



HAL
open science

Expansion-based Hill-climbing

Sara Tari, Matthieu Basseur, Adrien Goëffon

► **To cite this version:**

Sara Tari, Matthieu Basseur, Adrien Goëffon. Expansion-based Hill-climbing. Information Sciences, 2023, 649, pp.119635. <10.1016/j.ins.2023.119635>. <hal-04393827>

HAL Id: hal-04393827

<https://hal.science/hal-04393827v1>

Submitted on 1 Oct 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC 4.0 - Attribution - Non-commercial use - International License

Expansion-based Hill-climbing

Sara Tari^{a,*}, Matthieu Basseur^a, Adrien Goëffon^b

^a*Univ. Littoral Cote d'Opale, 50 rue Ferdinand Buisson, Calais, F-62228, France*

^b*Univ. Angers, 2 boulevard Lavoisier, Angers, F-49045, France*

Abstract

This paper investigates the influence of adaptive walks heuristics within local searches, by studying to what extent a wiser choice among improving neighbors influences the expected quality of the attained local optima. To this aim, we specifically focus on hill-climbers and introduce the maximum expansion pivoting rule, which selects the improving neighbor having the highest number of improving neighbors. Experiments show that having one step ahead of information allows for wiser neighbor choices, leading to better local optima. As the best improvement climber, a maximum expansion climber selects a particular search trajectory among all possible first improvement trajectories, however, it significantly increases the expected quality of the trajectory. On the other hand, the computational overhead makes this heuristic less valuable when included in an iterated search process, where repeating fast random hill-climbings from random starting points still allows a better exploration of the search space. This paper, therefore, extends previous studies on the relative efficiency of hill-climbing pivoting rules, by focusing on the original maximum expansion selection criterion.

Results suggest that the achievement of good local optima combined with the use of approximation techniques and problem specificities could lead to the design of effective advanced metaheuristics that exploit the maximum expansion principle.

Keywords: maximum expansion, hill-climbing algorithm, local search,

*Corresponding author

Email addresses: sara.tari@univ-littoral.fr (Sara Tari),
matthieu.basseur@univ-littoral.fr (Matthieu Basseur),
adrien.goeyffon@univ-angers.fr (Adrien Goëffon)

Preprint submitted to Information Sciences

August 30, 2023

1. Introduction

Within the field of metaheuristics, the quantity of so-called original methods drastically grows since the 90s, especially regarding metaheuristic hybridizations, applications, and definitions of alternative nature-inspired evolutionary or behavioral metaphors. One can sometimes wonder about the actual contribution of some papers to the community, in particular when the main motivation is to extend the bestiary of optimization algorithms. Such metaheuristics are usually classic local search or evolution strategies, that are often shallowly compared to a small subset of other algorithms without necessarily identifying the mechanisms that impact the sampling of solutions. Today, researchers regularly comment on this issue [36, 39, 15, 12, 38, 47] and grant particular attention to more pragmatic considerations around the scalability of the algorithms, the adaptive parameter tuning, the understanding of specific search processes and their relevance for tackling problem instances with identified specificities.

Other papers describe a manner to apply or lightly adapt a metaheuristic for beating the best-known solutions of a particular recently-introduced problem, with a reduced analysis of the contribution brought by each mechanism. However, determining the relevance of the components of metaheuristics is essential to improve our understanding of their global behavior and to predict their adaptability to different kinds of problems [48].

More precisely, the scope of this work concerns neighborhood-based search strategies and the dynamics behind exploitation mechanisms in local search algorithms. Indeed, designing effective metaheuristics requires a good understanding of basic operators and their combination for a particular search or a particular class of problem. Specific experimental studies can provide general observations, which can be facilitated by abstracting combinatorial optimization problem instances into fitness landscapes [49].

Naturally, highly multidimensional and combinatorial fitness landscapes are drastically different from continuous bi- or tri-dimensional surfaces that usually serve as a visual representation to explain local search and evolutionary concepts (as well as exploration/exploitation issues) but are sometimes misused as a mental visualization to guide the design of metaheuristics.

Wrong intuitions arise from such figures, where the idea of being trapped in a local area is over-estimated. In general, it is inexact that an intensifica-

tion local search process can only explore a reduced area around an arbitrary initial solution. There are often many ways to drive the search toward good local optima (including global optima) just by using a hill-climbing process [9]. However, the primary concern while using metaheuristics is usually to pay attention to the diversification mechanism(s).

Contribution

The general aim of this paper is to study and identify the principles which enable drive local searches more likely toward better solutions. We solely focus on hill-climbing processes, often denoted by *climber* algorithms (*descent*, for minimization problems), which correspond to basic iterative improvement rules. Although climber-like intensification operators are a central component of metaheuristics (including modern and sophisticated ones), their conceptual choices and appropriateness are rarely discussed (apart from neighborhood considerations) in comparison to crossover or perturbation operators, and parameter tuning issues.

Given a search space equipped with a neighborhood structure and a fitness function, the search trajectory of a local search algorithm, and in particular of a climber, is defined by the neighbor selection strategy, also called *pivoting rule* [50]. We investigate how the pivoting rule affects the ability of the search to reach better local optima. In previous studies [6, 7], we highlighted that empirical observations often contradict a dominant conception choice while solving combinatorial problems with metaheuristics, where the local search strategy generally favors best neighbors. Here, we are looking for a more precise and general selection criterion.

The question that drives this work is: starting from any point (solution) of a combinatorial fitness landscape, and by using only local information (fitness values of close solutions), how to move iteratively in a way that describes an ascending trajectory such that the expected fitness value of the last point (local optimum) is maximum? In other words, finding a hill-climbing strategy that maximizes the expected quality of the optimum reached. Let us specify that here we consider a strict definition of climbers, with no neutral moves, perturbations (decreasing moves), or jumps (non-neighboring moves). Thus, this definition does not include more sophisticated climbing-based metaheuristics like iterated local search [26], variable neighborhood search [31], β -hill climbing [3, 4] or late-acceptance hill-climbers [11].

To investigate the underlying question of this work we introduce an alternative pivoting rule, called *Maximum Expansion* (ME) that guides the tra-

jectory in order to maximize future possibilities of improvement. At each iteration, the ME climber selects the improving neighbor which has the largest number of improving neighbors, therefore reducing the risk of being trapped early in a local optimum. As highlighted in the experiments, ME is particularly efficient for reaching good local optima. This result provides interesting information on the properties and behavior of a landscape-aware local search. However, ME requires the evaluation of a larger neighborhood and is therefore computationally more expensive than basic climbing strategies.

Although this study primarily aims at identifying more intelligent search space trajectory rules, and not (yet) to present more efficient local search operators to include in metaheuristics, additional experiments focus on finding ways to use the ME principle within faster climbers in order to achieve tradeoffs between the expected quality of local optima and computational effort/quantity of information used.

The paper provides experimental comparisons of ME and basic hill-climbing pivoting rules *First* (ie. random), *Best*, and *Worst improvement*, as well as other natural variants. The experiments have been conducted on pseudo-boolean fitness functions, from NK and UBQP (Unconstrained Binary Quadratic Problem) problem instances, and permutation-based fitness functions, from Flow-shop and QAP (Quadratic Assignment Problem) instances. The empirical analysis includes statistical comparisons of the quality of local optima reached, and the fitness values attained given a fixed budget.

This paper is organized as follows. The next section recalls fundamental concepts related to fitness landscapes and climbers. It also introduces new notations and perspectives specifically adapted to the particular issues of examining search space trajectories. In section 3, we recall previous results related to climber guidance strategies and present the maximum expansion pivoting rule. In section 4, experiments on various combinatorial fitness landscapes assess the ability of this search strategy to reach good-quality local optima. As high computational budgets are required, we provide in section 5 a strategy that leads to obtaining interesting tradeoffs between the number of fitness evaluations needed and the quality of local optima obtained. This strategy is compared to other ones for a similar budget in order to highlight and discuss its limitations. In the last section, we discuss this work and some possible ways forward.

2. Combinatorial optimization, fitness landscapes

Solving a combinatorial optimization problem consists of determining, among a discrete set of potential solutions, the optimal solution with respect to an objective function that assigns an objective value to each solution. Search spaces are induced by instances of processed problems and are generally constituted with an exponential number of solutions w.r.t. the size of instance representation.

Most of the widely studied combinatorial optimization problems are NP-hard, and polynomial iterative approaches can be used to solve them in an approximated way, by exploring a subset of solutions of the search space. This subset leads to a solution determined by the strategy under consideration, with an objective value as close as possible to the optimal objective value. A usual technique is to define the subset by navigating through the search space, w.r.t. a neighborhood relation described by a function that associates a set of neighboring solutions to each solution of the search space. The neighborhood relation generally induces a reasonable size, to limit the number of solutions to evaluate.

Although no restriction is required while defining a neighborhood relation, it is intuitive to connect solutions that share most of their structural properties. In most cases, the distance between two solutions should depend upon their differences; the closer a solution is to another in terms of neighborhood distance, the more they should share similar structures. When solving a combinatorial optimization problem using local search-based algorithms, we assume that there is a significant level of correlation between the fitness values of solutions and their structural properties. Otherwise, this would make any search space exploration by a neighborhood-based strategy pointless.

2.1. Local search and Hill-climbing

A combinatorial problem instance is determined by a discrete set of solutions \mathcal{X} called *search space*, and an objective function $f : \mathcal{X} \rightarrow \mathbb{K} \subseteq \mathbb{R}$. Solving a combinatorial optimization problem (\mathcal{X}, f) in a maximization context consists of determining $\operatorname{argmax}_{x \in \mathcal{X}} f(x)$, which is the set of solutions x^* such that $\forall x \in \mathcal{X}, f(x) \leq f(x^*)$. In the following, we consider by default maximization problems, but all definitions and algorithms can naturally be extended to minimization problems without loss of generality.

A local search (LS) algorithm explores the search space by navigating through a set of solutions w.r.t. a neighborhood relation defined by a function $\mathcal{N} : \mathcal{X} \rightarrow 2^{\mathcal{X}}$, and an evaluation function $f : \mathcal{X} \rightarrow \mathbb{K}$. The neighborhood relation associates a subset of solutions $\mathcal{N}(x) \subseteq \mathcal{X}$ to each solution $x \in \mathcal{X}$. \mathcal{N} is commonly described by the possible outcomes of a particular transformation which links each solution to a set of neighbors with a similar structure. The evaluation function associates a scalar value to each solution x and implies a preference relation between couples of solutions.

Given a search space \mathcal{X} with a transition relation $\mathcal{T} \subseteq \mathcal{X}^2$ between solutions, we call *topological transition graph* the graph $(\mathcal{X}, \mathcal{T})$. If $\mathcal{T} = \{(x, y), y \in \mathcal{N}(x)\}$, then the topological transition graph $(\mathcal{X}, \mathcal{T})$ defined by the neighborhood function \mathcal{N} describes the possible trajectories of a local search within the search space. Given a topological transition graph, an LS algorithm is determined by an initial solution, a move policy, and a stopping criterion. A *move* $x \rightarrow y$ is an edge $(x, y) \in \mathcal{T}$ selected by the move policy. A sequence of s moves $x_0 \rightarrow x_1, x_1 \rightarrow x_2, \dots, x_{s-1} \rightarrow x_s$ gives an LS trajectory $(x_0, x_1, x_2, \dots, x_s)$, which is a path in the graph. Note that an LS process can be viewed as an intelligent way to sample the search space. A local search trajectory is determined by its pivoting rule (deterministic or stochastic) which chooses neighboring solutions depending on additional information, in most cases their evaluation $f(\cdot)$. The relative efficiency between local search algorithms is then based on their ability to generate trajectories in a search space that include a solution as best as possible (according to the objective function f). Algorithm 1 represents a local search process with such a view.

Algorithm 1 Local search algorithm

```

1: Choose  $x \in \mathcal{X}$ 
2:  $s \leftarrow 0$ 
3:  $x_s \leftarrow x$ 
4: while stop criterion not reached do
5:   Choose  $x \in \mathcal{N}(x_s)$ 
6:    $s \leftarrow s + 1$ 
7:    $x_s \leftarrow x$ 
8: end while
9: Choose  $x \in \{x_0, \dots, x_s\}$ 
10: return  $x$ 

```

A hill-climbing algorithm (*climber*) is a local search that only performs non deteriorating moves according to an evaluation function, ie. $\forall i < s : f(x_{i+1}) > f(x_i)$ (*strict climber*), or $f(x_{i+1}) \geq f(x_i)$ (*stochastic climber*, with side-walk mechanism). A strict climber determines a trajectory in the graph $(\mathcal{X}, \mathcal{T}_>)$, with $\mathcal{T}_> = \{(x, y), y \in \mathcal{N}(x), f(y) > f(x)\}$.

From an initial solution, a *first improvement climber* randomly walks in $(\mathcal{X}, \mathcal{T}_>)$ until no move is possible. Indeed a *local optimum* is a vertex in $(\mathcal{X}, \mathcal{T}_>)$ with no outgoing edge. Other hill-climbing strategies like *best improvement* or *worst improvement* choose neighbors according to a preference relation \succ , respectively $y_1 \succ y_2 \leftrightarrow f(y_1) > f(y_2)$ and $y_1 \succ y_2 \leftrightarrow f(y_1) < f(y_2)$, where y_1 and y_2 are two improving neighbors of a given solution x . In the following, for simplicities sake, we consider only symmetric neighborhood relations and thus bidirected (or undirected, depending on the context) topological transition graphs.

Famous LS metaheuristics like simulated annealing or tabu search use adaptive information to influence the neighbor choices and thus the trajectories. Such information includes the trajectory length so far, computational effort or visited solutions memory so that the move strategies are not only determined by *objective* information. Although these advanced LS techniques naturally outperform basic searches like climbers, their behavior is based on explicit diversification aiming to compensate for the lack of efficiency of guiding strategies.

