

On the Attainability of NK Landscapes Global Optima

Matthieu Basseur, Adrien Goëffon, Frédéric Lardeux, Frédéric Saubion and Vincent Vigneron

{firstname.lastname}@univ-angers.fr
LERIA, Université d'Angers, France

Abstract

In this paper, we aim at evaluating the impact of the starting point of a basic local search based on the first improvement strategy. We define the coverage rate of a configuration as the proportion of the search space from which a particular configuration can be reached by a strict hill-climbing with a non-zero probability. In particular, we compute the coverage rate of fitness landscapes global optima, in order to evaluate their attainability by hill-climbing algorithms. The experimental study is realized on NK landscapes, in which the size and ruggedness can be controlled.

Results indicate that the coverage rate of global optima is usually high, which means that a basic strictly improving hill-climbing with first improvement strategy is able to reach global optima, independently to the starting point considered. This confirms that it is more important to focus on an effective search strategy rather than worrying about the choice of the initial configurations.

Introduction

Hill-climbing (Selman and Gomes 2002), among fast local search techniques (Hoos and Stützle 2004), is often considered as an intensification mechanism which focuses on a tiny part of the search space, its efficiency being greatly dependent to the initial configuration considered. Globally, one can think that most configurations are not reachable since they belong to parts of the search space which are too far from the initial configuration. However, in most search landscapes, the neighborhood induces a reduced diameter of the associated transition graph¹. For instance, considering a bit-string landscape, the 1-flip neighborhood links configurations through an N -dimensional hypercube, N being the bit-string length. Obviously, neighborhood distances between pairs of configurations in such landscapes are $N/2$ in average — N in the (unique) extreme case — which is not particularly large in comparison with the number of moves usually performed during a local search.

In this work, we measure the impact of the initial solution of a basic first improvement local search. To achieve this,

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The transition graph simply consists of nodes that represent configurations and edges that link two neighbor configurations.

we aim at evaluating the coverage rate of particular configurations of fitness landscapes, which represents the proportion of the search space from which a configuration can be reached thanks to a hill-climbing process. This study can be viewed as an extension of the notion of basin of attraction (Ochoa et al. 2008; Ochoa, Vérel, and Tomassini 2010). In these previous works, authors focused on local optima networks and on the *probabilities* to reach given local optima from any initial configuration. Here, we propose to concentrate about *possibilities* to reach particular configurations — eg. the global optimum — from any configuration thanks to a first improvement hill-climbing. Experiments are realized on NK landscapes, in which size and ruggedness are tunable.

The main motivation of this paper is to provide a clear understanding of the real impact of well-known basic local search components in order to improve the design of more complex and more sophisticated algorithms, based on these intuitions. In particular, the management of the balance between intensification and diversification has been widely studied in local search, leading to algorithms that may even adapt their behavior according to the current state of the search (Battiti, Brunato, and Mascia 2008). Clearly, the design of such algorithm may benefit from more knowledge about, for instance, influence of the starting point in local search, in order to manage restart strategies, or probabilities of reaching local and global optima, in order to better use diversification by means of escape techniques (e.g., prohibition mechanisms).

The paper is organized as follow. The next section focuses on the definitions of fitness landscapes, hill-climbing algorithms and we also introduce the concept of coverage rate. The NK landscapes, which are considered in the experiments, are also presented. In section 3, we propose several experiments in order to evaluate coverage properties of NK landscapes. In the last section, we summarize our contributions and point out ways to exploit these results while designing metaheuristics.

Definitions

Fitness Landscapes and hill-climbing

A *combinatorial optimization problem instance* is a pair (X, f) , where X is a discrete set of *feasible solutions*, and $f : X \rightarrow \mathbb{R}$ a scalar *objective function* which has to be

maximized or minimized. In a maximization context, solving (X, f) consists in finding $x^* \in \operatorname{argmax}_{x \in X} f(x)$.

A *fitness landscape* is a triplet (X, \mathcal{N}, f) , where X is a set of configurations called *search space*, $\mathcal{N} : X \rightarrow 2^X$ is a *neighborhood function*, and f is a *fitness function*. $f(x)$ is the *fitness* of x . $x' \in \mathcal{N}(x)$ is a *neighbor* of a configuration $x \in X$. $\mathcal{N}(x)$ is the *neighborhood* of x . A *local optimum* is a configuration x such that $\forall x' \in \mathcal{N}(x), f(x') \leq f(x)$. A *global optimum* is a configuration $x^* \in \operatorname{argmax}_{x \in X} f(x)$.

