
- Kauffman SA, Weinberger ED (1989) The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* 141(2):211–245, DOI 10.1016/S0022-5193(89)80019-0
- Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220(4598):671–80, DOI 10.1126/science.220.4598.671
- Laarhoven PJM, Aarts EHL (1988) *Simulated Annealing: Theory and Applications*. Kluwer Academic Publishers, Norwell, MA, USA
- Lin S, Kernighan BW (1973) An Effective Heuristic Algorithm for the Traveling-Salesman Problem. *Operations Research* 21(2):498–516, DOI 10.1287/opre.21.2.498
- Lourenço HR, Martin OC, Stützle T (2003) Iterated Local Search. In: *Handbook of Metaheuristics*, Kluwer Academic Publishers, Boston, pp 320–353, DOI 10.1007/0-306-48056-5_11
- Lu G, Li J, Yao X (2014) Fitness Landscapes and Problem Difficulty in Evolutionary Algorithms: From Theory to Applications. In: Richter H, Engelbrecht A (eds) *Recent Advances in the Theory and Application of Fitness Landscapes, Emergence, Complexity and Computation*, Springer, Berlin and Heidelberg, pp 133–152, DOI 10.1007/978-3-642-41888-4
- Malan KM, Engelbrecht AP (2014) Fitness Landscape Analysis for Metaheuristic Performance Prediction. In: Richter H, Engelbrecht A (eds) *Recent Advances in the Theory and Application of Fitness Landscapes*, Springer, Berlin and Heidelberg, pp 103–129, DOI 10.1007/978-3-642-41888-4_9
- Ochoa G, Tomassini M, Vérel S, Darabos C (2008) A study of NK landscapes’ basins and local optima networks. In: Keijzer M (ed) *Proceedings of the 10th annual conference on Genetic and evolutionary computation - GECCO ’08*, ACM Press, Atlanta, GA, USA, pp 555–562, DOI 10.1145/1389095.1389204
- Ochoa G, Verel S, Tomassini M (2010) First-improvement vs. best-improvement local optima networks of nk landscapes. In: Schaefer R, Cotta C, Kolodziej J, Rudolph G (eds) *PPSN’10: Proceedings of the 11th International Conference on Parallel Problem Solving from Nature*, Springer, Kraków, Poland, vol I, pp 104–113
- Ochoa G, Verel S, Daolio F, Tomassini M (2014) 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, Springer, Berlin and Heidelberg, pp 233–262, DOI 10.1007/978-3-642-41888-4_9
- Ochoa G, Chicano F, Tinós R, Whitley D (2015) Tunnelling Crossover Networks. In: Silva S (ed) *Proceedings of the 2015 Genetic and Evolutionary Computation Conference - GECCO ’15*, ACM Press, Madrid, Spain, pp 449–456, DOI 10.1145/2739480.2754657

- Ochoa G, Veerapen N, Whitley D, Burke EK (2016) The Multi-Funnel Structure of TSP Fitness Landscapes: A Visual Exploration. In: Artificial Evolution: 12th International Conference, Evolution Artificielle, EA 2015, Springer International Publishing, Lyon, pp 1–13, DOI 10.1007/978-3-319-31471-6_1
- Perron O (1907) Zur Theorie der Matrices. *Mathematische Annalen* 1 64(1):248–263
- Pitzer E, Affenzeller M (2012) A Comprehensive Survey on Fitness Landscape Analysis. In: Fodor J, Klempous R, Suárez Araujo CP (eds) Recent Advances in Intelligent Engineering Systems, Studies in Computational Intelligence, vol 378, Springer, Berlin and Heidelberg, pp 161–191, DOI 10.1007/978-3-642-23229-9
- R Core Team (2014) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, URL <http://www.R-project.org/>
- Reidys CM, Stadler PF (2002) Combinatorial Landscapes. *SIAM Review* 44(1):3–54, DOI 10.1137/S0036144501395952
- Rothlauf F (2011) Design of modern heuristics: Principles and application. Springer, Berlin and Heidelberg
- Stadler PF (2002) Fitness landscapes. *Applied Mathematics and Computation* 117:187–207
- Stillinger FH (1995) A Topographic View of Supercooled Liquids and Glass Formation. *Science* 267(5206):1935–1939, DOI 10.1126/science.267.5206.1935
- Vérel S, Daolio F, Ochoa G, Tomassini M (2012) Local optima networks with escape edges. In: Artificial Evolution, Springer, Angers, France, pp 49–60, DOI 10.1007/978-3-642-35533-2_5
- Weinberger ED (1990) Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological cybernetics* 336:325–336
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In: Jones DF (ed) Proceedings of the 6th International Congress of Genetics, Morgan Kaufmann Publishers Inc., Ithaca, New York, pp 356–366