To study the practical efficiency of neighborhood-based search principles, we focus on *non-contextual* neighborhood selection strategies, that is pivoting rules defined such that the probability of a move $x \rightarrow y$ is not affected by any learning process or search information and only depends on pre-calculable, external information. In this way, we propose to extend search spaces to combinatorial fitness landscapes to abstract from the problem of determining efficient non-contextual neighborhood selection strategies.

2.2. Fitness landscapes

The fundamental concept of fitness landscapes takes its source in theoretical biology when Wright provided a visual representation of existing links between genotype and chances of reproductive success of individuals [49]. One can use such an abstract representation of evolutionary mechanisms within a genotypic space described by a hypercube whose vertices represent solutions of a pseudo-boolean problem and edges represent potential mutations. A fitness landscape can then be defined by associating fitness values




Figure 1: Hypercube of dimension 5 representing the landscape of a bit-string problem with 5 decision variables. Numbers correspond to the decimal values associated to each bitstring representing a solution. Solutions in the dark grey area correspond to the neighbors of 0 (distance from 0 equals to 1). The grey scale denotes the distance of solutions from 0.

to the hypercube vertices (Figure 1).

From this base, an analogy in soft computing was made between neighborhood searches and evolutionary genetics. A search space corresponds to a set of acceptable genotypes, a neighborhood relation describes possible mutations, and the objective value or evaluation of a solution refers to the quality of the phenotype of an individual (adaptation to the environment, chances of reproductive success, etc.). The triplet (search space, neighborhood relation, evaluation function) can then be abstracted by a triplet (genotype, mutation, phenotype/fitness).

More formally, a *fitness landscape* \mathcal{L} can then be defined as a triplet $(\mathcal{X}, \mathcal{N}, \mathbf{f})$ where \mathcal{X} , \mathcal{N} and \mathbf{f} respectively correspond to the aforementioned search space, neighborhood function, and evaluation function (denoted here as *fitness function*). A local optimum of a landscape $(\mathcal{X}, \mathcal{N}, \mathbf{f})$ is a solution




Figure 2: Depiction of a completely smooth landscape (left-hand side) and a particularly rugged landscape in which fitness are randomly distributed (right-hand side).

$x \in \mathcal{X}$ such that $\forall y \in \mathcal{N}(x), f(y) \leq f(x)$. A global optimum is a solution $x \in \mathcal{X}$ such that $\forall y \in \mathcal{X}, f(y) \leq f(x)$. From a landscape $\mathcal{L} = (\mathcal{X}, \mathcal{N}, f)$, we note $\mathcal{T}(\mathcal{L})$ the topological transition graph $(\mathcal{X}, \{(x, y), y \in \mathcal{N}(x)\})$, and $\mathcal{HC}(\mathcal{L})$ the (strict) hill-climbing transition graph, ie. the graph $(\mathcal{X}, \mathcal{T}_>)$ such that $\mathcal{T}_> = \{(x, y), y \in \mathcal{N}(x), f(y) > f(x)\}$.

Considering solutions placed depending on their relative genotypic distance and following the distribution of fitness values, landscapes are described as more or less *rugged*. A rugged landscape has more peaks, which are points of landscapes from which it is impossible to improve individuals with a single mutation (i.e. corresponding to the local optima). Figure 2 depicts these concepts and how they would translate in a continuous space. Note that this representation is deceptive when it comes to combinatorial spaces but helps to convey the idea of the climbing difficulty induced by these characteristics.

From an evolutionary point of view, we can easily conceptualize what makes a combinatorial landscape smooth or rugged. If we consider a hypercube representation and assume that each gene represents a bit, there are only two possible states, modifiable by a genetic mutation. If the state of a given gene does not influence the expression of other genes, then any mutation of this gene will have the same effect regardless of the genetic heritage of the individual. For instance, if mutating the first gene from 0 to 1 gives a fitness variation of one unit regardless of the rest of the genotype of individuals, then any evolutionary pressure will naturally set this gene to 1. This gene would not add any ruggedness to the landscape. Indeed, at any place in the landscape, an identical mutation of this gene will always have the same effect, and the improvement potential of the remaining bits will not change before or after this particular mutation. If every gene is independent,

then the fitness landscape has no ruggedness and contains only a single local optimum. In such a landscape, from any initial individual, any ordering of beneficial mutation application will lead to the individual with the highest fitness value. However, in most cases, genes are in interaction with some other genes and the mutation effect of a gene is under the influence of the value of other genes. The ruggedness of a landscape directly results from this phenomenon, called *epistasis* [10]. The more there exist interactions between genes the less the fitness variation resulting from a gene mutation is directly associated with the value of this gene.

The combinatorial fitness landscape model of [22, 21] is an appropriate tool to apprehend the solving of combinatorial optimization problems using neighborhood search processes, by studying search landscapes [18, 28, 27]. Once defined, such landscapes can be studied independently of the problem, the solutions representation, the neighborhood relation, or any objective or evaluation function.

Using fitness landscape abstraction to describe the behavior of neighborhood-based and evolutionary techniques can lead to a better understanding of the capacity of a strategy to handle specific problem instances. Working on a more rugged landscape naturally makes the search and the discovery of good solutions more difficult. However, a highly rugged landscape comes from a low correlation between genotype and fitness. In this case, one can often observe large fitness differences between two neighboring genotypes. It can result from the hardness of the instance, in which it is difficult to find a strong link between the solution structure and its fitness value. The use of an inadequate neighborhood relation also leads to this issue. The neighborhood structure is a key aspect of designing efficient search algorithms but is not discussed in this paper.

2.3. Fitness Landscapes features

Fitness landscapes can be characterized by many features describing their dimension, neutrality, and ruggedness. The size and connectivity of a fitness landscape primarily determine its dimension, which reflects its topological structure. The size of a landscape is the cardinality of \mathcal{X} and naturally contributes to the difficulty to reach a given point, in particular, the global optimum, using a sampling strategy. The connectivity of a landscape depends upon \mathcal{N} and affects the distance between solutions and the trajectory possibilities. High connectivity induces small distances between solutions but leads search strategies to face a huge number of possible moves. Low

connectivity induces high distances between solutions and a reduced number of possible moves for search strategies. Classic graph theory metrics and properties are among possible indicators which can help to characterize the connectivity of a landscape.

Neutrality appears in fitness landscapes if there is a significant proportion of neighboring solutions with the same fitness value. This phenomenon can strongly influence the behavior of search algorithms [37, 46]. However, we do not consider neutrality in this study on landscapes and strategies, which helps us to simplify the analysis of the results and identify features that influence search strategies behavior more clearly.

Ruggedness appears when the effect of one mutation is conditioned by other solution information and is a direct consequence of the epistasis phenomenon discussed previously. The ruggedness level of landscapes is a primary characteristic that influences their difficulty to be explored efficiently. Establishing a precise and formal definition of such a concept, or an accurate indicator, is not apparent. In many works, the ruggedness of a landscape is estimated with an autocorrelation measure [30, 5, 24]. [29] distinguishes the level of variable interdependency, hence the epistasis, and the ruggedness, only associated with the number and distribution of local optima. However, we think that it is relevant to directly link the ruggedness indicator of a landscape to the number of epistatic interactions. In a hill-climber study context, sign epistasis is a particularly relevant feature. Sign epistasis occurs when the presence of a given gene changes the beneficial or deteriorating nature of a given mutation. High rates of sign epistasis induce a high number of local optima and increase the probability of quickly meeting a local optimum during the search process. Since in a hill-climbing context the selection of a local optimum stops the search process, high rates of sign epistasis increase the difficulty to reach good local optima. Indeed, a significant number of local optima necessarily involves a high level of variable interdependency and an accurate measure of the difficulty to climb landscapes. Moreover, the ruggedness of fitness landscapes as a concept is originally presented as the consequence of epistasis.

In this paper, we use the *1-ruggedness* indicator [8] described hereafter (see also Fig. 3). Let us consider four distinct solutions x , x' , y , and y' such that $\{x', y\} \subseteq \mathcal{N}(x) \cap \mathcal{N}(y')$ and $y' \notin \mathcal{N}(x)$. Thus $\{x, x', y', y\}$ forms a (non-clique) square in the graph $\mathcal{T}(\mathcal{L})$. $\{x, x', y', y\}$ describes a sign epistasis iff $(f(x') > f(x)) \neq (f(y') > f(y))$, ie. $(f(x') - f(x)) \times (f(y') - f(y)) < 0$. The 1-ruggedness of a landscape is the proportion of squares in $\mathcal{T}(\mathcal{L})$ which




Figure 3: Epistasis occurs when the presence of a mutation (\uparrow) affects the effect of another mutation (\rightarrow). We observe sign epistasis once a deteriorating mutation becomes beneficial when occurring after another mutation (right-hand side). Left-hand side: no sign epistasis.

involve sign epistasis. Since in general, it is not conceivable to determine this proportion accurately, we estimate this 1-ruggedness indicator by sampling the search space. Note that this definition is restricted to undirected fitness landscapes ($x \in \mathcal{N}(y) \leftrightarrow y \in \mathcal{N}(x)$) but can be extended to non-symmetric neighborhood functions. In this paper, we focus on a global characterization of various landscapes.

Another possible way to analyze the topology of fitness landscapes consists of studying their local optima. These studies include works on basins of attraction [45], local optima networks [32], and saddle points [40]. The basin of attraction of a local optimum x_{opt} is the set of solutions such that a given climber starting from each of these solutions return x_{opt} . The study of these basins and their properties can bring useful information for optimizers (size, localization). Studying local optima networks consists of constructing graphs in which vertices are local optima and edges are possible transitions between them. Saddle points are solutions in the search space that are identified as key solutions when a simple move can determine if the search will attain either local optimum.

2.4. Navigating Through Fitness Landscapes

When the topological transition graph of a fitness landscape is strongly connected, a random walk can potentially reach any solution of the search space from any starting solution in most cases. This global attainability is verified with most classical solution representations and neighborhood operators. For instance, considering an N -dimensional hypercube fitness landscape \mathcal{L} , then $\mathcal{T}(\mathcal{L})$ is strongly connected and has $N \times 2^{N-1}$ bidirectional transitions. A first improvement climber on the landscape \mathcal{L} is equivalent to a

random walk on $\mathcal{HC}(\mathcal{L})$. In general, $\mathcal{HC}(\mathcal{L})$ loses all bidirectional transitions of the associated topological transition graph $\mathcal{T}(\mathcal{L})$. However, when considering a non-neutral landscape with a symmetrical neighborhood, its associated topological transition graph contains exactly 50% of improving transitions. In particular, regardless of the defined non-neutral fitness function f , $\mathcal{T}(\mathcal{L})$ and $\mathcal{HC}(\mathcal{L})$ contains respectively $N \times 2^N$ and $N \times 2^{N-1}$ directed edges ($\mathcal{HC}(\mathcal{L})$ is then weakly connected). Climbers are often viewed as intensification mechanisms that are only able to attain a tiny part of the search space. This vision is inaccurate since hill-climbing transition graphs are half-dense and potentially still have a high level of connectivity.

In [9], we introduced the *coverage rate* of a particular solution x , as the proportion of the search space from which x can be reached by a hill-climbing process. Experiments on enumerable NK landscapes show that the best solutions (including the global optimum) are generally reachable from a significant proportion of starting points. This study leads us to think there exist ways to improve the efficiency of local searches, especially of climbers. For instance, the definition of alternative pivoting rules could help to increase the probability to reach better local optima.

3. Maximum Expansion: a climber driven by the expansion score

3.1. Motivations

As stated previously, a pivoting rule is a mechanism that directly influences the search trajectory of climbers. Among basic pivoting rules, first and best improvement are widely used within metaheuristics. Recall that the (random) first improvement pivoting rule consists of choosing the first improving neighbor encountered at each step of the search (neighbors being considered in random order), whereas the best improvement pivoting rule consists of selecting the improving neighbor with the best fitness value at each step of the search.

One of the advantages of first improvement compared to best improvement is it requires fewer evaluations as it avoids the systematic evaluation of the whole neighborhood. Note that when using an incremental evaluation is possible, the difference in computational cost between the two strategies is less significant. Previous studies [33, 6] investigated the effect of these two classical pivoting rules on the ability of climbers to reach good quality local optima w.r.t. some landscape properties. In particular, experiments on a panel of large fitness landscapes highlighted that first improvement led

toward better quality local optima than best improvement, except on very smooth landscapes [6]. The superiority of first improvement over best improvement in most landscapes is not a direct consequence of the stochasticity of the mechanism. Most best-improving steps tend to lead the search toward close local optima which can be avoided by a more stochastic process. It was also shown that when first improvement outperforms best improvement, then worst improvement (selecting the neighbor which brings the smallest fitness improvement) outperforms first improvement [7], and vice versa. This observation indicates a correlation between the expected quality of reached local optima and the relative quality of neighbors selected by the climbing process. Combined experimental studies [33, 6, 7] indicate that on homogeneous randomly generated landscapes, this correlation is positive on highly smooth and highly rugged landscapes, and negative on mildly rugged landscapes. It makes sense that on smooth landscapes, the small number of local optima decreases the necessity to bypass them. Moreover, with a reduced amount of epistasis, it is more likely that most steep slopes lead toward better local optima. At each step, the probability of reaching a local optimum is thus significantly reduced when performing larger steps compared to more rugged landscapes, allowing a sufficient convergence for the best improvement.

On landscapes with *very* high ruggedness levels, with no apparent correlation between neighborhood distance and fitness variation, the probability of reaching higher local optima is increased when the fitness of the solution to be selected is maximized at each step. Such landscapes have a considerable proportion of local optima, leading to a faster convergence of hill climbers, and eventually to favor best improvement moves [33]. However, they do not reflect the properties of typical search landscapes tackled in academic or real-life contexts. Such landscapes are not considered in this work since their lack of structure makes any local search process inefficient.