A *local search algorithm* aims at finding the best configuration of X (thanks to f) while exploring a subset of X relatively to (X, \mathcal{N}, f) . A *hill-climbing* algorithm is a basic local search strategy which navigates through the search space in allowing only non-deteriorating moves. Given an initial configuration called *starting point*, a basic hill-climbing algorithm iteratively moves to better neighbors, until it reaches a local optimum. We consider here the most stochastic basic hill-climbing strategy, *first improvement*, which accepts at each iteration the first evaluated neighbor which satisfies the moving condition (see algorithm 1). Other well-known strategy, *best improvement*, is not considered here, since its (globally) deterministic aspect will not distinguish possibilistic and probabilistic cases. Moreover, in (Basseur and Goëffon 2013), it has been shown that first improvement often outperforms best improvement on a wide variety of landscapes.

input : a fitness landscape (X, \mathcal{N}, f) , a starting point $x_0 \in X$

output: a local optimum x_{opt} and its fitness F_{opt}

$F \leftarrow f(x_0)$

$N \leftarrow \mathcal{N}(x_0)$

while $N \neq \emptyset$ **do**

 Randomly selects $x' \in N$

$F' \leftarrow f(x')$

if $F' > F$ **then**

$x \leftarrow x'$

$F \leftarrow F'$

$N \leftarrow \mathcal{N}(x)$

end

else

$N \leftarrow N \setminus \{x'\}$

end

end

$x_{\text{opt}} \leftarrow x$

$F_{\text{opt}} \leftarrow F$

return $(x_{\text{opt}}, F_{\text{opt}})$

Algorithm 1: Basic first improvement hill-climbing.

More details about hill-climbing and various stochastic local search techniques can be found in (Hoos and Stützle 2004).

Basins of attraction and coverage rate

Given a fitness landscape (X, \mathcal{N}, f) and a hill-climbing strategy (first or best improvement), the *basin of attraction* of a local optimum x_{opt} is the set $B_{x_{\text{opt}}} = \{x \in$

$X, p_{x_{\text{opt}}}(x) > 0\}$, where $p_{x_{\text{opt}}}(x)$ represents the probability to reach x_{opt} from x as starting point and using the given hill-climbing strategy (Ochoa, Vérel, and Tomassini 2010). Note that this definition is restricted to local optima only but can actually be generalized to any configuration of the search space.

In (Ochoa, Vérel, and Tomassini 2010), the size of a basin of attraction $B_{x_{\text{opt}}}$ is defined as the sum of probabilities $\sum_{x \in X} p_{x_{\text{opt}}}(x)$. Authors do not define the size of a basin of attraction as its cardinality, in order to represent the probability for a given local optimum to be attained by a hill-climbing, which starts from a random configuration (the global probability is actually $\frac{\text{size}(B_{x_{\text{opt}}})}{\#X}$). In this paper, we will focus on the cardinality of basins of attraction in order to determine the *possibility* of reaching a given local optimum. Thus, we introduce the *coverage rate* $\gamma(x)$ of a configuration $x \in X$ as the proportion of the search space that belongs to its basin of attraction:

$$\gamma(x) = \frac{\#B_x}{\#X}$$

NK Landscapes

The NK family of landscapes (Kauffman and Weinberger 1989) is a problem-independent model for constructing multimodal landscapes. NK-landscapes use a basic search space, with binary strings as configurations and bit-flip as neighborhood (two configurations are neighbors iff their Hamming distance is 1). Characteristics of an NK landscape are determined by two parameters N and K . N refers to the size of binary string configurations, which defines the search space size ($|X| = 2^N$). K specifies the ruggedness level of the landscape; indeed, the fitness value of a configuration is given by the sum of N terms, each one depending on $K + 1$ bits of the configuration. Thus, by increasing the value of K from 0 to $N - 1$, NK landscapes can be tuned from smooth to rugged.

In NK landscapes, the fitness function $f : \{0, 1\}^N \rightarrow [0, 1]$ to be maximized is defined as follows.

$$f(x) = \frac{1}{N} \sum_{i=1}^N c_i(x_i, x_{i_1}, \dots, x_{i_K}) \quad (1)$$

where $c_i : \{0, 1\}^{K+1} \rightarrow [0, 1]$ defines the component function associated with each variable x_i , $i \in \{1, \dots, N\}$, and where $K < N$.

NK landscapes instances are determined by the $(K + 1)$ -uples $(x_i, x_{i_1}, \dots, x_{i_K})$ as well as the $2^N(K + 1)$ c_i result values corresponding to a fitness-contribution matrix C whose values are randomly generated in $[0, 1]$ and following the random neighborhood model (Kauffman and Weinberger 1989).

