Local Search Pivoting Rules and the Landscape Global Structure

Sara Tari
Univ. Littoral Côte d'Opale
LISIC, F-62100 Calais, France
sara.tari@univ-littoral.fr

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

ABSTRACT

In local search algorithms, the pivoting rule determines which neighboring solution to select and thus strongly influences the behavior of the algorithm and its capacity to sample good-quality local optima. The classical pivoting rules are first and best improvement, with alternative rules such as worst improvement and maximum expansion recently studied on hill-climbing algorithms. This article conducts a thorough empirical comparison of five pivoting rules (best, first, worst, approximated worst and maximum expansion) on two benchmark combinatorial problems, NK landscapes and the unconstrained binary quadratic problem (UBQP), with varied sizes and ruggedness. We present both a performance analysis of the alternative pivoting rules within an iterated local search (ILS) framework and a fitness landscape analysis and visualization using local optima networks. Our results reveal that the performance of the pivoting rules within an ILS framework may differ from their performance as single climbers and that worst improvement and maximum expansion can outperform classical pivoting rules.

CCS CONCEPTS

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

KEYWORDS

Local Search, Iterated Local Search, Pivoting Rules, Local Optima Networks

ACM Reference Format:
Sara Tari and Gabriela Ochoa. 2021. Local Search Pivoting Rules and the Landscape Global Structure. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3449639.3459295

1 INTRODUCTION

Local search algorithms are widely-used to tackle combinatorial optimization problems. Their ability to find good-quality solutions varies according to several factors such as their components, the neighborhood-relation, the budget allotted to the search, and the instance to optimize. A way to improve our understanding of local search behavior is to empirically study them using fitness landscapes as a tool for their analysis. This concept helps to visualize and measure characteristics of the search space according to the neighborhood-relation under consideration. Among the various techniques available with fitness landscapes, *local optima networks* (LONs) are some of the most capable of displaying the global landscape structure and are well-suited to analyze the behavior induced by different components. One of the critical components of a local search is its *pivoting rule* that determines which neighbor to select at each step of the search. This rule plays an essential role in guiding the search and, therefore, in its ability to find good-quality solutions. This ability may vary according to the instance to tackle, but also when the pivoting rules interact with other local search components. As these interactions can be numerous and complex, deconstructing algorithms can help to build them better in the future. For example, knowing the impact of pivoting rules in straightforward local searches as hill-climbers is helpful, as these can be whole mechanisms in more elaborate optimization algorithms.

Two widely-used classical pivoting rules are the first and best improvement. They have both been empirically compared on various contexts [6, 19, 26], but also within single hill-climbers on large NK landscapes of different ruggedness levels in [2]. This work show that performing smaller steps during the climbing process with first improvement often leads to better-quality local optima, which led to the proposition of the *worst improvement* that performs the slightest improvement at each step of the search [3]. Within single hill-climbers the worst improvement outperforms classical pivoting rules on non-smooth landscapes, and seem to avoid being trapped prematurely into low-quality local optima. This work led to the proposition of the *maximum expansion* that explicitly seeks to keep the most improving possibilities during the climbing process [23]. The maximum expansion outperforms the aforementioned pivoting rules on large NK and UBQP landscapes, including classical pivoting rules using the same amount of knowledge to navigate the landscape (i.e., the same extended neighborhood vision). The approximation of the worst improvement proposed in [3] was compared to the first improvement within an iterated local search framework (ILS) for the same evaluation budget to observe whether alternative pivoting rules can be competitive for the same budget of evaluations [24] While first improvement performs better on large NK landscapes, the approximation of worst improvement outperforms first improvement on large UBQP landscapes. A landscape analysis of the ruggedness levels highlighted that UBQP landscapes are locally rugged but globally smooth, in contrast to the uniform ruggedness repartition of NK landscapes.

The work in this paper follows these previous studies on pivoting rules. As hill-climbers can be components within more powerful local search, the natural follow-up in the study of alternative pivoting rules is to observe what happens when hill-climbers interact with another component. Here, we focus on an iterated local search framework to observe the differences in the relative efficiency of pivoting rules when climbers are combined with perturbations. Moreover, we use LONs to understand better the behavior induced by these different pivoting rules and compare the global structures of some NK and UBQP landscapes of the aforementioned studies.

Indeed, LONs are often more useful than an estimation of ruggedness performed by indicators to display the differences in the global structure of landscapes. The contributions of the paper are (1) an experimental comparison of classical and alternative pivoting rules within iterated local search algorithms, (2) a comparison of the local optima networks built with these different pivoting rules, (3) a comparison of the local optima networks of NK landscapes and UBQP landscapes.

The paper is organized as follows. Section 2 presents the concepts related to fitness landscapes and local optima networks, and the landscapes used in this study. We then present the iterated local search framework we focus on, as well as the pivoting rules in section 3. In section 4 we provide the experimental protocol and the results of the experimental analysis of pivoting rules within an ILS. The local optima network analysis is provided in section 5. Finally, we discuss this work and point out some perspectives.

2 FITNESS LANDSCAPES

Fitness landscapes were introduced by Wright [28] in the biology field to depict the evolution process and are nowadays used in several fields to study complex systems [1, 12]. This concept is also widely used in evolutionary computation, where the analogy between an individual with a fitness and a solution with a score is easy to make. Fitness landscapes then provide a way to abstract from the problem instances and to study the behavior of neighborhood-based methods in a theoretical or empirical way.