Although there is some evidence or intuitions to explain the positive/negative correlations between the quality of a hill-climbing (regarding the expected fitness of the reached local optima) and its fitness-based pivoting rule, we present in the following complementary studies aiming to provide a better understanding of what is driving a climber efficient.

First of all, let us recall that in theory, hill-climbing pivoting rules of maximal efficiency exist. Given $\mathcal{HC}(\mathcal{L})$, the optimal hill-climbing process consists of following the smallest path leading to the highest solution reachable from the current solution. More generally, a similar local search process can be defined from $\mathcal{T}(\mathcal{L})$ without the improving moves constraint (leading to an exact

algorithm if $\mathcal{T}(\mathcal{L})$ is strongly connected). However, in most cases, such a rule cannot be described explicitly and for all landscapes in polynomial time or space (unless $P=NP$). In the same way, it is unlikely to define polynomially a *perfect* neighborhood which would only connect the global optimum to all other solutions.

In practice, we seek the best tradeoff between reducing the average number of possible moves and increasing the coverage rate of the highest local optima. Indeed, with an intensification strategy, local optima can be avoided by deleting appropriate edges of the hill-climbing transition graph. The deletion of these edges reduces the coverage rates of some local optima and then increases the probability to reach other local optima. Then, the goal of an efficient pivoting rule is mainly to delete edges that drive toward low-quality local optima.

Of course, the degree of refinement of a non-contextual neighborhood selection strategy depends on the available quantity of information. We categorize as *first-level*, *second-level*, or *l-level pivoting rule* a strategy requiring fitness information of solutions at a maximum distance of 1, 2, or l from the current solution. First, best and worst improvement are first-level pivoting rules, while local searches HC2 (Hill-Climber two-step) [44] and VNS (*variable neighborhood search* [31]) are respectively second and l -level pivoting rules. However, since HC2 authorizes deteriorating moves, and VNS considers jumps to non-neighboring solutions (relatively to a fixed neighborhood), these local searches cannot be considered exactly as climbers within a specific fitness landscape. In the following, we investigate second-level pivoting rules for determining relevant information to guide the search towards higher local optima, by applying improving moves only in a fixed (unvarying) fitness landscape.

3.2. Description of the Maximum Expansion pivoting rule

As recalled previously, the best improvement hill-climbing strategy tends to reach good local optima mainly on smooth and small landscapes. Oppositely, the worst improvement rule enables the search to reach higher local optima in many cases, especially on more rugged landscapes recognized as difficult to climb. Such a strategy avoids the search being trapped prematurely in a local optimum and tends to perform more steps than other classical pivoting rules.

We intuitively deduced that worst improvement is less likely to be trapped in low-quality local optima since it keeps more improvement possibilities

during the search. Since on regular fitness landscapes the fitness of solutions is correlated with their number of improving neighbors, we investigated if the dynamics of the worst improvement rule improves the number of remaining moves or paths. In this case, the relative efficiency of first-level fitness-based pivoting rules (such as first improvement vs. best improvement) could then be mainly explained by their ability to keep many search paths accessible. These previous studies led us to consider a *maximum expansion* pivoting rule (ME) which aims at keeping the most improvement possibilities during the search. ME is guided by the *expansion score* of solutions, corresponding to their number of improving neighbors. ME is a second-level pivoting rule since the expansion score requires the evaluation of an additional level of neighborhood.

The maximum expansion hill-climbing consists then of selecting, at each step, the improving neighbor having the highest expansion score (depicted in Figure 4). Here, we consider two variants. In the first one (ME_F, or simply ME), ties are broken randomly, except once all improving neighbors' expansion scores are equal to zero. In this case, since all improving neighbors are local optima, the natural choice consists of choosing the best one (according to the fitness value). In the second variant (denoted as ME_B), the fitness value is used as a secondary criterion (highest fitness first) to discriminate solutions of equal expansion score.

Naturally, using ME induces a huge computation cost mostly due to the evaluation of two levels of neighborhood at each step of the search. However, let us recall that our focus here is to verify if such a criterion can be more relevant to determine better hill-climbing trajectories than directly fitness-based criteria. Practically, implementation choices will determine which effort can be devoted to collecting information that can help guide the search toward better solutions, to reach an ideal tradeoff between efficiency and computation cost. Such an approach has been used for the worst improvement climber which unfortunately requires a high computational effort, but once approximated can lead to efficient tradeoffs [7].

Table 1 recalls the 5 pivoting rules considered in the experiments reported in the next section. This table also describes the selection processes, considering $\mathcal{I}_x = \{y \in \mathcal{N}(x) : f(y) > f(x)\}$ being the set of strictly improving neighbors of a solution x . F, B, W, ME, and ME_B pivoting rules consist then in choosing a neighbor within the sets as defined in Table 1. Let us notice that sets can be empty, and correspond to the end of the climbing process (a local optimum is attained).




Figure 4: Depiction of the behavior of B, W and ME from a solution x . Only improving moves are depicted. F can select any solution from the \mathcal{A}_1 area.

Table 1: Pivoting rules selection processes.

| Pivoting rule | Set of selectable solutions |
|---------------|---|
| F | \mathcal{I}_x |
| B | $\operatorname{argmax}_{y \in \mathcal{I}_x} f(y)$ |
| W | $\operatorname{argmin}_{y \in \mathcal{I}_x} f(y)$ |
| ME | $\operatorname{argmax}_{y \in \mathcal{I}_x} \#\mathcal{I}_y$ |
| ME_B | $\operatorname{argmax}_{z \in \mathcal{E}} f(z)$, with $\mathcal{E} = \operatorname{argmax}_{y \in \mathcal{I}_x} \#\mathcal{I}_y$ |




Figure 5: Fictive landscape. Vertices represent solutions with their associated fitness values, edges indicate neighborhood relations (8 neighbors for each solution). S_{ij} refers to the solution belonging to the i^{th} row and the j^{th} of the figure. Starting from S_{34} (squared), improving neighbors are in gray. Fitness of local optima appear underlined and in bold (4 local optima here).

3.3. Numerical Example

Figure 5 represents a portion of a landscape, reporting the fitness values of 40 solutions, and their neighborhood being represented by solutions connected with edges. In the following, let us consider that the starting solution is S_{34} of fitness 10, which has six improving neighbors.

Using the first improvement pivoting rule (F), the move to perform is selected uniformly at random among the six improving neighbors (S_{24} , S_{33} , S_{35} , S_{43} , S_{44} , S_{45}).

The best improvement pivoting rule (B) selects S_{43} , which has the best fitness (17) among all improving neighbors. The next and final step leads to S_{42} , a local optimum of fitness 18.

The worst improvement pivoting rule (W) selects S_{24} , which has the worst fitness (11) among improving neighbors. The next steps will reach solutions S_{14} , S_{13} , S_{22} , S_{32} , S_{43} , and finally the local optimum S_{42} (fitness 18).

The maximum expansion pivoting rule (ME) selects either S_{33} or S_{35} , the two improving neighbors that have the most improving neighbors (five each). From S_{33} , the next steps are S_{22} , S_{32} , S_{21} , S_{12} , and finally S_{11} (fitness 20). From S_{35} , the next steps are S_{46} , S_{36} , and finally either S_{45} or S_{37} (with respective fitness of 16 and 17).

The other variant of the maximum expansion pivoting rule (ME_B) selects

S_{33} , since $f(S_{33}) > f(S_{35})$. Then from S_{33} , the next steps are S_{22} , S_{32} , S_{21} , S_{12} , and finally S_{11} (fitness 20).

4. Comparison of the Maximum Expansion with other pivoting rules

4.1. Experimental protocol

The present experiments are mainly realized on NK landscapes to measure the ability of climbers to reach good local optima w.r.t. the considered pivoting rule and the properties (size and ruggedness) of landscapes. The goal is to assess the capacity of the maximum expansion to attain good local optima.

The NK landscape model [22] can be used to generate random combinatorial fitness landscapes with tunable size and ruggedness. The NK model is widely used in evolutionary optimization to study the behavior of search methods [2, 35, 14, 25, 1] and in the context of more applied work [17, 16, 13]. An N -size NK landscape \mathcal{L} is an N -dimensional hypercube fitness landscape, ie. with a search space $\mathcal{X} = \{0, 1\}^N$ and a neighborhood \mathcal{N} which links solutions with a Hamming distance of 1. The properties of NK landscapes are determined through two parameters N and K . N refers to the length of the bit string solutions and $K \in \{0, \dots, N - 1\}$ specifies the number of variable interdependencies, thus the ruggedness of the landscape. An NK function is based on N pseudo-Boolean sub-functions $\gamma_i : \{0, 1\}^{K+1} \rightarrow [0, 1]$. The objective function to be maximized is $f(x) = \frac{1}{N} \sum_{i=1}^N \gamma_i(x_i, x_{l_{i1}}, \dots, x_{l_{iK}})$, provided that each bit i is linked with K other bits $l_{i1}, l_{i2}, \dots, l_{iK}$. Setting K to 0 leads to a completely smooth landscape containing a single local (then global) optimum. Increasing K induces an increase in the ruggedness levels of generated landscapes.

We use 28 randomly generated NK landscapes of various log sizes and variable interdependency levels which are reported in Table 2

Since NK functions naturally induce a specific landscape structure, we also consider another type of binary landscapes with the same hypercube topology (bit-flip neighborhood) but using UBQP fitness functions (Unconstrained Binary Quadratic Programming problem [20]). Two classes of permutation-based landscapes are also included: QAP (Quadratic Assignment Problem [23] fitness functions, with the neighborhood defined by the swap operator) and FSP (Flow-shop Scheduling Problem fitness functions, with the neighborhood defined by the insertion operator) to have a more

Table 2: Instances parameters for each problem.

| Problem | #inst. | # decision variables | other |
|---------|--------|---|-----------------------------------|
| NK | 28 | $N \in \{128, 256, 512, 1024\}$ | $K \in \{1, 2, 4, 6, 8, 10, 12\}$ |
| UBQP | 24 | $N \in \{128, 256, 512, 1024, 2048, 4096\}$ | $d \in \{10, 25, 50, 75, 100\}$ |
| FSP | 118 | $n \in \{20, 30, 50\}$ | $m \in \{5, 10, 15, 20\}$ |
| QAP | 92 | $n \in [12, 64]$ | - |

precise view of the behavior of climber variants. UBQP, QAP, and FSP are well-known academic NP-hard combinatorial optimization problems.

A UBQP instance is composed of an $N \times N$ matrix Q of constants q_{ij} which can be positive or negative. A solution is a binary vector x of size N where $x_i \in \{0, 1\}$ corresponds to the i -th element of x . The objective function to be maximized is $f(x) = \sum_{i=1}^N \sum_{j=1}^N q_{ij} x_i x_j$. We use a set of 24 instances obtained with the Palubeckis instance generator [34]. Instances have different sizes and densities reported in Table 2. The density d affects the percentage of values equal to zero in the matrix Q : $d = 0$ leads to a matrix full of zero except on the diagonal whereas a $d = 100$ leads to a matrix with no zero (except possibly on the diagonal).

The Flow-shop Scheduling Problem is an NP-complete problem that consists of scheduling n jobs on m machines [19]. Each job J_i is composed of m consecutive operations $(t_{ik}, k \in \{1, \dots, m\})$ with a specific processing time $(p_{ik}, k \in \{1, \dots, m\})$. The k -th operation of a job must be performed on the k -th machine in a specified order. In the variant under consideration, jobs must be scheduled in the same order on all machines. The objective function to be minimized refers to the total completion time, also called *makespan*: $f(\pi) = \max_{i \in \{1, \dots, n\}} \{s_{im} + p_{im}\}$. s_{ij} corresponds to the moment where operation t_{ij} is performed in order to respect job order $\pi \in S_n$, with S_n the set of permutations of size n . As landscape fitness functions, we use 118 Taillard Flow-shop instances [41] with sizes varying from 20 to 50 jobs and 5 to 20 machines.

QAP aims at assigning n facilities to a set of n locations. A distance d_{ij} is specified for each pair (i, j) of locations while a flow f_{rs} is given for each pair (r, s) of facilities. The fitness function to be minimized is the sum of distances multiplied by the corresponding flows: $f(\pi) = \sum_{i=1}^n \sum_{j=1}^n d_{ij} f_{\pi_i \pi_j}$. π_i represents the i -th element of the solution $\pi \in S_n$, with S_n the set of permutations of size n . We consider 92 QAP landscapes, using instances from the QAPLIB of size varying from 12 to 64.

In this section, we perform two sets of experiments. The first one aims to assess ME efficiency against classic pivoting rules, while the second one aims to assess its efficiency against pivoting rules using the same level of knowledge (two neighborhood levels) to select an improving neighbor. Specific details regarding each set of experiments are given in the corresponding section. In both cases, experiments consist of 100 hill-climbing executions for each couple (pivoting rule, landscape) starting from identical sets of 100 randomly generated initial solutions. This protocol allows us to perform accurate statistical analysis and reduces the possible stochastic bias resulting from random initializations. Note that for each problem, each instance described in Table 2 corresponds to one landscape. We use a *1-flip* neighborhood operator for NK and UBQP instances, the *insert* operator for FSP, and the *swap* operator for QAP, which exchanges two elements that are not necessarily adjacent.