In the following, we study NK landscapes to extract information about coverage rates. Let us recall that a coverage rate is defined with respect to a landscape and a move strategy. In this paper, we focus on the first improvement strategy.

Experiments will be carried out on various landscapes classified with respect to their size and ruggedness. We

first focus on studying coverage rate properties of NK -landscapes on small instances in order to provide exact results, which will help us to extrapolate such properties on larger landscapes.

Coverage rate of NK landscape global optima

This section aims at calculating the coverage rate of NK landscape global optima by performing a first improvement hill-climbing. A high coverage rate means that an effort for determining an *appropriate* initial solution is not necessarily useful. Indeed, this would mean that a randomly chosen solution would have a high probability of having a chance to reach the global optimum.

Methodology

The coverage rate computation of a global optimum first requires to solve the optimization problem itself. For each NK landscape instance considered in the results section, the global optimum is obtained by means of an exhaustive search. To compute exactly the coverage rate of the global optimum, we determine all possible reverse hill-climbing paths from the optimum as starting point.

```

input : a fitness landscape  $(X, \mathcal{N}, f)$  and its global
         optimum  $s^*$ 
output: the coverage rate of  $s^*$ 
cardcov  $\leftarrow 1$ 
 $\tilde{f} \leftarrow \lfloor f(s^*).H \rfloor$  //  $\tilde{f}$  is the integer transformation of  $f$  in
the interval  $\{0..H-1\}$ , where  $H$  is the hash table size
InsertSol (HashTable,  $\tilde{f}, s^*$ )
for  $i \leftarrow \tilde{f}$  to 0 do
  foreach  $s \in \text{SolSet}(\text{HashTable}, i)$  do
    foreach  $s' \in \mathcal{N}(s)$  do
       $j \leftarrow \lfloor f(s').H \rfloor$  // a solution is inserted if
the fitness is deteriorated and if this solution
is not already recorded;
      if  $j < i$  then
        if not FindSol (HashTable,  $j, s'$ ) then
          InsertSol (HashTable,  $j, s'$ )
          cardcov  $\leftarrow$  cardcov + 1
        end
      end
    end
  end
  DeleteSolSet (HashTable,  $i$ )
end
return cardcov/ $2^N$ ;

```

Algorithm 2: Computation of a coverage rate.

A basic approach is to recursively generate deteriorating neighbors until reaching a local minimum, while counting the number of distinct configuration considered. However such an approach is intractable since most solutions belong to numerous hill-climbing paths. The use of a hash table which considers fitness as hash function will allow us to effectively ensure that each solution is handled only

once. Combined with a treatment of solutions by decreasing fitness, this representation allows to free memory during the search, by deleting solutions which cannot appear afterwards (see algorithm 2).

Results

Table 1 presents the coverage rates obtained on small NK landscapes instances. By focusing on the column LO_1 , which corresponds to the coverage rates of global optima, one may observe that the average rate is greater than 70% for all (N, K) parameterizations. The smallest coverage result (70.17%) is obtained on a rugged instance ($K = 8$), which can be classified as difficult. Nevertheless, the average global optima coverage rate on rugged NK landscapes remains relatively high. Considering average rates according to the landscapes ruggedness, we observe that the greater values are obtained for $K = 4$. This result is quite surprising since one would expect to get the best coverage for the least rugged instances ($K = 2$), even if standard deviations are higher.

An interesting point is that global optima coverage rates are high and their variation between NK instances are relatively low. It means that global optima can be attained with basic hill-climbings from most initial configurations. This result seems to be not dependent to the size of the search space, and we can expect that this remains valid for wider landscapes.

During experiments, we observed the distribution of solutions according to their fitness, considering if they belong or not in the coverage of the global optimum. By this way, we are able to visualize the correlation between coverage rates and configuration fitnesses, as shown in figure 1. This figure emphasizes that the global optimum is almost always attainable from the worst configurations of the search space, which is intuitively understandable since it allows the generation of many possible improving paths. At the contrary, it is less often possible to reach the global optimum when starting from a configuration with a high fitness.

Let us notice that a high coverage does not mean that the global optimum is easily attainable, but refers to the possibility to reach it. In the next section, we focus on the probability of reaching global optima from randomly selected solutions.

Probability of reaching global and local optima

Previous results show that the global optimum is achievable with a hill-climbing from most solutions of the search space. However, this does implies a significant probability to reach it. This section aims at computing the probabilities to reach each local optimum, ie. the size of their basin of attraction.