2.1 Definitions

A fitness landscape is a triplet $(\mathcal{X}, \mathcal{N}, \mathfrak{f})$, where $\mathcal{X}$ refers to the search space, $\mathcal{N} : \mathcal{X} \rightarrow 2^{\mathcal{X}}$ to the neighborhood relation that connects a set of neighboring solutions to each candidate solution, and $\mathfrak{f}$ is a fitness function which, by associating a value to each solution, provides a measure of its quality. A fitness landscape corresponds to a graph where each vertex corresponds to a solution in the search space, and edges represent the neighborhood relations between the different solutions.

In a maximization context, the best solution of a fitness landscape, called global optimum, has the highest fitness value and corresponds in a pictorial way to the highest peak (definition 2.1). The peaks of a fitness landscape are called local optima and correspond to solutions with no improving neighbors. Among those solutions, those without same-fitness (or neutral) neighbors are called strict local optima (definition 2.2). The *basin of attraction* of a local optimum is defined by all the solutions from which the local optimum is reachable by applying improving mutations only. Thus, the probability of reaching a given local optimum correlates with the size of its basin.

Definition 2.1. A global optimum of a landscape $P = (\mathcal{X}, \mathcal{N}, \mathfrak{f})$ is a solution such that $x^* \in \text{argmax}_{x \in \mathcal{X}} \mathfrak{f}(x)$.

Definition 2.2. A (strict) local optimum of a landscape is a solution x such that $\forall x' \in \mathcal{N}(x), \mathfrak{f}(x')(<) \leq \mathfrak{f}(x)$.

While it is common to use algebraic properties of the problem in fitness landscape analysis [22, 27], another primary approach consists of extracting statistical properties to characterize a landscape from a sample of solutions. Some of the most studied properties with this approach are ruggedness, neutrality, and dimension. Ruggedness mainly refers to the number and distribution of local optima within the landscape and the size of their basins of attraction. A *rugged* landscape has many local optima with small basins of attraction; a *smooth* landscape has a few local optima with large basins of attraction. As a result, finding good-quality solutions is often easier on smooth landscapes and harder on rugged landscapes. One of the best-known indicators used to estimate the ruggedness levels is the autocorrelation function [25]. Various other properties are used to characterize landscapes, most of which are described and discussed in [14, 15, 21].

Those different indicators give information on some properties of landscapes; however, the resulting output generally does not take into account the structural variations of different landscape areas that may induce difficulties for some optimization algorithms. Network-based models such as local optima networks can highlight such differences.

2.2 Local Optima Networks

Local optima networks (LONs) [16] are a model of fitness landscapes suited to characterizing their global structure. LONs convey a compressed view of a landscape as a network (or graph) where nodes are local optima, according to a given neighborhood, and edges are possible transitions among optima according to a perturbation or escape operator. In this article, we use a recently proposed LON variant, the monotonic LON (MLON) model [18], in which edges are restricted to non-deteriorating transitions. To specify the model, we need to define the nodes and edges. The relevant definitions are given below, and the process for sampling and empirically constructing the models is described later in Section 5.

Nodes. The nodes correspond to local (and global) optima according to a given neighborhood $\mathcal{N}$, as defined above (Definitions 2.1 and 2.2). The set of nodes is denoted as L.

Monotonic perturbation edges. There is an edge from local optimum l_1 to local optimum l_2, if l_2 can be obtained after applying a random perturbation ($\mathcal{P}$ 1-flips in this paper) to l_1 followed by local search, and $f(l_2) \geq f(l_1)$. These edges are called *monotonic* as they record only non-deteriorating transitions between local optima. Edges are weighted with frequencies of transition computed during the construction process (described in Section 5). The weight is the number of times a transition between two local optima occurred. The set of edges is denoted by E.

Monotonic LON. is the directed graph MLON = (L, E), with node set L, and edge set E as defined above.

2.3 Problems

Here, we consider large landscapes derived from two pseudo-Boolean combinatorial optimization problems: NK landscapes [10, 11] that display tunable ruggedness levels, and the unconstrained binary quadratic problem (UBQP) that can be used to reformulate a wide range of real problems [5]

NK landscapes. These benchmark problems are determined by means of two parameters: N and K. N determines the number of decision variables and directly affects the size of the search

space (2^N solutions), while K determines the interdependency levels between variables. When $K = 0$, the resulting landscape is entirely smooth, and there is no variable interdependency. When $K = N - 1$, the resulting landscape is entirely rugged and corresponds to a landscape whose peaks are randomly determined.

We use a set of instances with various sizes $N \in \{128, 256, 512\}$ and various levels of variable interdependency $K \in \{2, 4, 6, 8, 10\}$.

UBQP. An instance of UBQP is composed of a matrix Q containing $N \times N$ integers. A solution is a bit-string x of size N, where each $x_i \in \{0, 1\}$ corresponds to the i^{th} bit of x.

We use a set of instances with $N \in \{128, 256, 512\}$ and different densities generated with the generator proposed in [20]. The density d affects the rate of zeros in the matrix Q: $d = 0$ leads to a matrix filled with zeros, except on its diagonal, whereas $d = 100$ leads to a matrix without zeros (except on its diagonal).

For both problems, the neighborhood relation under consideration is the 1-flip operator, which consists of flipping the value of a single bit in the solution.

3 ALGORITHMS

A local search algorithm consists of iteratively modifying a single solution following a given move-policy until a stopping criterion is reached. This class of metaheuristics is naturally more exploitative, whereas population-based metaheuristics tend to explore more.

Four aspects mainly determine local search algorithms: the initial solution, the neighborhood relation, the move policy, and the stopping criterion. The move policy determines which neighboring solution to select and thus strongly influences the behavior of local search algorithms. As the neighborhood relation, the move policy strongly influences the algorithm capacity to determine a good-quality search space sample. Many local search types exist, and their difference mainly lies in the move policy, which can be more or less sophisticated.