As the aim of this section is to assess ME ability to attain good local optima, we only focus on the quality of the local optima reached. Therefore, on each landscape, we report the average fitness of local optima over 100 runs for each pivoting rule, and the performance rank of each pivoting rule. The performance rank reported for a given pivoting rule P is equal to $1+D$, where D is the number of pivoting rules that statistically dominate P w.r.t. a binomial test (p -value < 0.001). Thus, the method ranked 1 is the one that statistically dominates the most methods, while the method ranked last is statistically the most dominated. For experiments on FSP and QAP landscapes, we only report the global average rankings since the 1-ruggedness indicator originally proposed for binary landscapes is not directly extendable to permutation-based landscapes.

4.2. Comparison of F , B , W , ME , and ME_B

In this first experiment, we compare the five climber variants, all following strictly improving trajectories, stopping when a local optimum is reached. Note that every ME, ME_B , best and worst improvement trajectory is a potential first improvement one. We recall that it is not the case of previously mentioned VNS or HC2 local searches which perform either deteriorating or 2-step moves.

Table 3 reports for each couple (pivoting rule, NK landscape) and Table 4 (pivoting rule, UBQP landscape) the average fitness value of the 100 local optima reached as well as the *performance rank* of each pivoting rule.

Table 3: Statistical comparison of hill-climbing strategies on NK fitness landscapes. Results are detailed for each instance. The best average fitness for each instance appears in bold.

| Landscape | N K | 1-rug. | [F]irst | | [B]est | | [W]orst | | ME | | ME _B | |
|-----------------------|---------|--------|---------|----------|--------|----------|---------|----------|------|---------------|-----------------|---------------|
| | | | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. |
| NK landscapes (N K) | 128 1 | 0.53 | 4 | 0.7020 | 2 | 0.7098 | 5 | 0.6920 | 2 | 0.7093 | 1 | 0.7147 |
| | 128 2 | 2.40 | 4 | 0.7042 | 3 | 0.7076 | 5 | 0.6990 | 1 | 0.7201 | 1 | 0.7234 |
| | 128 4 | 3.13 | 3 | 0.7266 | 3 | 0.7251 | 3 | 0.7275 | 1 | 0.7539 | 1 | 0.7542 |
| | 128 6 | 5.23 | 4 | 0.7176 | 5 | 0.7137 | 3 | 0.7228 | 1 | 0.7432 | 1 | 0.7435 |
| | 128 8 | 7.38 | 4 | 0.7176 | 5 | 0.7128 | 3 | 0.7223 | 1 | 0.7421 | 1 | 0.7438 |
| | 128 10 | 9.60 | 3 | 0.7056 | 3 | 0.7058 | 3 | 0.7112 | 1 | 0.7363 | 1 | 0.7366 |
| | 128 12 | 11.83 | 3 | 0.7011 | 3 | 0.6977 | 3 | 0.701 | 1 | 0.7287 | 1 | 0.7275 |
| | 256 1 | 0.25 | 4 | 0.7016 | 2 | 0.7078 | 5 | 0.6992 | 2 | 0.7081 | 1 | 0.7125 |
| | 256 2 | 0.65 | 4 | 0.7077 | 3 | 0.7109 | 5 | 0.7011 | 1 | 0.7213 | 1 | 0.7237 |
| | 256 4 | 1.49 | 3 | 0.7218 | 3 | 0.7209 | 3 | 0.7272 | 1 | 0.7511 | 1 | 0.7504 |
| | 256 6 | 2.64 | 4 | 0.7264 | 4 | 0.7252 | 3 | 0.7333 | 1 | 0.7524 | 1 | 0.7528 |
| | 256 8 | 4.06 | 4 | 0.7157 | 4 | 0.7127 | 3 | 0.7261 | 1 | 0.7429 | 1 | 0.7416 |
| | 256 10 | 5.29 | 4 | 0.7107 | 4 | 0.7079 | 3 | 0.7215 | 1 | 0.7368 | 1 | 0.7365 |
| | 256 12 | 6.51 | 4 | 0.7066 | 5 | 0.7041 | 3 | 0.7117 | 1 | 0.7303 | 1 | 0.7291 |
| | 512 1 | 0.12 | 4 | 0.6902 | 2 | 0.6959 | 5 | 0.6845 | 3 | 0.6944 | 1 | 0.6998 |
| | 512 2 | 0.32 | 4 | 0.7129 | 3 | 0.7173 | 5 | 0.7106 | 2 | 0.7280 | 1 | 0.7311 |
| | 512 4 | 0.79 | 4 | 0.7214 | 4 | 0.7203 | 3 | 0.7263 | 1 | 0.7460 | 1 | 0.7454 |
| | 512 6 | 1.39 | 4 | 0.7248 | 5 | 0.7212 | 3 | 0.7318 | 1 | 0.7501 | 1 | 0.7495 |
| | 512 8 | 2.14 | 4 | 0.7193 | 4 | 0.7165 | 3 | 0.7299 | 1 | 0.7460 | 1 | 0.7447 |
| | 512 10 | 2.67 | 4 | 0.7139 | 4 | 0.7116 | 3 | 0.7244 | 1 | 0.7381 | 1 | 0.7379 |
| | 512 12 | 3.55 | 4 | 0.7096 | 5 | 0.7023 | 3 | 0.7187 | 1 | 0.7317 | 1 | 0.7304 |
| | 1024 1 | 0.07 | 4 | 0.6969 | 2 | 0.7033 | 5 | 0.6885 | 3 | 0.7014 | 1 | 0.7067 |
| | 1024 2 | 0.15 | 4 | 0.7155 | 3 | 0.7192 | 5 | 0.7118 | 2 | 0.7307 | 1 | 0.7349 |
| | 1024 4 | 0.39 | 4 | 0.7242 | 4 | 0.7235 | 3 | 0.7300 | 1 | 0.7492 | 1 | 0.7491 |
| | 1024 6 | 0.67 | 4 | 0.7240 | 4 | 0.7211 | 3 | 0.7344 | 1 | 0.7501 | 1 | 0.7493 |
| | 1024 8 | 1.08 | 4 | 0.7218 | 5 | 0.7175 | 3 | 0.7330 | 1 | 0.7465 | 1 | 0.7459 |
| | 1024 10 | 1.40 | 4 | 0.7166 | 5 | 0.7121 | 3 | 0.7266 | 1 | 0.7392 | 1 | 0.7387 |
| | 1024 12 | 1.83 | 4 | 0.7111 | 5 | 0.7063 | 3 | 0.7206 | 1 | 0.7323 | 1 | 0.7319 |

Table 4: Statistical comparison of hill-climbing strategies on UBQP landscapes and permutation-based fitness landscapes. Results on UBQP landscapes are detailed for each instance. On permutation-based landscapes, we only report the average performance rank. The best average fitness for each instance appears in bold.

| Landscape | $N d$ | 1-rug. | [F]irst | | [B]est | | [W]orst | | ME | | ME _B | |
|-----------------|----------|--------|-----------|-----------|-------------|-----------|-------------------|------------|------------------|------------------|-----------------|------------------|
| | | | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. |
| UBQP landscapes | 128 25 | 2.34 | 3 | 23398.1 | 3 | 23388.9 | 3 | 23519.9 | 1 | 23730.33 | 1 | 23724.9 |
| | 128 50 | 3.15 | 4 | 32784.6 | 5 | 32458.5 | 3 | 32913.4 | 1 | 33331.97 | 1 | 33342.2 |
| | 128 75 | 3.59 | 4 | 45435.3 | 5 | 45282.5 | 3 | 45571.5 | 2 | 45782.3 | 1 | 45754.3 |
| | 128 100 | 4.36 | 4 | 50104.6 | 4 | 49896.8 | 1 | 50475.6 | 1 | 50559.87 | 1 | 50474.5 |
| | 256 25 | 1.46 | 4 | 69183.4 | 5 | 69034.0 | 3 | 69518.9 | 1 | 69977.5 | 2 | 69899.8 |
| | 256 50 | 2.13 | 3 | 102037.5 | 3 | 101866.6 | 3 | 102121.2 | 1 | 102628.42 | 1 | 102635.0 |
| | 256 75 | 2.66 | 4 | 132076.8 | 4 | 131494.5 | 3 | 133027.2 | 2 | 133273.14 | 1 | 133370.9 |
| | 256 100 | 2.85 | 4 | 144187.0 | 5 | 143724.4 | 3 | 144817.0 | 1 | 145316.5 | 1 | 145262.5 |
| | 512 10 | 0.58 | 3 | 132133.4 | 3 | 131881.8 | 3 | 131954.0 | 1 | 133258.7 | 1 | 133362.8 |
| | 512 50 | 1.48 | 4 | 284049.7 | 4 | 283560.3 | 1 | 286723.2 | 1 | 287135.7 | 1 | 286784.0 |
| | 512 75 | 1.84 | 4 | 361205.4 | 4 | 360518.1 | 3 | 362802.6 | 1 | 363653.7 | 1 | 363497.2 |
| | 512 100 | 2.26 | 3 | 419889.8 | 5 | 418458.4 | 3 | 421017.2 | 1 | 422248.6 | 1 | 422114.4 |
| | 1024 10 | 0.55 | 4 | 350574.0 | 4 | 349469.4 | 3 | 351940.9 | 1 | 352975.3 | 1 | 352923.4 |
| | 1024 50 | 1.03 | 4 | 797468.9 | 5 | 795302.1 | 3 | 800153.2 | 1 | 801953.0 | 2 | 801510.5 |
| | 1024 75 | 1.14 | 4 | 998493.2 | 5 | 995156.7 | 1 | 1001466.5 | 1 | 1001920.6 | 1 | 1001766.6 |
| | 1024 100 | 1.46 | 4 | 1130314.9 | 5 | 1127151.2 | 3 | 1134202.9 | 1 | 1136000.9 | 1 | 1136132.8 |
| | 2048 10 | 0.29 | 4 | 991463.9 | 5 | 988944.7 | 3 | 994765.6 | 1 | 998034.0 | 1 | 997796.7 |
| | 2048 25 | 0.66 | 4 | 1626645.6 | 5 | 1623603.0 | 3 | 1630519.19 | 1 | 1633614.0 | 2 | 1632281.3 |
| | 2048 50 | 0.87 | 4 | 2377355.3 | 5 | 2373279.2 | 3 | 2382892.94 | 1 | 2386536.3 | 1 | 2386386.4 |
| | 2048 100 | 1.08 | 4 | 3058752.1 | 5 | 3045996.1 | 1 | 3074422.1 | 1 | 3074440.5 | 1 | 3070783.2 |
| 4096 10 | 0.18 | 4 | 2775627.4 | 4 | 2771590.4 | 3 | 2788231.89 | 1 | 2791760.4 | 2 | 2790432.0 | |
| 4096 25 | 0.35 | 4 | 4552524.9 | 4 | 4545855.0 | 1 | 4568967.64 | 1 | 4570193.9 | 1 | 4568413.8 | |
| 4096 50 | 0.51 | 4 | 6470526.6 | 5 | 6457299.0 | 1 | 6495855.44 | 1 | 6495310.9 | 1 | 6492171.0 | |
| 4096 100 | 0.66 | 4 | 9014549.8 | 5 | 8997240.0 | 1 | 9053431.85 | 2 | 9044430.4 | 2 | 9039280.8 | |
| FSP land. | - | - | 3.44 | - | 3.83 | - | 3.00 | - | 1.02 | - | 1.01 | - |
| QAP land. | - | - | 3.17 | - | 3.22 | - | 2.94 | - | 1.03 | - | 1.02 | - |

Results on NK landscapes (Table 3) highlight that fitness maximization is a useful criterion to guide the search on smooth landscapes. Indeed, best improvement (B) often reaches better local optima than F, W, and ME on very smooth landscapes, while ME_B statistically dominates B. The expansion score criterion is also relevant to guide the search on such landscapes, yet the fitness criterion must be used as a secondary criterion. On landscapes with a higher number of local optima (more rugged and/or larger), ME mostly dominates other methods, showing that the expansion score is a relevant criterion to determine adequate NK landscapes climbing trajectories. On these landscapes, ME and ME_B obtain similar results, since expansion scores are more diversified, leading to a reduced use of the second criterion.

On UBQP landscapes, ME and ME_B are also, in general, the most effective pivoting rules for reaching higher local optima. Among first-level pivoting rules, worst improvement is the most efficient. On several UBQP landscapes, especially the larger ones, a worst improvement climber finds local optima of quality comparable to ME. Permutation-based landscapes results (FSP, QAP) (Table 4) confirm the efficiency of ME to attain good local optima whatever the neighborhood structure or landscape characteristics.

Although these comparisons show that the expansion score is an interesting criterion to guide the search toward good local optima, we keep in mind that ME uses the neighborhood’s neighbors’ knowledge to select a solution, contrary to the first-level climbers. The next section compares ME with pivoting rules using the same level of knowledge (two levels of neighborhood) to assess the relevance of the expansion score criterion from a more operational perspective.

4.3. Extended experiments: second level pivoting rules

Here, we propose to use two pivoting rules derived from classical best and first improvement but using the knowledge of two levels of neighborhood. The extended best pivoting rule (noted as B_{ext}) selects the improving neighbor with the best-improving solution in its neighborhood. The extended first (F_{ext}) selects an improving neighbor which has at least one improving neighbor. The only difference of F_{ext} compared to a first-level first improvement climber is that a local optimum cannot be selected if an alternative exists. Such as ME and ME_B , B_{ext} and F_{ext} perform the last step using the best improvement (when all improving neighbors are local optima). All climbers ultimately reach a local optimum that indicates the end of the search process.