3.1 Local Search

Strict hill-climbing algorithms correspond to a class of local search algorithms that only select improving neighbors at each step of the search.The neighbor x' of a solution x is improving if its fitness value is strictly superior. Strict hill-climbers naturally stop when the current solution is a local optimum. Note that non-strict hill-climbers also accept neutral neighbors, but in that case, determining the stopping criterion can be more challenging.

Only accepting improving moves is an incomplete move-policy as there are often several improving neighbors for a given solution. Thus, many hill-climbers can be distinguished according to their sub-move-policy on a non-deteriorating neighbor. We elaborate this aspect in section 3.2.

As hill-climbers are straightforward, they are often used as a component of more sophisticated metaheuristics such as iterated local search (ILS) [8, 13]. A generic ILS is determined by a local search, a perturbation phase, and an acceptance criterion. The local search returns a solution, and the acceptance criterion determines if this solution is stored, then the perturbation phase applies changes to the last solution stored. The process is repeated until the stopping criterion is met.

In this paper, we focus on a straightforward ILS described in algorithm 1. This ILS alternates a strict hill-climber and a perturbation phase that applies $\mathcal{P}$ random moves to the last-encountered local optimum, meaning that the last stored solution is always replaced (except when consecutive hill-climbers reach the same local optimum). While hill-climbers can be viewed as an exploitative process, the perturbation phase adds diversity to the search by allowing the deterioration of the current solution. It can help escape the basins of attraction of local optima to reach promising fitness landscape areas. In the following, we study the impact of different pivoting rules on the behavior of this ILS.

Algorithm 1 Iterated local search used in this work.

```
Choose x_0 ∈ X (initialization)
x ← x_0
x* ← x
x ← HillClimber(x)
if f(x) > f(x*) then
    x* ← x
end if
while stopping criterion not reached do
    x ← applyRandomMoves(P, x)
    x ← HillClimber(x)
    if f(x) > f(x*) then
        x* ← x
    end if
end while
return x*
```

3.2 Pivoting Rules

To be complete, the move policy of a hill-climbing algorithm must determine the improving neighbor to select at each step of the search. This selection strategy, also called *pivoting rule*, strongly influences a given hill-climber ability to reach good-quality solutions. In this work, we focus on the following pivoting rules:

- The *best improvement* (B) is a deterministic pivoting rule that consists of selecting the improving neighbor with the highest fitness value.
- The *first improvement* (F) is a stochastic pivoting rule that consists of selecting the first evaluated improving neighbor.
- The *worst improvement* [3] (W), mentioned under the name of least ascent in [9], is deterministic and consists of selecting the improving neighbor with the lowest fitness value.
- The *approximated worst improvement* [3] (W_κ) is a stochastic strategy halfway between the first and worst improvement. It consists of selecting the solution with the lowest fitness value among κ randomly generated improving neighbors.
- The *maximum expansion* [23] (E) consists of selecting the improving neighbor with the highest *expansion score*, i.e. with the highest number of improving solutions.

The first and best improvement are widely used in the literature and correspond to what we can call *classic pivoting rules*. The best improvement is an intuitive way of selecting a neighbor and requires to evaluate the whole neighborhood of the current solution. The first improvement is particularly competitive as it often requires fewer evaluations than best improvement to reach local

optima. The other three pivoting rules are rarely encountered in the literature and are referred to as *alternative pivoting rules* is the following. The worst improvement requires to evaluate the whole neighborhood of the current solution and often requires more steps than classical pivoting rules. The approximated worst improvement was proposed to keep the principle and efficiency of the worst improvement while reducing the run duration. The maximum expansion was proposed to observe the effect of a move policy that explicitly seeks to keep the most improving capacities. The calculation of the expansion score requires the evaluation of solutions distant from 1 and 2 mutations of the current solution. This rule was previously competed against classical pivoting rules and adaptations of classical pivoting rules in an extended neighborhood context to ensure the relevance of the expansion criterion. Note that in the case of maximum expansion score equality, we consider the version that selects the improving neighbor with the best fitness value among those of maximal expansion score.

4 PERFORMANCE ANALYSIS

4.1 Experimental Protocol

These experiments investigate the efficiency of the pivoting rules presented in section 3.2 within the ILS process described by algorithm 1. In particular, we focus on the first, best, and worst improvement, the approximated worst using $\kappa = 2$ ($\mathtt{W}_2$) and the maximum expansion. To observe the differences induced by the number of perturbations of the ILS, we use $\mathcal{P} = \{5, 10\}$ for both problems.

For each triplet ($\mathcal{P}$, pivoting rule, landscape), we conducted 30 runs from the same set of 30 randomly generated solutions. We use a given number of climbing processes as stopping criterion: 10 000 on NK landscapes and 5 000 on UBQP landscapes, as these landscapes were found to be easier to solve [24]. Indeed, most of the ILS variants considered here attain the known global optimum on UBQP landscapes, which is not the case on NK landscapes. Note that we use the number of climbing processes instead of a number of evaluations to compare the pivoting rules to (1) construct LONs from the same number of local optima for each pivoting rule, (2) to observe the differences for the same number of perturbations phases. While such results are less interesting in practice than running the ILS for the same amount of time, it allows us to observe the relative efficiency differences between the pivoting rules with perturbations compared to the one observed within single hill-climbers.

4.2 Results

NK landscapes. Results of the experiments are reported in table 1. Following these results, the worst improvement is systematically less efficient than first improvement to reach good-quality local optima. This particular result differs from the results on single climbing processes over most of these landscapes. $\mathtt{W}_2$ is usually outperformed by F but outperforms W. This result is coherent with the variation in efficiency of the two aforementioned pivoting rules. F was already more efficient than $\mathtt{W}_2$ within these ILS for the same number of evaluations and than in this configuration, the higher number of evaluations performed by worst improvement does not induce a different relative efficiency. A hypothesis would be that using worst improvement within climbing processes is more efficient when starting for a low or medium-quality solution. The maximum expansion outperforms most of the pivoting rules on these landscapes, as previously observed within single hill-climbers. Due to the evaluation of two neighborhood levels, this strategy requires a significantly higher number of evaluations than other pivoting rules to achieve 10 000 climbing processes. However, W requires a significantly higher number of evaluations than F but achieves inferior results in this context. In climbers only, the efficiency of the maximum expansion and the worst improvement was almost always positively correlated. Our results show limitations to the hypothesis that correlates the behavior of these two pivoting rules.

When using 5 perturbations, B is outperformed by several pivoting rules, except when $N = 128$ and on smooth landscapes $K = 2$. However, when using 10 perturbations, this pivoting rule outperforms several of the other pivoting rules and is the only one to generally perform better with more perturbations. B seems to require more perturbations to escape from the basins of attractions than the other pivoting rules.

UBQP landscapes. Results are reported in table 2. When using 5 perturbations W and $\mathtt{W}_2$ always reach the global optimum. It is the case of E and F on most landscapes. On the contrary, it is seldom the case for B. These results are globally consistent with previous results using climbers only or comparing F and $\mathtt{W}_2$ within an ILS for the same number of perturbations. On these landscapes, W and its variant remain among the most efficient pivoting rules. When using $\mathcal{P} = 10$ perturbations, the different pivoting rules always reach the global optimum except B when $N > 128$. Nevertheless, using a larger number of perturbations also seems beneficial for this last pivoting rule, as well as F, B, $\mathtt{W}_2$ and E which all require less climbing processes, while it is the opposite for W.

5 LOCAL OPTIMA NETWORK ANALYSIS

To further understand the performance differences among the five pivoting rules, we studied the induced fitness landscapes using the monotonic local optima network model (MLON) defined in section 2.2. The MLONs were sampled and constructed for a representative set of instances with $N = 256$ for the two studied problems. The number of perturbations considered was $\mathcal{P} = 5$. For each of these instances a MLON was constructed by aggregating all the unique nodes and edges encountered across 30 independent runs of the ILS algorithm (Alg. 1) with each pivoting rule.

5.1 Network Metrics

We aim to identify fitness landscape features that correlate with and help to explain the performance differences among the pivoting rules. Many network properties can be computed from LONs, the most basic are the number of nodes and edges, but a variety of metrics have been studied [16–18]. We focus here on four network properties that help us to asses the landscapes global structure. These metrics are summarized in Table 3. For characterizing the multi-funnel structure, we measure the size of the connected component containing the best found solution (*best-comp*). Connected components have been associated to funnel structures [17]. A funnel refers to a grouping of local optima, forming a coarse-level gradient towards a high quality solution at the end. Funnels can be considered as basins of attraction at the level of local optima. When sub-optimal funnels exist, search can get trapped and fail to

Table 1 Best local optimum found and pivoting rules performance within ILS on each NK landscapes over 30 runs. For each pivoting rule, the average fitness and a rank are given. The ranking for each (landscape, $\mathcal{P}$) is obtained according to a Mann-Whitney statistical test with $\alpha = 0.05$ and a Bonferroni correction. Grey cells indicate a statistical dominance between the same pivoting rule using a different number of perturbations. Purple underline values indicate whether the best local optimum is found on at least one run.

land.	best LO	$\mathcal{P} = 5$					$\mathcal{P} = 10$				
		F	B	W	W_2	E	F	B	W	W_2	E
128 2	0.74237	0.7402 2	0.7409 1	0.7363 5	0.7384 4	0.7401 2	0.7405 2	0.7417 1	0.7370 5	0.7391 4	0.7407 2
128 4	0.79586	0.7943 2	0.7910 3	0.7918 4	0.7929 3	0.7958 1	0.7933 3	0.7951 2	0.7890 5	0.7906 4	0.7957 1
128 6	0.80045	0.7908 2	0.7872 3	0.7855 4	0.7884 3	0.7983 1	0.7871 3	0.7929 2	0.7792 5	0.7826 4	0.7979 1
128 8	0.80257	0.7851 2	0.7838 2	0.7786 5	0.7817 3	0.7952 1	0.7784 3	0.7850 2	0.7758 4	0.7766 3	0.7920 1
128 10	0.79445	0.7748 2	0.7752 2	0.7697 4	0.7712 4	0.7863 1	0.7684 3	0.7737 2	0.7690 3	0.7690 3	0.7823 1
256 2	0.74477	0.7406 3	0.7426 2	0.7322 5	0.7380 4	0.7443 1	0.7402 3	0.7440 2	0.7333 5	0.7384 4	0.7443 1
256 4	0.79403	0.7887 2	0.7826 5	0.7865 4	0.7877 3	0.7931 1	0.7880 3	0.7893 2	0.7842 5	0.7863 4	0.7933 1
256 6	0.80741	0.7944 2	0.7831 5	0.7898 4	0.7929 2	0.8011 1	0.7913 3	0.7926 2	0.7819 5	0.7865 4	0.8025 1
256 8	0.79799	0.7839 2	0.7702 5	0.7750 4	0.7811 3	0.7923 1	0.7751 3	0.7802 2	0.7675 5	0.7721 4	0.7887 1
256 10	0.79147	0.7722 2	0.7624 4	0.7640 4	0.7685 3	0.7826 1	0.7642 3	0.7721 2	0.7592 5	0.7616 4	0.7777 1
512 2	0.75185	0.7435 3	0.7471 2	0.7365 5	0.7400 4	0.7507 1	0.7441 3	0.7499 2	0.7373 5	0.7414 4	0.7509 1
512 4	0.78740	0.7769 2	0.7667 5	0.7741 4	0.7752 3	0.7845 1	0.7772 2	0.7757 2	0.7722 5	0.7750 3	0.7855 1
512 6	0.80260	0.7896 2	0.7676 5	0.7856 4	0.7882 3	0.7939 1	0.7888 2	0.7781 5	0.7800 4	0.7841 3	0.7992 1
512 8	0.79766	0.7845 2	0.7601 5	0.7789 4	0.7827 3	0.7881 1	0.7811 2	0.7700 4	0.7687 4	0.7752 3	0.7927 1
512 10	0.78438	0.7742 2	0.7529 5	0.7671 4	0.7722 3	0.7799 1	0.7702 2	0.7616 3	0.7563 5	0.7631 3	0.7797 1