Table 5: Statistical comparison of second-level pivoting rules (NK landscapes).

| Landscape | | F_{ext} | | B_{ext} | | ME | | ME_B | | |
|---------------|------|------------------|----------|------------------|----------|--------|----------|---------------|----------|---------------|
| N | K | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. | |
| NK landscapes | 128 | 1 | 4 | 0.7047 | 2 | 0.7108 | 2 | 0.7093 | 1 | 0.7147 |
| | 128 | 2 | 4 | 0.7032 | 3 | 0.7084 | 1 | 0.7201 | 1 | 0.7234 |
| | 128 | 4 | 3 | 0.7267 | 3 | 0.7276 | 1 | 0.7539 | 1 | 0.7542 |
| | 128 | 6 | 3 | 0.7157 | 3 | 0.7185 | 1 | 0.7432 | 1 | 0.7435 |
| | 128 | 8 | 3 | 0.7133 | 3 | 0.7139 | 1 | 0.7421 | 1 | 0.7438 |
| | 128 | 10 | 3 | 0.7099 | 3 | 0.7104 | 1 | 0.7363 | 1 | 0.7366 |
| | 128 | 12 | 3 | 0.7016 | 3 | 0.7027 | 1 | 0.7287 | 1 | 0.7275 |
| | 256 | 1 | 4 | 0.7032 | 2 | 0.7086 | 2 | 0.7081 | 1 | 0.7125 |
| | 256 | 2 | 4 | 0.7065 | 3 | 0.7125 | 1 | 0.7213 | 1 | 0.7237 |
| | 256 | 4 | 3 | 0.7222 | 3 | 0.7242 | 1 | 0.7511 | 1 | 0.7504 |
| | 256 | 6 | 3 | 0.7263 | 3 | 0.7264 | 1 | 0.7524 | 1 | 0.7528 |
| | 256 | 8 | 3 | 0.7184 | 3 | 0.7156 | 1 | 0.7429 | 1 | 0.7416 |
| | 256 | 10 | 3 | 0.7118 | 3 | 0.7095 | 1 | 0.7368 | 1 | 0.7365 |
| | 256 | 12 | 3 | 0.7046 | 3 | 0.7034 | 1 | 0.7303 | 1 | 0.7291 |
| | 512 | 1 | 4 | 0.6918 | 2 | 0.6967 | 3 | 0.6944 | 1 | 0.6998 |
| | 512 | 2 | 4 | 0.7148 | 3 | 0.7186 | 2 | 0.7280 | 1 | 0.7311 |
| | 512 | 4 | 3 | 0.7207 | 3 | 0.7218 | 1 | 0.7460 | 1 | 0.7454 |
| | 512 | 6 | 3 | 0.7219 | 3 | 0.7226 | 1 | 0.7501 | 1 | 0.7495 |
| | 512 | 8 | 3 | 0.7194 | 3 | 0.7188 | 1 | 0.7460 | 1 | 0.7447 |
| | 512 | 10 | 3 | 0.7135 | 3 | 0.7111 | 1 | 0.7381 | 1 | 0.7379 |
| | 512 | 12 | 3 | 0.7072 | 3 | 0.705 | 1 | 0.7317 | 1 | 0.7304 |
| | 1024 | 1 | 4 | 0.6982 | 2 | 0.7039 | 3 | 0.7014 | 1 | 0.7067 |
| | 1024 | 2 | 4 | 0.7168 | 3 | 0.7208 | 2 | 0.7307 | 1 | 0.7349 |
| | 1024 | 4 | 3 | 0.7237 | 3 | 0.7247 | 1 | 0.7492 | 1 | 0.7491 |
| | 1024 | 6 | 3 | 0.7235 | 3 | 0.7225 | 1 | 0.7501 | 1 | 0.7493 |
| | 1024 | 8 | 3 | 0.7206 | 3 | 0.7199 | 1 | 0.7465 | 1 | 0.7459 |
| | 1024 | 10 | 3 | 0.7146 | 3 | 0.7131 | 1 | 0.7392 | 1 | 0.7387 |
| | 1024 | 12 | 3 | 0.709 | 3 | 0.7074 | 1 | 0.7323 | 1 | 0.7319 |

Table 6: Statistical comparison of second-level pivoting rules (UBQP, FSP, and QAP).

| | Landscape | F_{ext} | | B_{ext} | | ME | | ME_B | |
|-----------------|-----------|------------------|-----------|------------------|-----------|------------------|------------------|-----------|------------------|
| | $N d$ | rank | avg fit. | rank | avg fit. | rank | avg fit. | rank | avg fit. |
| UBQP landscapes | 128 25 | 3 | 23421.59 | 3 | 23400.4 | 1 | 23730.33 | 1 | 23724.86 |
| | 128 50 | 3 | 32992.88 | 3 | 32725.77 | 1 | 33331.97 | 1 | 33342.19 |
| | 128 75 | 3 | 45484.38 | 3 | 45298.95 | 1 | 45782.3 | 1 | 45754.33 |
| | 128 100 | 3 | 50202.02 | 4 | 49915.75 | 1 | 50559.87 | 2 | 50474.46 |
| | 256 25 | 3 | 69355.88 | 3 | 69048.22 | 1 | 69977.5 | 1 | 69899.82 |
| | 256 50 | 3 | 102094.19 | 3 | 101854.72 | 1 | 102628.42 | 1 | 102635.03 |
| | 256 75 | 3 | 132599.24 | 4 | 131614.81 | 2 | 133273.14 | 1 | 133370.87 |
| | 256 100 | 3 | 144491.74 | 4 | 143805.1 | 2 | 145316.47 | 1 | 145262.47 |
| | 512 10 | 3 | 132201.41 | 3 | 132011.84 | 1 | 133258.68 | 1 | 133362.84 |
| | 512 50 | 3 | 285219.33 | 4 | 283268.66 | 1 | 287135.71 | 1 | 286783.99 |
| | 512 75 | 3 | 361503.89 | 3 | 360814.9 | 1 | 363653.65 | 1 | 363497.17 |
| | 512 100 | 3 | 420359.44 | 4 | 418785.77 | 1 | 422248.57 | 1 | 422114.42 |
| | 1024 10 | 3 | 350575.43 | 3 | 349424.62 | 1 | 352975.29 | 1 | 352923.39 |
| | 1024 50 | 3 | 797640.26 | 3 | 795075.6 | 1 | 801952.96 | 1 | 801510.47 |
| | 1024 75 | 3 | 998008.74 | 3 | 996121.54 | 1 | 1001920.6 | 1 | 1001766.6 |
| | 1024 100 | 3 | 1131859.5 | 4 | 1127576.3 | 1 | 1136000.9 | 1 | 1136132.8 |
| | 2048 10 | 3 | 992468.08 | 4 | 988779.04 | 1 | 998033.95 | 1 | 997796.65 |
| | 2048 25 | 3 | 1626688 | 4 | 1623637.6 | 1 | 1633614 | 1 | 1632281.3 |
| | 2048 50 | 3 | 2378693.7 | 4 | 2373846.9 | 1 | 2386536.3 | 1 | 2386386.4 |
| | 2048 100 | 3 | 3061028.9 | 4 | 3048932.6 | 1 | 3074440.5 | 1 | 3070783.2 |
| 4096 10 | 3 | 2777470.3 | 4 | 2770370.5 | 1 | 2791760.4 | 1 | 2790432 | |
| 4096 25 | 3 | 4553630.3 | 4 | 4544473.2 | 1 | 4570193.9 | 1 | 4568413.8 | |
| 4096 50 | 3 | 6473133.6 | 4 | 6462722.7 | 1 | 6495310.9 | 1 | 6492171 | |
| 4096 100 | 3 | 9021410.5 | 4 | 9000265.7 | 1 | 9044430.4 | 1 | 9039280.8 | |
| FSP landscapes | | 3.18 | - | 2.83 | - | 1.00 | - | 1.05 | - |
| QAP landscapes | | 3.14 | - | 2.68 | - | 1.08 | - | 1.05 | - |

Experimental results on NK landscapes (Table 5) show that ME statistically outperforms the two other pivoting rules, except on very smooth landscapes ($K = 1$). F_{ext} and B_{ext} lead the search toward only slightly better local optima than classical best and first improvement pivoting rules despite their additional knowledge.

On UBQP landscapes, ME also outperforms F_{ext} and B_{ext} . Note that F_{ext} is always statistically equivalent or better than B_{ext} , making their relative efficiency similar to the one of the original first and best improvement.

Results on permutation-based landscapes (swap/QAP, insert/FSP) (Table 6) confirm the superiority of ME to attain good local optima. This comparison with fitness-based second-level pivoting rules highlights the relevance of the expansion score to guide the search within landscapes of various structures.

5. Maximum Expansion From a Computational Budget Perspective

5.1. ME approximation with an expansion score estimator

A central aim of this work is to investigate how neighboring information can intelligently drive intensification searches toward better areas. The experiments described in the previous section clearly show the impact of the pivoting rule. In particular, we highlight the relevance of favoring moves to solutions of higher expansion scores. The only issue remains the computation cost induced by the evaluation of two levels of neighborhood, which reduces the interest in incorporating ME within operational metaheuristics. The number of evaluations of ME relatively to the first improvement is reported in Table 7 for various landscapes.

Following a previous study on the approximation of the worst improvement principle [7], in this section we propose to approximate the maximum expansion strategy to seek an adequate balance between computational effort and search efficiency. The idea is to still use the expansion criterion for choosing improving neighbors, while only evaluating partially the neighborhoods of solutions.

Therefore, we propose an ME approximated strategy using two parameters k and l ($k, l \in \{1, \dots, |\mathcal{N}|\}$, $|\mathcal{N}|$ being the size of the neighborhood), which fix the maximum number of improving neighbors to be considered for each level of neighborhood. The algorithm, denoted as ME_k^l works as follows. First, the neighborhood is evaluated until at most k improving neighbors have

Table 7: Computational budget of B, W and ME hill-climbers in comparison to a first improvement (F), on various landscapes. The column #eval indicates the average number of function evaluations performed by the least consuming one, ie. systematically strategy F. Then, we report for each strategy the factor between its computational budget and the F one ($\times 1.0$ recalls that F climbings always need fewer evaluations and are thus used here as a budget reference).

| landscape | #eval | [F]irst | [B]est | [W]orst | ME |
|--------------|-------|--------------|---------------|-----------------|-------------------|
| NK 128 1 | 0.7 | $\times 1.0$ | $\times 10.3$ | $\times 13.6$ | $\times 568.3$ |
| NK 128 4 | 0.8 | $\times 1.0$ | $\times 5.8$ | $\times 38.0$ | $\times 819.1$ |
| NK 128 12 | 0.7 | $\times 1.0$ | $\times 3.6$ | $\times 87.4$ | $\times 702.9$ |
| NK 256 1 | 1.6 | $\times 1.0$ | $\times 17.6$ | $\times 23.9$ | $\times 1931.5$ |
| NK 256 4 | 2.0 | $\times 1.0$ | $\times 9.3$ | $\times 66.3$ | $\times 2598.7$ |
| NK 256 12 | 2.0 | $\times 1.0$ | $\times 5.2$ | $\times 169.6$ | $\times 1757.7$ |
| NK 512 1 | 3.6 | $\times 1.0$ | $\times 31.3$ | $\times 43.7$ | $\times 6790.6$ |
| NK 512 4 | 4.8 | $\times 1.0$ | $\times 15.5$ | $\times 116.0$ | $\times 8638.9$ |
| NK 512 12 | 5.4 | $\times 1.0$ | $\times 7.3$ | $\times 339.5$ | $\times 5448.7$ |
| NK 1024 1 | 8.1 | $\times 1.0$ | $\times 55.6$ | $\times 77.0$ | $\times 24185.9$ |
| NK 1024 4 | 11.5 | $\times 1.0$ | $\times 26.4$ | $\times 205.6$ | $\times 29987.9$ |
| NK 1024 12 | 13.8 | $\times 1.0$ | $\times 12.2$ | $\times 617.4$ | $\times 16677.1$ |
| UBQP 128 25 | 1.0 | $\times 1.0$ | $\times 7.8$ | $\times 59.8$ | $\times 816.0$ |
| UBQP 256 25 | 2.7 | $\times 1.0$ | $\times 11.7$ | $\times 188.8$ | $\times 2359.2$ |
| UBQP 512 10 | 6.7 | $\times 1.0$ | $\times 20.4$ | $\times 242.1$ | $\times 7676.1$ |
| UBQP 1024 10 | 17.9 | $\times 1.0$ | $\times 30.0$ | $\times 845.1$ | $\times 21992.4$ |
| QAP sko42 | 5.1 | $\times 1.0$ | $\times 6.6$ | $\times 232.7$ | $\times 25558.9$ |
| QAP sko49 | 7.7 | $\times 1.0$ | $\times 6.9$ | $\times 281.6$ | $\times 38760.7$ |
| QAP sko56 | 10.7 | $\times 1.0$ | $\times 8.3$ | $\times 407.6$ | $\times 62706.5$ |
| QAP sko64 | 15.0 | $\times 1.0$ | $\times 8.8$ | $\times 498.5$ | $\times 97971.5$ |
| QAP sko72 | 21.3 | $\times 1.0$ | $\times 9.3$ | $\times 600.9$ | $\times 130497.7$ |
| QAP lipa40a | 3.7 | $\times 1.0$ | $\times 4.8$ | $\times 122.6$ | $\times 18214.8$ |
| QAP lipa40b | 4.5 | $\times 1.0$ | $\times 4.3$ | $\times 671.2$ | $\times 19118.7$ |
| QAP lipa60a | 9.7 | $\times 1.0$ | $\times 6.4$ | $\times 262.3$ | $\times 69060.8$ |
| QAP lipa60b | 12.1 | $\times 1.0$ | $\times 5.1$ | $\times 1883.4$ | $\times 58282.3$ |
| FSP 20 5 | 0.9 | $\times 1.0$ | $\times 4.1$ | $\times 62.8$ | $\times 2097.6$ |
| FSP 20 20 | 1.2 | $\times 1.0$ | $\times 4.0$ | $\times 107.5$ | $\times 3076.1$ |
| FSP 30 5 | 1.7 | $\times 1.0$ | $\times 5.1$ | $\times 125.7$ | $\times 7606.3$ |
| FSP 30 20 | 3.1 | $\times 1.0$ | $\times 5.1$ | $\times 157.0$ | $\times 9977.7$ |
| FSP 50 5 | 5.3 | $\times 1.0$ | $\times 6.5$ | $\times 181.8$ | $\times 26067.8$ |
| FSP 50 20 | 9.2 | $\times 1.0$ | $\times 7.2$ | $\times 227.0$ | $\times 42005.7$ |

been encountered. Parameter k directly affects the number of evaluations per step since at most k second-level neighborhoods will be evaluated to compute the expansion scores. Then, neighborhoods of these k improving neighbors are simultaneously evaluated until a neighbor reaches an expansion score of l . The first neighbor to reach the expansion score of l is selected. If all improving neighbors have an expansion score lower than l , then the strategy selects a neighbor with the highest expansion score. Parameter l reflects the degree of approximation for determining the solution with the maximum expansion score.