reach the global optimum despite a large computational effort. The centrality of good solutions has been found to correlate with search difficulty [7]. As a measure of the centrality and reachability of the global optimum, we computed the normalized incoming strength (weighted degree) of the global optimal solution (*strength*). This is computed as the sum of the incoming strengths of the globally optimal nodes divided by the maximum incoming strengths for all nodes.

Figure 1 shows the MLONs metrics for each pivoting rule on the different NK landscapes of size $N = 256$. Except for best improvement, the number of unique nodes and the average fitness of all nodes correlates negatively with the pivoting rules relative efficiency. Despite finding fewer nodes with a good fitness value, best improvement never reaches the best local optimum found on each landscapes for this number of perturbations and achieves poor values for the incoming strength and the size of the best component. For most pivoting rules, the incoming strength and the size of the best component values are low, which indicates that these are difficult problems. However, they are often high for the maximum expansion, which is always the most efficient, and higher than zero on one landscape for the first improvement. Figure 2 shows the MLONs metrics on UBQP landscapes of varied densities. The most efficient pivoting rules find more different local optima than other pivoting rules. The relative efficiency of pivoting rules correlates positively to this number, except between first improvement and the maximum expansion. This difference seems linked to ability of worst improvement of escaping the basins of attraction of the different local optima encountered during the search, which is beneficial to reach the best local optima. The average fitness value is almost similar for each pivoting rule, which is not surprising as the global optimum is often reached. The incoming strength correlates well with the pivoting rule performance, and the size of the best component is fully correlated with the relative efficiency of pivoting rules. A notable observation for this last metric is that best improvement is significantly below the values achieved with other pivoting rules despite the metric normalization.

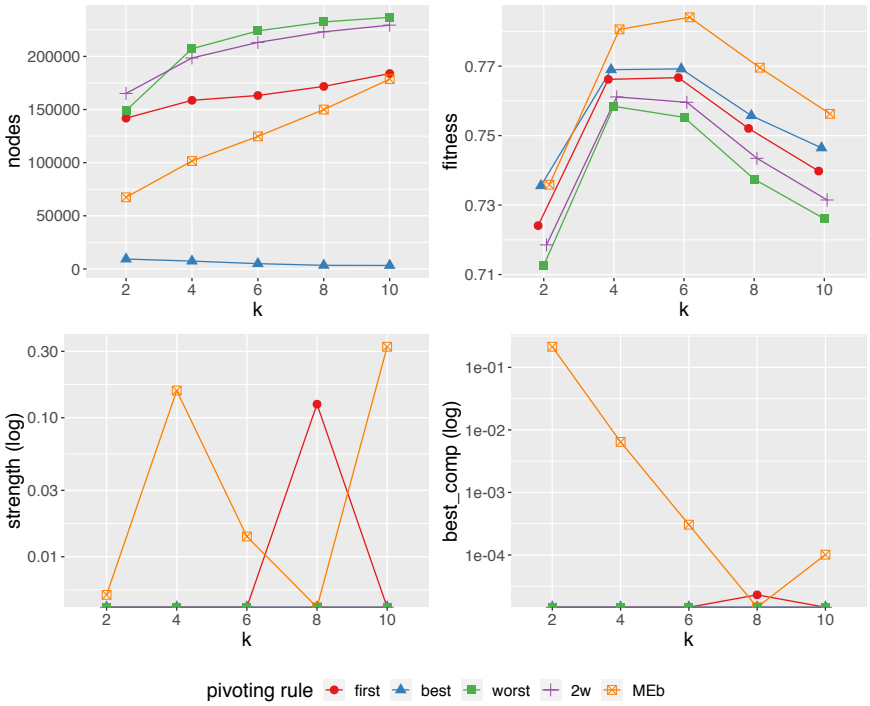

Figure 1: MLON metrics as described in Table 3, for NK landscapes with $N = 256$ and varied K.

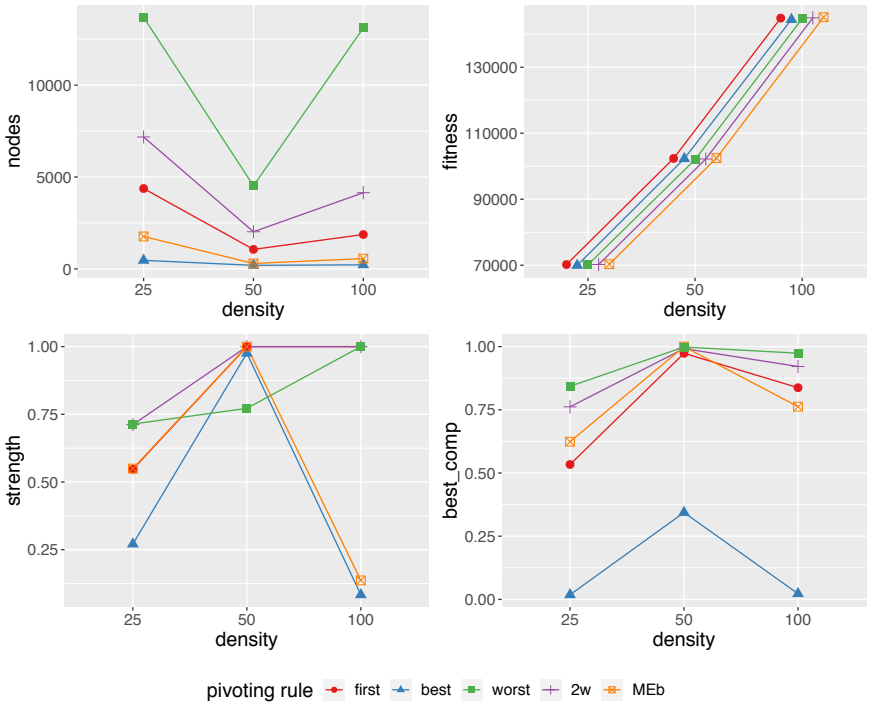

Figure 2: MLON metrics as described in Table 3, for UBQP landscapes with $N = 256$ and varied densities.

Table 2 Global optimum and pivoting rules performance for ILS on each UBQP landscape. For each pivoting rule, the value above corresponds to the number of runs where the global optimum is found; the values below report the number of climbing processes required for this aim.

land.	best LO	$\mathcal{P}=5$					$\mathcal{P}=10$				
		F	B	W	W_2	E	F	B	W	W_2	E
128 25	24087.0	**30**	29	**30**	**30**	**30**	**30**	**30**	**30**	**30**	**30**
		78±44	1519±1060	39±34	43±**30**	29±26	39±31	407±441	45±44	31±27	19±16
128 50	33440.0	**30**	**30**	**30**	**30**	**30**	**30**	**30**	**30**	**30**	30
		16±11	410±549	7±6	8±5	3±3	8±4	66±83	7±4	10±7	3±2
128 100	51130.0	**30**	21	**30**	**30**	**30**	**30**	**30**	**30**	**30**	**30**
		44±33	1427±1386	6±5	13±16	7±7	27±27	472±807	7±6	9±9	4±4
256 25	70861.0	21	3	**30**	**30**	25	**30**	7	**30**	**30**	**30**
		817±1351	1443±1051	133±99	363±569	538±1071	341±414	1155±1118	137±151	134±114	172±209
256 50	102914.0	**30**	14	**30**	**30**	**30**	**30**	18	**30**	**30**	**30**
		126±195	2562±1427	52±43	66±53	65±90	38±49	885±729	112±105	63±46	26±29
256 100	146377.0	23	1	**30**	**30**	27	**30**	4	**30**	**30**	**30**
		1528±1487	2759	58±50	272±280	1264±1261	495±512	1156±486	44±41	94±104	249±286
512 10	134112.0	**30**	5	**30**	**30**	**30**	**30**	10	**30**	**30**	**30**
		694±684	417±580	203±175	271±274	184±227	315±281	2217±1966	121±88	163±162	71±67
512 50	288521.0	**30**	0	**30**	**30**	**30**	**30**	4	**30**	**30**	**30**
		154±209	-	16±16	54±69	47±91	68±110	2174±1189	19±17	24±16	11±13
512 100	424728.0	18	2	**30**	**30**	24	**30**	1	**30**	**30**	**30**
		1084±1611	455±283	25±19	289±371	744±1237	311±358	279	20±16	45±42	128±217

Table 3 Description of network metrics.

nodes	Number of nodes (unique local optima).
fitness	Average fitness of all the nodes.
strength	Incoming strength (weighted degree) of the best local optimum, normalized by the maximum incoming strength.
best-comp	Size (number of nodes) of the component containing the best local optimum, normalized by the total number of nodes.

5.2 Network Visualization

Visualization plays a fundamental role in network analysis, often revealing structural features that are difficult to asses by computing statistical metrics only. To produce network plots we use the R statistics package and the graph layout methods implemented in the igraph library [4]. Our visualization used varied color and sizes to highlight relevant aspects of the LONs.

Figure 3 (top) shows the the network visualization for the NK instance with $N = 256$ and $k = 2$, where the sampling process used a perturbation size $\mathcal{P} = 5$. Since for NK landscapes the number of nodes is too large for visualization purposes (see *nodes* in Fig.1), the networks were pruned to only show local optima with fitness value above the 90 percentile. The plots for the two pivoting rules producing the best and worst performance in this instance, respectively *MEb* and *best*, are shown. The MLON constructed with best improvement does not contain the best local optimum found during the experiments. The 10% of best local optima of the MLON form several small components and are not connected to the best component. The MLON constructed with maximum expansion displays several small components and a large one containing most of the best local optima including the best one.

Figure 3 (bottom) shows the the network visualization for the UBQP instance with $N = 256$ and *density* = 50, where the sampling process used a perturbation size $\mathcal{P} = 5$. The plots for the two pivoting rules producing the best and worst performance in this instance, respectively W_2 and *best*, are shown. The MLON obtained with best improvement displays several small components. The one containing the global optimum is slightly larger and the global optimum is the solution with the highest incoming strength. For worst improvement, the local optima are located in a single funnel. This result illustrates the ability of worst improvement to efficiently navigate through several good local optima, which seems necessary to find the global optimum on most runs.