5.2. Experiments

As for the previous section, experiments are conducted on binary/pseudo-boolean (NK, UBQP) and permutation-based (QAP, FSP) landscapes. Note that for the permutation-based problems we only considered a subset of the instances used in the previous experiments: we used 9 of the largest instances from the `lipa` and `sko` families of the QAPlib, as well as 12 instances from Taillard. For these, we kept one instance (of id 01) for each couple (n, m) . We compare 32 climbers variants:

- First and best improvement.
- 5 worst improvement variants: W_k , where $k \in \{2, 4, 8, 16, |\mathcal{N}|\}$ is the first-level parameter which fixes the maximum number of improving neighbors to consider at each iteration ($W_{|\mathcal{N}|}$ corresponds to the regular worst improvement climber).
- 25 maximum expansion variants: ME_k^l , with $k, l \in \{2, 4, 8, 16, |\mathcal{N}|\}$ ($ME_{|\mathcal{N}|}^{|\mathcal{N}|}$ is the regular ME climber).

As in our previous set of experiments, we conduct for each couple (landscape, method) 100 hill-climbings starting from the same initial set of solutions, leading to 3,200 local optima (not necessarily distinct) per landscape. All local optima are ranked according to their fitness value.

We obtain as a result a set of 100 ranks per couple (landscape, method). Figure 6 reports the ranking repartition of runs for two large NK landscapes: a smooth one ($K = 1$) and a rugged one ($K = 12$). Figures 7 and 8 show for 6 NK and 2 UBQP landscapes the Pareto fronts of tradeoffs between the expected fitness of local optima and the expected number of evaluations. Figures 9 and 10 show these Pareto fronts for landscapes derived from state-of-the-art instances of QAP and FSP. Recall that QAP and FSP are minimization problems, while NK and UBQP are maximization ones.




Figure 6: Ranking repartition of the local optima found by each hill-climbing variant on two NK landscapes of size $N=1024$ (top: $K=1$ / bottom: $K=12$).

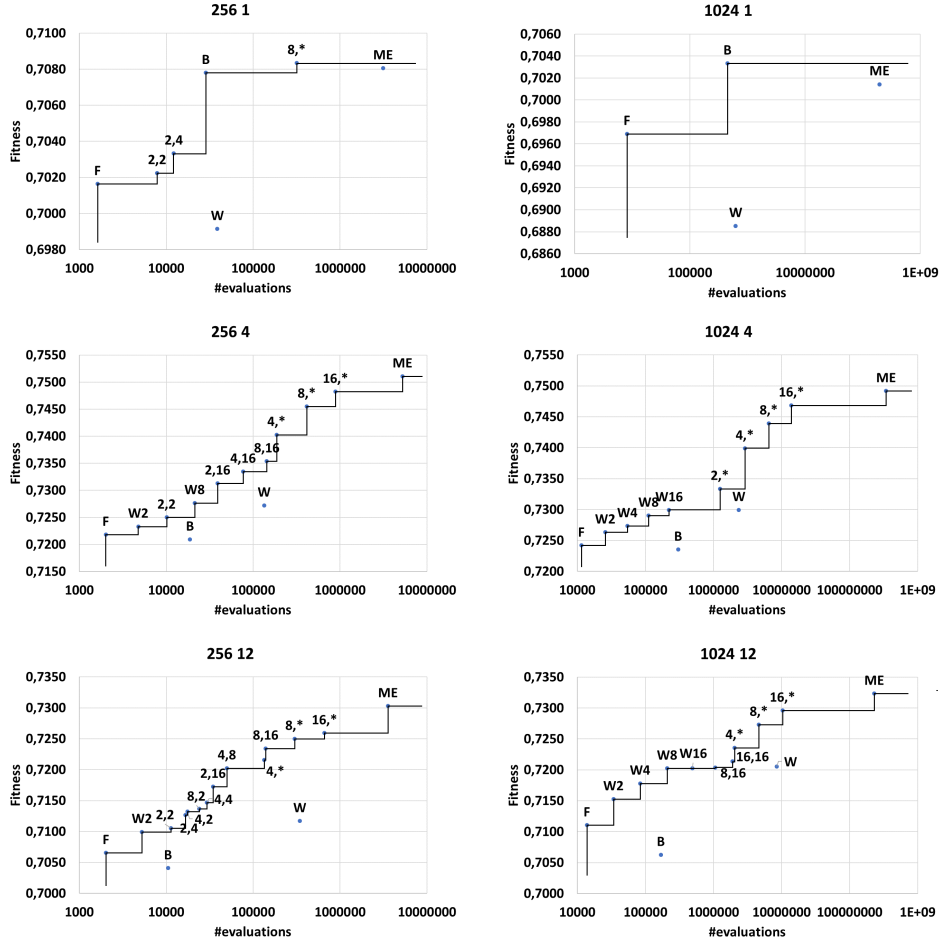


Figure 7: Comparative performance of hill-climbing strategies on various NK landscapes. The log scale x-axis represents the average number of evaluations ($\times 1000$), and the y-axis the average fitness of reached local optima. Methods offering the best tradeoffs between computation cost and local optima quality are represented by circles and named. ME approximations are only labeled by their parameters k,l (* means $|\mathcal{N}|$, ie. no expansion score approximation). We also systematically report classical pivoting rules ([F]irst, [B]est, [W]orst) and classical ME even if they do not appear in the best quality/cost compromise front.

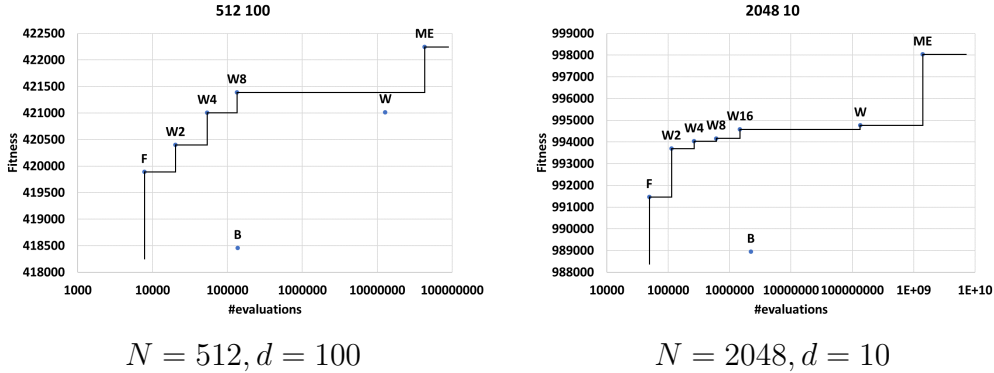


Figure 8: Comparative performance of hill-climbing strategies on UBQP landscapes (see figure 7 caption).

First, experiments highlight that in general, ME_k^l leads to better local optima than classical pivoting rules on NK landscapes. However, we observe that on very smooth NK landscapes ($K = 1$), best improvement (B) is more efficient. This corroborates Section 3.2 results, considering that here the approximation study is based on ME and not ME_B . Nevertheless, as previously stated, first and best improvement remain relevant strategies to tackle such landscapes. Keep in mind that the underlying optimization problems are not challenging, contrary to more rugged ones. On the remaining landscapes, Pareto fronts illustrate that the different ME approximation variants offer several levels of balance between quality and computational effort, that is between collected knowledge and guidance relevance.

Let us recall that in ME approximations ME_k^l , k tunes the size of the neighborhood to evaluate (especially during the first steps of the search, since it limits the number of improving neighbors) and l samples the second level neighborhood to approximate the k (at most) expansion scores. One can see that the parameter k is particularly relevant to regulate the balance between efficiency and computational effort. However, reducing the precision of the expansion score using parameter l is not counterbalanced by sufficient computational saving. On larger landscapes, W_k remains efficient when the evaluation credit is small.

For larger computational budgets, $ME_k^{|\mathcal{N}|}$ are the most efficient variants. It indicates that approximations should be mainly made in the neighborhood of current solutions by considering a partial neighborhood when sufficient. Results also reveal the relevance of the expansion score for selecting solutions,

since the additional quantity of information brought by the neighborhood directly affects the guidance ability of the climber.

Pareto fronts are of similar shape on all considered UBQP landscapes. On such landscapes, the approximations of the maximum expansion do not appear on the Pareto fronts. Methods resulting from approximations of worst improvement are systematically on the Pareto fronts, meaning they are more efficient than the approximation variants of the maximum expansion. These figures show that the worst improvement requires a high number of evaluations. Considering it only evaluates the first level of neighborhood, the convergence of the worst improvement seems slower than the convergence of ME, in terms of the number of steps required to reach a local optimum. The worst improvement strategy seems more efficient to avoid the search being prematurely trapped into local optima on UBQP landscapes. Approximating the principle of the worst improvement leads to efficient methods that outperform the first and best improvement in a significantly reduced number of evaluations. It also provides different tradeoffs between the quality of local optima and computational time, better than those provided by the approximations of the maximum expansion.

On FSP landscapes, approximations of ME and worst improvement offer numerous tradeoffs between local optima quality and number of evaluations. When the expansion principle determination is highly approximated, it leads toward better results than W_8 and W_{16} .

On QAP, ME approximations seem efficient on `lipa-a` landscapes, while on `lipa-b` and `sko` landscapes, approximating the worst improvement climber appear to be more efficient (ME approximations are rarely among the non-dominated tradeoffs), as observed on UBQP landscapes.

Depending on the landscapes under consideration, these approximations of ME can lead toward good tradeoffs between local optima quality and number of evaluations. However, such pivoting rules still require a higher number of evaluations than other considered pivoting rules.

5.3. Fitness and expansion correlation

Here, we investigate the correlation between the fitness rank and expansion scores of solutions within the same neighborhood. We conduct this study to observe whether solutions of lower fitness tend to have a better expansion score, as the maximum expansion principle was inspired by the tendency of the worst improvement to avoid prematurely low-quality local optima. More

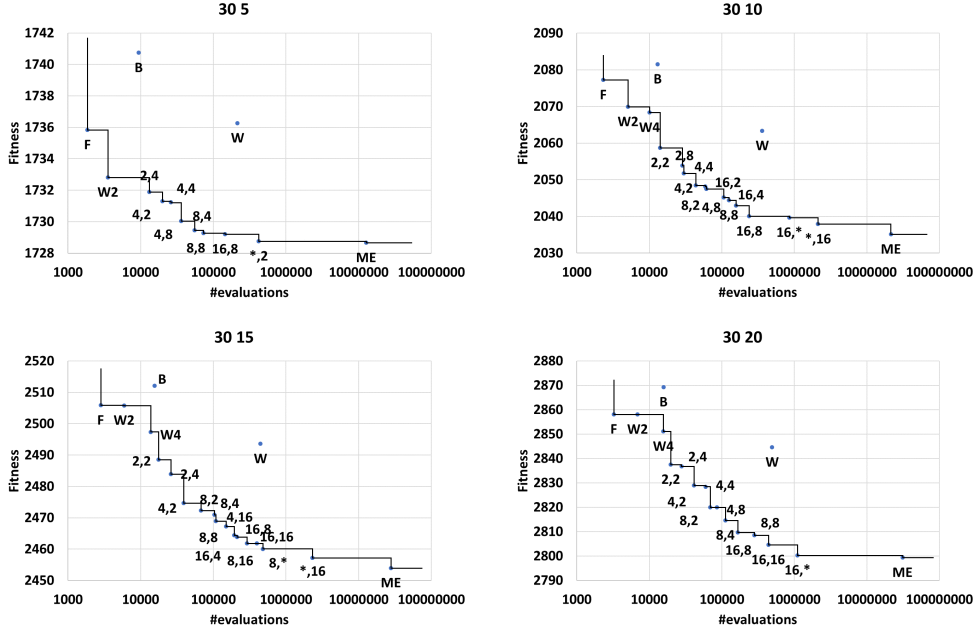


Figure 9: Comparative performance of hill-climbing strategies on FSP landscapes of size $j = 30$ and with a number of machines $m \in \{5, 10, 15, 20\}$ (see Figure 7 caption for more details).

precisely, we seek whether this correlation exists for improving solutions in a neighborhood.

For each fitness landscape considered in the previous section, we studied the ranks of the neighbors selected by ME at each step of the search. The notion of rank here refers to the fitness quality of the selected solution among all improving neighbors.