The MLON models with the most efficient pivoting rules on each problem highlight a *big valley* structure on UBQP landscapes, which would explain why these landscapes are easier to solve for a given same size than NK landscapes. It also corroborates the results of the previous ruggedness study on these landscapes which are locally rugged but globally smooth.

6 CONCLUSIONS

In this paper, we empirically compared classical and alternative pivoting rules within an iterated local search framework on large pseudo-Boolean fitness landscapes. We observed that for the same number of climbing processes, alternative pivoting rules can be efficient. On NK landscapes, the pivoting rules relative efficiency differs from the one observed within single-climbing processes. However, the maximum expansion remains the most efficient to reach good-quality solutions. On UBQP landscapes, the worst improvement and its approximation are still the most efficient pivoting rules within the iterated local search. The analysis of the MLONs revealed that the efficiency of the pivoting rules correlates well

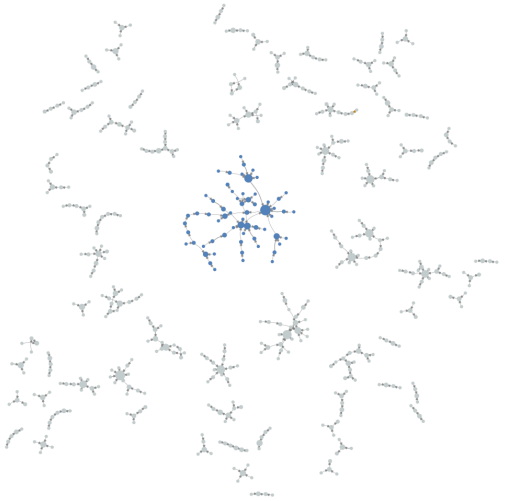

(a) NK, $K = 2$, pivot = best

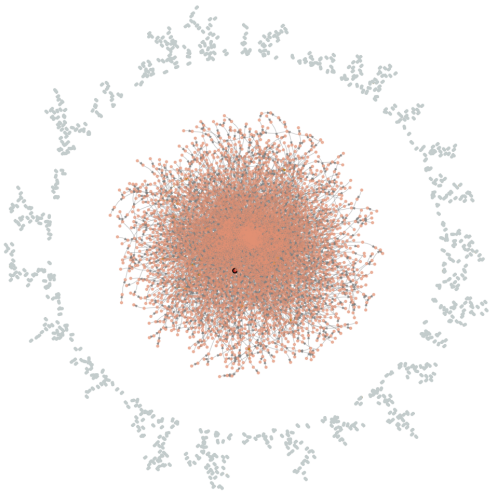

(b) NK, $K = 2$, pivot = MEb

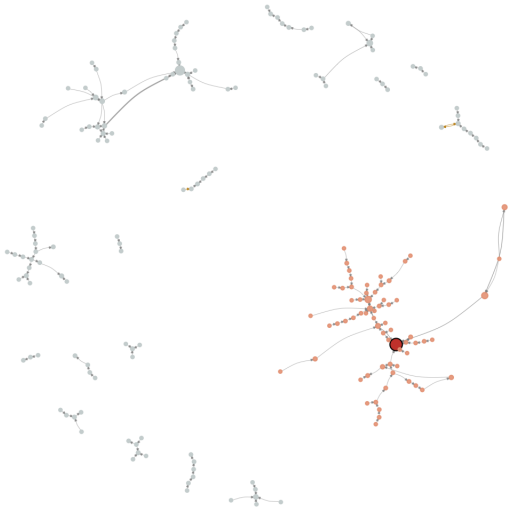

(c) UBQP, density = 50, pivot = best

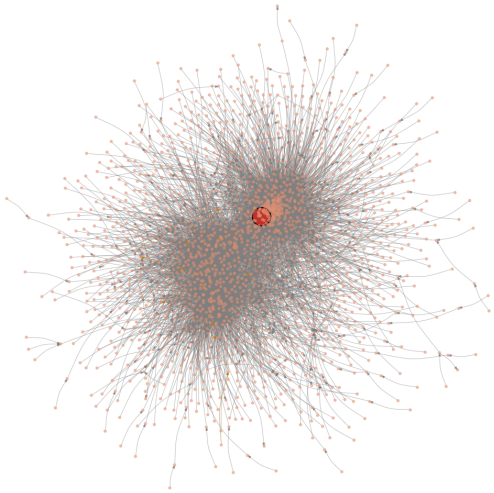

(d) UBQP, density = 50, pivot = W_2

Figure 3: MLONs for representative NK and UBQP instances with $N = 256$. The pivoting rule is shown in the sub-captions. Since the networks are too large for NK landscapes, they were pruned to show only the nodes with fitness above the 90 percentile. The best-found optimum is indicated in red, while the local optima belonging the same connected component of the best are visualized in pink. Blue nodes represent the largest connected sub-optimal connected component in the network. In the NK landscape, the *best* pivoting rule failed to locate the best solution. The reminder nodes are visualized in light gray. The size of nodes is proportional to their incoming weighted degree (strength).

with the number of nodes, the average fitness value of the network, the strength, and the size of the component containing the best local optimum. The MLON models on UBQP landscapes reveal a big valley structure, which makes them easier to solve.

Similar experiments should be conducted on several iterated local search frameworks to further understand the behavior induced by classic and alternative pivoting rules, using different kinds of perturbations and acceptance criteria. Ideally, this should be done according to different stopping criteria, including several budgets of evaluations and of climbing processes. Conducting a similar study on permutation-based problems, such as the flow-shop scheduling problem or the quadratic assignment problem, would help to improve the understanding of classic and alternative pivoting rules.

REFERENCES

[1] Ciro Amitrano, Luca Peliti, and Mohammed Saber. 1989. Population dynamics in a spin-glass model of chemical evolution. *Journal of Molecular Evolution* 29, 6 (1989), 513–525.