For a given landscape, we compute σ , the average rate of improving neighbors better than the selected neighbor, during the whole climbing process. We report the average σ value across 100 climbing processes. At the i -th step of a climbing process, the rate corresponds to:

$$\sigma_i = \frac{r - 1}{|\mathcal{N}_>| - 1}$$

where $\mathcal{N}_>$ refers to the set of improving neighbors and r is the rank of the solution, ranging from 1 (best solution) to $|\mathcal{N}_>|$ (worst solution). Basically, we compute, the proportion of the $|\mathcal{N}_>| - 1$ non-selected improving neighbors

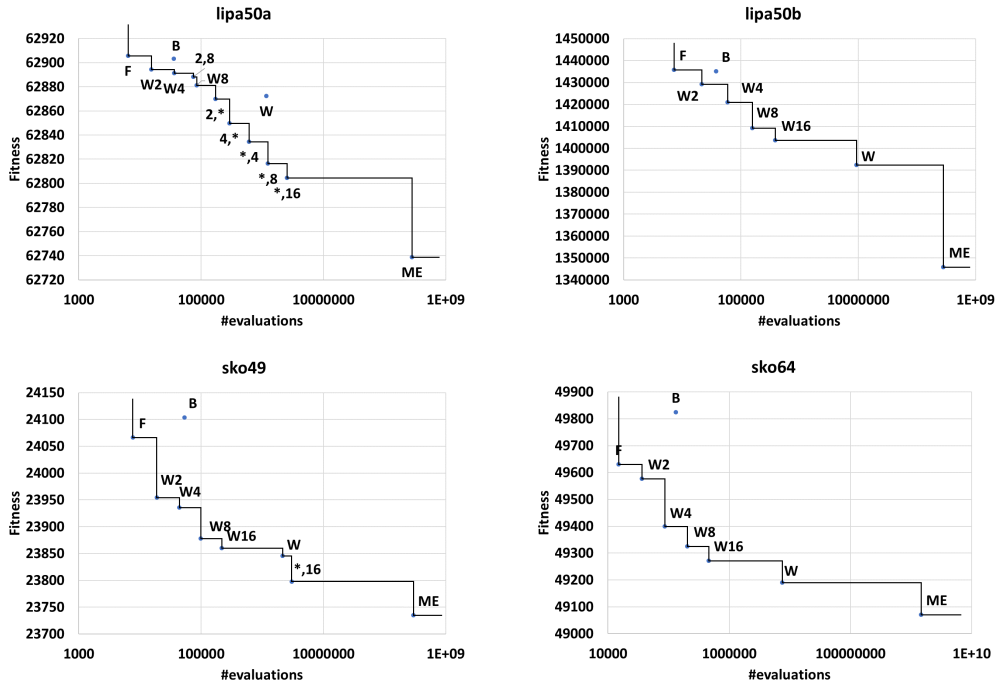


Figure 10: Comparative performance of hill-climbing strategies on QAP landscapes of various size from `lipa` and `sko` classes (see figure 7 caption).

that have a better fitness value than the solution selected by ME. Note that when $|\mathcal{N}_>| = 1$ the rate cannot be computed, the step is then ignored. This is not an issue since all pivoting rules select the same unique improving neighbor.

Results on all landscapes are reported in Figure 11. Let us notice that for each instance, the standard deviations of the average rates are very low among the 100 executions while being high when computed within each climbing process. Thus, we do not report these standard deviations to ease the readability of the figure 11.

On NK landscapes, the neighbors selected by ME are always on average among the worst 50% of the improving solutions. The more rugged the landscapes are, the lower the average rank of the selected neighbors. This suggests that when ruggedness increase, the correlation between fitness rank and expansion score increases.

On UBQP landscapes, the average fitness rank of selected improving

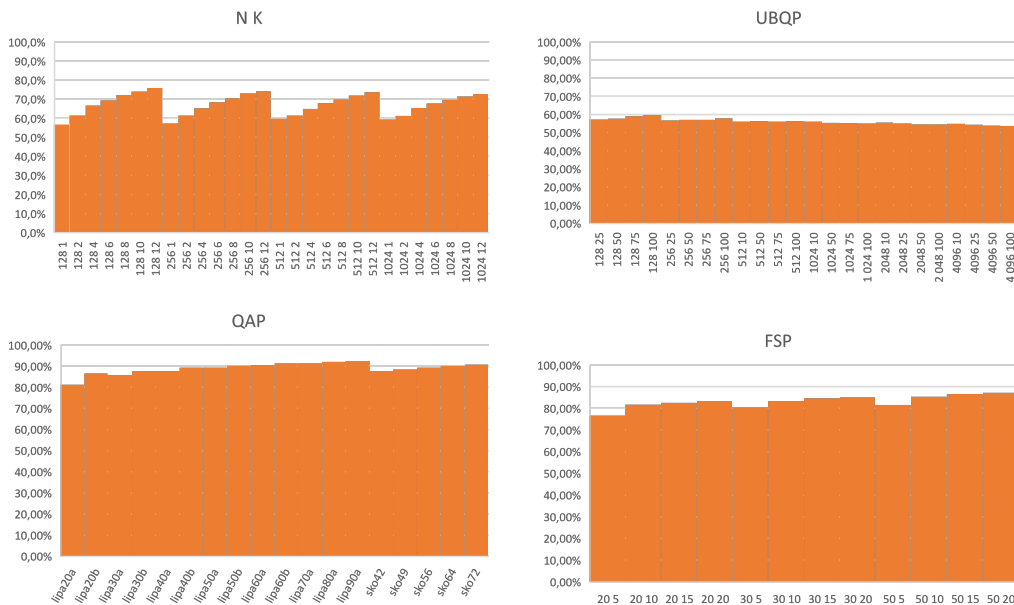


Figure 11: Average rate of improving neighbors with a better fitness than the solutions selected by ME. Values above than 50% indicate that ME tends to choose less good neighbors, which suggests a selection policy closer to a worst improvement than a best improvement.

neighbors is similar across all landscapes and does not seem to adopt a strong fitness-based policy like best or worst improvement. Across all problems, these landscapes are the ones with the lowest fitness and expansion score correlation.

On QAP and FSP landscapes, the average rank of solutions selected by ME is among the worst 25% of the improving solutions, like on the most rugged NK landscapes.

Except for UBQP landscapes and smooth NK landscapes, improving neighbors with the highest expansion score tend to be among the neighbors that improve fitness the least. Note that while the 1-ruggedness scores of UBQP landscapes and medium rugged NK landscapes are similar, the first step of autocorrelation¹ of UBQP landscapes is similar to those of very

¹The first step of autocorrelation $\rho(1)$ is a classic ruggedness indicator. In this work, we focused on the 1-ruggedness as it is based on sign-epistasis, which particularly makes sense in the context of climbers.

smooth NK landscapes [43]. The differences between 1-ruggedness and autocorrelation measures tend to indicate that UBQP landscapes are locally rugged but globally smooth, as already suggested in [43]. The fitness rank and expansion score correlation seem to be higher on landscapes shown to be rugged according to the autocorrelation score.

The negative correlation between fitness and expansion score of improving neighbors shows similarities in the behavior of ME and W. Paired with the results of previous sections, this leads to thinking that the worst improvement strategy naturally seems to be an effective way of approximating the maximum expansion principle. Additionally, the results of approximated pivoting rules convey that W approximations are often among the best tradeoffs between local optima quality and number of evaluations, even more often than ME approximations. W approximations seem either closer or similar in principle to ME than ME approximations and require a lower number of evaluations, leading them to be among the best tradeoffs. This hypothesis should be investigated further but requires its own dedicated study.

5.4. *Limitations and possible ways forward*

Results presented in the previous sections show the potential interest in considering climbers that often reach higher local optima while requiring a more expensive computational budget. Consequently, the real value of integrating ME or W principle in advanced metaheuristics could be questionable. Here, we try to partially answer this issue by (1) providing a new set of experiments, and (2) discussing possible ways to integrate the maximum expansion principle within metaheuristics.

The experiments proposed here aim to evaluate whether the gain in quality of local optima attained counterbalances the additional computation cost required by the estimation of expansion scores. To this aim, we compare three pivoting rules on a fixed computational budget within a random restart hill-climber framework. When the stopping criterion is reached, the best-encountered solution is returned. We consider the first improvement (RR-F), which corresponds to the reference climber, and the variant ME_2^2 (RR- ME_2^2), which is the maximum expansion variant requiring the lowest number of evaluations. We also include W_2 (RR- W_2), since a solution with a low fitness is intuitively more likely to have a high expansion score than a solution with a

high fitness². The selection of a solution among the least improving neighbors tends to favor solutions with high expansion scores. The advantage of such a pivoting rule is that it requires a smaller computational effort than RR-ME and its approximations since only one level of neighborhood is used. One could then consider RR-W₂ as an alternative way to approximate RR-ME₂^N and consequently RR-ME.

In the results reported hereafter, the stopping criterion corresponds to 100 million solution evaluations. Let us notice that the statistical results obtained are very similar if we extend the budget to 200 million solution evaluations. A synthesis of the results obtained on NK and UBQP landscapes is given in Table 8. We only consider these two problems as the neighborhood size is smaller and thus more favorable to the maximum expansion principle. Indeed, two levels of neighborhood represent a budget of $O(N^2)$ evaluations on pseudo-boolean problems, and $O(n^4)$ evaluations on permutation problems³.

On NK and UBQP, we observe that RR-ME₂² is not able to outperform RR-F, and RR-ME₂² is not statistically outperformed by RR-F on around half of the considered landscapes. RR-W₂ regularly outperforms RR-F on several landscapes of both problems.

More precisely, on NK landscapes, RR-ME₂² is outperformed on less rugged landscapes while on rugged landscapes ($K \in \{8, 10, 12\}$), there is no statistically significant difference between strategies, except on larger instances ($N = 1024$). The number of hill-climbing iterations, ie. the number of not necessarily distinct local optima attained for each iterated climber variant, is reported in tables 9 and 10. Over the same budget, RR-ME₂² performs around five times fewer climbings than RR-F on all landscapes, even if the ratio slightly varies according to N and K . As mentioned above, the sta-

²If we consider a fitness landscape with random fitness values (from a uniform distribution in $[0, 1]$), then the expected average expansion score of a solution x equals to $(1 - f(x)) \times |\mathcal{N}|$, where $|\mathcal{N}|$ is the number of neighbors of x . In this case, fitness is negatively correlated with the expansion score. Intuitively, such correlation can be verified on most landscapes.

³Without considering incremental or multiple evaluations, nor model parameters (K, d, m), a single evaluation costs respectively $O(N)$ for NK, $O(N^2)$ for UBQP, $O(n)$ for FSP, and $O(n^2)$ for QAP. Since the comparisons are performed considering a fixed budget expressed with the same number of evaluations for all problems, ie. within a black-box perspective, we point out here the complexity issues which specifically come from the exploration of the second-level neighborhoods.

Table 8: Summary of the comparative performance of a random-restart first improvement hill-climbing (RR-F), a random-restart based on an approximated worst improvement (RR-W₂²), and a random-restart expansion based hill-climbing (RR-ME₂²), considering a fixed budget of 100 million evaluations. For each hill-climber on each problem, we report the number of instances where the climber dominates or is dominated by another climber. Symbol ”+” for a column A \succ B means that method A is statistically better than method B. Symbol ”-” means that B is statistically better than A. Symbol ”~” means no significant dominance.

| Prob. | dom. | RR-W ₂ \succ RR-F | RR-ME ₂ ² \succ RR-F | RR-W ₂ \succ RR-ME ₂ ² |
|-------|------|--------------------------------|--|---|
| NK | + | 6 | 0 | 11 |
| | - | 8 | 17 | 3 |
| | ~ | 14 | 11 | 14 |
| UBQP | + | 7 | 0 | 12 |
| | - | 0 | 13 | 0 |
| | ~ | 17 | 11 | 12 |

Table 9: Number of hill-climbing iterations performed by RR-F (#F) on NK landscapes for different problem sizes (N), in comparison to RR-W₂ (#W) and RR-ME₂² (#ME). Allocated budget: 100 million evaluations.

| K | #F | | | | Ratio | |
|-----|-----------|-----------|-----------|------------|-------|--------|
| | $N = 128$ | $N = 256$ | $N = 512$ | $N = 1024$ | #F/#W | #F/#ME |
| 1 | 144153 | 61859 | 28122 | 12295 | 1.89 | 4.87 |
| 2 | 134519 | 57027 | 24205 | 10548 | 2.07 | 4.91 |
| 4 | 124984 | 49492 | 20663 | 8723 | 2.25 | 5.05 |
| 6 | 125918 | 48339 | 19244 | 7871 | 2.37 | 5.18 |
| 8 | 129712 | 48914 | 18818 | 7504 | 2.46 | 5.30 |
| 10 | 136366 | 50269 | 18957 | 7368 | 2.54 | 5.43 |
| 12 | 144484 | 52162 | 19330 | 7335 | 2.60 | 5.55 |

tistical dominance between the three variants does not change when runs are performed for twice as many evaluations. A possible explanation of the difference found on larger landscapes concerns the level of approximation of RR-ME₂². Indeed, on larger landscapes, using $k = 2$ and $l = 2$ is not sufficient to approximate the ME principle as the number of neighbors $|\mathcal{N}|$ is high, while it requires more budget anyway, limiting the number of restarts. Moreover, on results already reported in Figure 7, we can see that for unique climbing, the pivoting rules F, W₂, and ME₂² are among the best tradeoffs on smaller NK landscapes ($N = 512$), while it is only the case for F and W₂ on large landscapes ($N = 1024$).

Using the expansion score induces a reduction of the impact of local ruggedness on climber efficiency. This could explain the fact that RR-ME₂² is more effective on rugged landscapes. RR-W₂ outperforms RR-F on 6 large and rugged landscapes. RR-W₂ is statistically outperformed on smooth landscapes ($K \in \{1, 2\}$), which was expected since [42] reports that the worst improvement hill-climbing tends to reach high local optima only on rugged landscapes. On UBQP landscapes, RR-ME₂² and RR-F are of similar efficiency on the smallest instances ($N \in \{128, 256, 512\}$) whereas RR-ME₂² is systematically outperformed on the largest instances. RR-W₂ is never statistically outperformed by RR-F on the considered instances. Note that, RR-W₂ statistically outperforms RR-F on most of the largest instances ($N \in \{1024, 2048, 4096\}$). Such a result leads us to think that while expansion scores are expensive to compute, it should be interesting to consider methods that naturally tend to keep improving possibilities during the search. Those methods could be mostly fitness-based as the worst improvement variants, or expansion-based at least partially.