[2] Matthieu Basseur and Adrien Goëffon. 2013. Hill-climbing strategies on various landscapes: an empirical comparison. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 479–486.

[3] Matthieu Basseur and Adrien Goëffon. 2014. On the efficiency of worst improvement for climbing NK-landscapes. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 413–420. https://doi.org/10.1145/2576768.2598268

[4] G. Csardi and T. Nepusz. 2006. The igraph software package for complex network research. *InterJournal* Complex Systems (2006), 1695.

[5] Michael R. Garey. 1979. A Guide to the Theory of NP-Completeness. *Computers and intractability* (1979).

[6] Pierre Hansen and Nenad Mladenović. 2006. First vs. best improvement: An empirical study. *Discrete Applied Mathematics* 154, 5 (2006), 802–817.

[7] Sebastian Herrmann, Gabriela Ochoa, and Franz Rothlauf. 2017. PageRank centrality for performance prediction: the impact of the local optima network model. *Journal of Heuristics* (12 May 2017).

[8] Holger H. Hoos and Thomas Stützle. 2004. *Stochastic local search: Foundations and applications*. Elsevier.

[9] Terry Jones et al. 1995. *Evolutionary algorithms, fitness landscapes and search*. Ph.D. Dissertation.

[10] Stuart A. Kauffman. 1992. The Origins of Order: Self-Organization and Selection in Evolution. In *Spin Glasses and Biology*. Series on Directions in Condensed Matter Physics, Vol. Volume 6. WORLD SCIENTIFIC, 61–100. https://doi.org/10.1142/9789814415743_0003

[11] Stuart A. Kauffman and Edward D. Weinberger. 1989. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* 141, 2 (Nov. 1989), 211–245. https://doi.org/10.1016/S0022-5193(89)80019-0

[12] Michael R. Lissack. 1999. Complexity: The science, its vocabulary, and its relation to organizations. *Emergence* 1, 1 (1999), 110–126.

[13] Helena R Lourenço, Olivier C Martin, and Thomas Stützle. 2003. Iterated local search. In *Handbook of metaheuristics*. Springer, 320–353.

[14] Katherine Mary Malan. 2021. A Survey of Advances in Landscape Analysis for Optimisation. *Algorithms* 14, 2 (2021), 40.

[15] Katherine M. Malan and Andries P. Engelbrecht. 2013. A survey of techniques for characterising fitness landscapes and some possible ways forward. *Information Sciences* 241 (Aug. 2013). https://doi.org/10.1016/j.ins.2013.04.015

[16] G. Ochoa, M. Tomassini, S. Verel, and C. Darabos. 2008. A study of NK landscapes' basins and local optima networks. In *Genetic and Evolutionary Computation Conference - GECCO 2008*. ACM, 555–562. https://doi.org/10.1145/1389095.1389204

[17] Gabriela Ochoa and Nadarajen Veerapen. 2016. Deconstructing the Big Valley Search Space Hypothesis. In *Evolutionary Computation in Combinatorial Optimization - 16th European Conference, EvoCOP 2016, Porto, Portugal, March 30 - April 1, 2016, Proceedings (Lecture Notes in Computer Science)*, Vol. 9595. Springer, 58–73.

[18] Gabriela Ochoa, Nadarajen Veerapen, Fabio Daolio, and Marco Tomassini. 2017. Understanding Phase Transitions with Local Optima Networks: Number Partitioning as a Case Study. In *Evolutionary Computation in Combinatorial Optimization, (EVOCOP) (LNCS)*, Vol. 10197. Springer, 233–248.

[19] Gabriela Ochoa, Sébastien Verel, and Marco Tomassini. 2010. First-Improvement vs. Best-Improvement Local Optima Networks of NK Landscapes. In *Parallel Problem Solving from Nature, PPSN XI (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 104–113. https://doi.org/10.1007/978-3-642-15844-5_11

[20] Gintaras Palubeckis. 2004. Multistart tabu search strategies for the unconstrained binary quadratic optimization problem. *Annals of Operations Research* 131, 1-4 (2004), 259–282.

[21] Erik Pitzer and Michael Affenzeller. 2012. A Comprehensive Survey on Fitness Landscape Analysis. *Recent Advances in Intelligent Engineering Systems* 378 (2012), 161–191.

[22] Peter F. Stadler and Günter P. Wagner. 1997. Algebraic theory of recombination spaces. *Evolutionary computation* 5, 3 (1997), 241–275.

[23] Sara Tari, Matthieu Basseur, and Adrien Goëffon. 2016. Toward the Design of Efficient Pivoting Rules for Local Search. In *Genetic and Evolutionary Computation Conference - GECCO 2016*. ACM, 55–56.

[24] Sara Tari, Matthieu Basseur, and Adrien Goëffon. 2018. Worst Improvement Based Iterated Local Search. In *European Conference on Evolutionary Computation in Combinatorial Optimization*. Springer, 50–66.

[25] Edward Weinberger. 1990. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological cybernetics* 63, 5 (1990), 325–336.

[26] Darrell Whitley, Adele Howe, and Doug Hains. 2013. Greedy or not? best improving versus first improving stochastic local search for MAXSAT. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 27.

[27] L. Darrell Whitley, Andrew M. Sutton, and Adele E. Howe. 2008. Understanding Elementary Landscapes *(Genetic and Evolutionary Computation Conference - GECCO 2008)*. ACM, New York, NY, USA, 585–592. https://doi.org/10.1145/1389095.1389208

[28] Sewall Wright. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the sixth international congress of genetics*, Vol. 1. 356–366.