These results show that the additional computational cost induced by directly estimating the expansion score is not sufficiently counterbalanced by its ability to reach higher local optima, at least on the considered problems. Note that these results were obtained for a fixed number of evaluations, not

Table 10: Number/ratios of hill-climbing iterations performed by random-restart hill-climbers on UBQP landscapes. Allocated budget: 100 million evaluations.

| N | 128 | 256 | 512 | 1024 | 2048 | 4096 |
|--------|-------|-------|-------|------|------|------|
| #F | 91816 | 34479 | 12850 | 4762 | 1760 | 598 |
| #F/#W | 2.49 | 2.47 | 2.40 | 2.41 | 2.35 | 2.27 |
| #F/#ME | 4.55 | 4.73 | 4.88 | 4.97 | 5.07 | 5.08 |

for a similar execution time. This experimental design was considered to facilitate the experimentations (more runs were done simultaneously) and also to keep the same setting on the two problems. However, on some of these problems using incremental evaluation is possible for the neighborhoods under consideration. Incremental evaluation is particularly beneficial in the context of pivoting rules as the maximum expansion and its variants, as it significantly reduces the computational impact of an evaluation (its complexity). Indeed, in such cases performing more evaluations has a lowered impact on the overall complexity and duration of an algorithm. This would make ME more competitive in a hill-climbing context, and could potentially lead to an increase in the climbing processes within non-strict local search algorithms relative to other pivoting rules.

Most importantly, this study was done in a black-box context, for which problem specificities are not taken into account. While we believe this context was required for a first in-depth study, our findings on the competitiveness of the maximum expansion principle mostly highlight that using this principle efficiently is not an easy task. The non-conclusive aspect of it cannot be generalized to all settings, for example, the way of using the expansion principle or the problems under consideration. In the first case, the relative efficiency of pivoting rules, including the maximum expansion, has yet to be studied in other frameworks. Such frameworks include a hill-climber framework in which neutral moves are authorized, as well as other neighborhood-based metaheuristics. In the second case, we believe that the maximum expansion could be beneficial outside of the black-box context. It could be used efficiently either as an indicator to guide the search or as a pivoting rule, probably by integrating problem specificities.

6. Conclusion

This paper is part of an effort to improve knowledge about the dynamics of hill-climbing principles. Indeed, this work was devoted to exploring how to maintain a better balance between exploration and exploitation within strict hill-climbing algorithms. In particular, we proposed the maximum expansion pivoting rule which selects the improving neighbor that has the highest number of improving neighbors. Unlike classical pivoting rules, ME is not directly based on fitness values to rank solutions but uses the expansion score to determine trajectories that we expect to drive more likely toward higher local optima.

In the first part of this study, experiments conducted on landscapes derived from four combinatorial optimization problems highlighted that the expansion principle leads toward better local optima than classic pivoting rules while a second set of experiments highlighted that its efficiency is not solely explained by the larger amount of knowledge it uses to select a neighbor. The maximum expansion principle was shown to be a relevant criterion to guide the search toward better local optima, regardless of the ruggedness levels of landscapes.

The second part of the study aims to explore the limitations of an expansion-based hill-climbing. In particular, we discuss its high computation cost and propose a way of approximating this principle to achieve good local optima while requiring fewer evaluations. This approximation yields good tradeoffs between local optima quality and computational cost but still requires a higher number of fitness queries than other pivoting rules. At equivalent budget, the random-restart first improvement climber remains, in general, the most efficient iterated hill-climbing strategy.

However, we believe that these approximations of the maximum expansion are an interesting way to consider the expansion indicator while conceiving local search based solving techniques and could be used more efficiently. For example, the expansion principle could be incorporated directly in search algorithms in a context where problem specificities and/or incremental evaluation are taken into account. Another possibility is to use the expansion score as an indicator to improve the understanding of search process behavior. This could be done by investigating the basins of attraction of local optima, as well as exploring further the correlation between fitness and expansion score.

Acknowledgements

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] Aboutaib, B., Verel, S., Fonlupt, C., Derbel, B., Liefoghe, A., Ahiod, B., 2020. On stochastic fitness landscapes: Local optimality and fitness landscape analysis for stochastic search operators, in: International Conference on Parallel Problem Solving from Nature, Springer. pp. 97–110.

- [2] Aguirre, H., Tanaka, K., 2003. A study on the behavior of genetic algorithms on NK-landscapes: Effects of selection, drift, mutation, and recombination. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* 86, 2270–2279.
- [3] Al-Betar, M.A., 2017. β -hill climbing: an exploratory local search. *Neural Computing and Applications* 28, 153–168.
- [4] Al-Betar, M.A., Aljarah, I., Awadallah, M.A., Faris, H., Mirjalili, S., 2019. Adaptive β -hill climbing for optimization. *Soft Computing* 23, 13489–13512.
- [5] Angel, E., Zissimopoulos, V., 2001. On the landscape ruggedness of the quadratic assignment problem. *Theoretical computer science* 263, 159–172.
- [6] Basseur, M., Goëffon, A., 2013. Hill-climbing strategies on various landscapes: an empirical comparison, in: *Genetic and Evolutionary Computation Conference (GECCO)*, ACM, New York, NY, USA. pp. 479–486.
- [7] Basseur, M., Goëffon, A., 2014. On the efficiency of worst improvement for climbing NK-landscapes, in: *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, ACM, New York, NY, USA. pp. 413–420.
- [8] Basseur, M., Goëffon, A., 2015. Climbing combinatorial fitness landscapes. *Applied Soft Computing* 30, 688–704.
- [9] Basseur, M., Goëffon, A., Lardeux, F., Saubion, F., Vigneron, V., 2014. On the attainability of NK landscapes global optima, in: *Seventh Annual Symposium on Combinatorial Search*, pp. 28–34.
- [10] Bateson, W., Waunders, E.R., Punnett, R.C., 1909. Experimental studies in the physiology of heredity. *Molecular and General Genetics MGG* 2, 17–19.
- [11] Burke, E.K., Bykov, Y., 2017. The late acceptance hill-climbing heuristic. *European Journal of Operational Research* 258, 70–78.
- [12] Camacho-Villalón, C.L., Dorigo, M., Stützle, T., 2019. The intelligent water drops algorithm: why it cannot be considered a novel algorithm. *Swarm Intelligence* 13, 173–192.

- [13] Chen, T., Huang, G., Olanipekun, A.O., 2018. Simulating the evolution mechanism of inner innovation in large-scale construction enterprise with an improved NK model. *Sustainability* 10, 4221.
- [14] Chicano, F., Whitley, D., Ochoa, G., Tinós, R., 2017. Optimizing one million variable NK landscapes by hybridizing deterministic recombination and local search, in: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 753–760.
- [15] Del Ser, J., Osaba, E., Molina, D., Yang, X.S., Salcedo-Sanz, S., Camacho, D., Das, S., Suganthan, P.N., Coello, C.A.C., Herrera, F., 2019. Bio-inspired computation: Where we stand and what’s next. *Swarm and Evolutionary Computation* 48, 220–250.
- [16] Faggini, M., Parziale, A., et al., 2017. Fitness landscape and tax planning: NK model for fiscal federalism. *International Journal of Economics and Finance* 9, 14–31.
- [17] Fan, I.Y., Lee, R.W., 2012. Design of a weighted and informed NK model for intellectual capital-based innovation planning. *Expert Systems with Applications* 39, 9222–9229.
- [18] Fonlupt, C., Robilliard, D., Preux, P., Talbi, E.G., 1999. Fitness landscapes and performance of meta-heuristics, in: *Meta-Heuristics*. Springer, pp. 257–268.
- [19] Garey, M.R., Johnson, D.S., Sethi, R., 1976. The complexity of flowshop and jobshop scheduling. *Mathematics of operations research* 1, 117–129.
- [20] Gary, M.R., Johnson, D.S., 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. WH Freeman and Company, New York.
- [21] Kauffman, S.A., 1993. *The origins of order: Self organization and selection in evolution*. Oxford University Press, USA.
- [22] Kauffman, S.A., Weinberger, E.D., 1989. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* 141, 211 – 245.

- [23] Koopmans, T., Beckmann, M.J., 1955. Assignment Problems and the Location of Economic Activities. Cowles Foundation Discussion Papers 4. Cowles Foundation for Research in Economics, Yale University.
- [24] Kouyos, R.D., Leventhal, G.E., Hinkley, T., Haddad, M., Whitcomb, J.M., Petropoulos, C.J., Bonhoeffer, S., 2012. Exploring the complexity of the HIV-1 fitness landscape. *PLoS genetics* 8, e1002551.
- [25] Li, K., Liang, Z., Yang, S., Chen, Z., Wang, H., Lin, Z., 2019. Performance analyses of differential evolution algorithm based on dynamic fitness landscape. *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)* 13, 36–61.
- [26] Lourenço, H.R., Martin, O.C., Stützle, T., 2003. Iterated local search, in: *Handbook of metaheuristics*. Springer, pp. 320–353.
- [27] Malan, K., Ochoa, G., 2021. Recent advances in landscape analysis for optimisation and learning, in: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 899–917.
- [28] Malan, K.M., 2021. A survey of advances in landscape analysis for optimisation. *Algorithms* 14.
- [29] Malan, K.M., Engelbrecht, A.P., 2013. A survey of techniques for characterising fitness landscapes and some possible ways forward. *Information Sciences* 241, 148–163.
- [30] Merz, P., Freisleben, B., 2000. Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. *IEEE Transactions on Evolutionary Computation* 4, 337–352.
- [31] Mladenović, N., Hansen, P., 1997. Variable neighborhood search. *Computers & Operations Research* 24, 1097–1100.
- [32] Ochoa, G., Tomassini, M., Verel, S., Darabos, C., 2008. A study of NK landscapes’ basins and local optima networks, in: *Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pp. 555–562.
- [33] Ochoa, G., Verel, S., Tomassini, M., 2010. First-improvement vs. best-improvement local optima networks of NK landscapes, in: *International*

- Conference on Parallel Problem Solving from Nature, Springer. pp. 104–113.
- [34] Palubeckis, G., 2004. Multistart tabu search strategies for the unconstrained binary quadratic optimization problem. *Annals of Operations Research* 131, 259–282.
 - [35] Pelikan, M., 2010. NK landscapes, problem difficulty, and hybrid evolutionary algorithms, in: *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pp. 665–672.
 - [36] Piotrowski, A.P., Napiorkowski, J.J., Rowinski, P.M., 2014. How novel is the “novel” black hole optimization approach? *Information Sciences* 267, 191–200.
 - [37] Reidys, C.M., Stadler, P.F., 2001. Neutrality in fitness landscapes. *Applied Mathematics and Computation* 117, 321–350.
 - [38] Sala, R., Müller, R., 2020. Benchmarking for metaheuristic black-box optimization: perspectives and open challenges, in: *2020 IEEE Congress on Evolutionary Computation (CEC)*, IEEE. pp. 1–8.
 - [39] Sörensen, K., 2015. Metaheuristics — the metaphor exposed. *International Transactions in Operational Research* 22, 3–18.
 - [40] Stadler, P.F., 2002. Fitness landscapes, in: *Biological evolution and statistical physics*. Springer, Berlin, pp. 183–204.
 - [41] Taillard, E., 1993. Benchmarks for basic scheduling problems. *European Journal of Operational Research* 64, 278–285.
 - [42] Tari, S., Basseur, M., Goëffon, A., 2018. Worst improvement based iterated local search, in: *European Conference on Evolutionary Computation in Combinatorial Optimization*, Springer. pp. 50–66.
 - [43] Tari, S., Basseur, M., Goëffon, A., 2022. Partial neighborhood local searches *International Transactions in Operational Research* 29, 2761–2788.
 - [44] Verel, S., Collard, P., Clergue, M., 2004. Scuba search: when selection meets innovation, in: *Congress on Evolutionary Computation (CEC)*, IEEE. pp. 924–931.

- [45] Verel, S., Daolio, F., Ochoa, G., Tomassini, M., 2011. Local optima networks with escape edges, in: International Conference on Artificial Evolution (Evolution Artificielle), Springer. pp. 49–60.
- [46] Verel, S., Ochoa, G., Tomassini, M., 2010. Local optima networks of NK landscapes with neutrality. *IEEE Transactions on Evolutionary Computation* 15, 783–797.
- [47] Villalón, C.C., Stützle, T., Dorigo, M., 2022. An analysis of why cuckoo search does not bring any novel ideas to optimization. *Computers & Operations Research* 142.
- [48] Warren G. Jackson, E.Ö., John, R.I., 2018. Move acceptance in local search metaheuristics for cross-domain search. *Expert Systems with Applications* 109, 131–151.
- [49] Wright, S., 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution.
- [50] Yannakakis, M., 1990. The analysis of local search problems and their heuristics, in: Annual Symposium on Theoretical Aspects of Computer Science, Springer. pp. 298–311.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Sara Tari: Writing- Original draft preparation, Software, Investigation
All authors: Writing- Reviewing and Editing, Conceptualization,
Methodology, Validation, Formal Analysis

Highlights

- Climbers reach high-quality local optima when guided by the expansion score.
- The ruggedness levels affect the efficiency of classic selection strategies but not of the maximum expansion.
- Approximating the maximum expansion principle significantly reduces the computational cost but remain relevant to achieve good tradeoffs.