Methodology

The computation of the probability of achieving a particular configuration (eg. the global optimum) cannot be achieved by exhaustively testing all possible paths from all possible starting points. Here, we propose to compute the probability to reach a configuration from a single random starting point. An estimation of the expected probability will be obtained

Instance	$N = 16$			$N = 24$		
	$K = 2$	$K = 4$	$K = 8$	$K = 2$	$K = 4$	$K = 8$
1	94.37%	91.49%	68.38%	63.19%	95.95%	88.89%
2	55.72%	86.68%	74.82%	91.93%	96.12%	88.90%
3	90.95%	86.90%	74.26%	66.19%	95.50%	91.64%
4	67.83%	92.95%	66.99%	71.55%	95.83%	88.70%
5	54.65%	88.79%	69.38%	92.42%	94.99%	90.60%
6	74.98%	87.76%	68.52%	72.44%	97.66%	86.64%
7	68.91%	87.81%	69.40%	63.63%	94.34%	87.61%
8	74.03%	87.19%	66.75%	96.92%	96.58%	89.60%
9	96.78%	91.07%	71.34%	80.10%	92.98%	90.55%
10	78.67%	87.83%	71.86%	82.39%	96.13%	91.00%
Avg	75.69%	88.85%	70.17%	78.08%	95.61%	89.41%
SD	14.86%	2.19%	2.82%	12.57%	1.28%	1.57%
Global	Avg: 78.23%; SD: 11.66%			Avg: 87.70%; SD: 10.24%		

Table 1: Coverage rates of randomly generated NK landscapes ($N \in \{16, 24\}$ and $K \in \{2, 4, 8\}$).

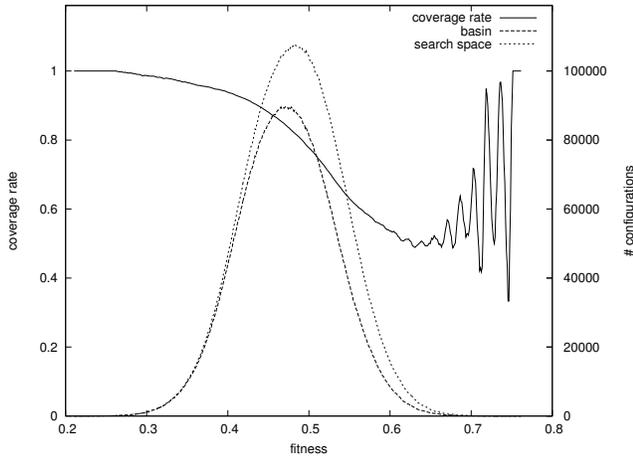


Figure 1: Coverage rate of an NK landscape global optimum ($N = 24$, $K = 2$, instance #9, global cov. rate: 80.10%) with respect to fitness. We indicate, for all fitness values (range size 10^{-3}), the number of configurations belonging to the basin of attraction of the global optimum (dashed line), and the total number of configurations of the search space (dotted line).

by averaging on a sample of starting points. The algorithm works as follow.

Computing the probability that a configuration x_0 reaches an optimum x_{opt} requires to determine the probability of selection of all improving paths that drive to x_{opt} . In particular, if no search path leads to x_{opt} , the probability will be set to 0. The probability of following a path will be evaluated starting from x_0 and taking into account all possible local searches ways. In a first improvement context, we consider that all improving neighbors have an equivalent probability to be selected. Thus, if a configuration x has a probability p to be reached and has k improving neighbors, then the

probability of choosing each of these paths will be p/k . If a solution can be reached from two distinct paths, then we add the corresponding probabilities.

Algorithm 3 describes how the probability of achieving an optimum x_{opt} starting from a solution x_0 is computed.

input : a fitness landscape (X, \mathcal{N}, f) , a starting point $s_0 \in X$ and a target configuration $s^* \in X$
output: the probability to reach s^* from s_0

```

 $\tilde{f}_0 \leftarrow \lfloor f(s_0) \cdot H \rfloor$ 
 $\tilde{f}^* \leftarrow \lfloor f(s^*) \cdot H \rfloor$ 
InsertSol (HashTable,  $\tilde{f}, s_0$ )
Pr[ $s_0$ ]  $\leftarrow 1$  // Probability to reach  $s_0$  from  $s_0$ 
for  $\tilde{f} \leftarrow \tilde{f}_0$  to  $\tilde{f}^*$  do
  foreach  $s \in \text{SolSet}(\text{HashTable}, \tilde{f})$  do
    foreach  $s' \in \mathcal{N}(s)$  do
       $\tilde{f}' \leftarrow \lfloor f(s') \cdot H \rfloor$ 
      if  $\tilde{f}' > \tilde{f}$  then
        if not FindSol (HashTable,  $\tilde{f}', s'$ ) then
          InsertSol (HashTable,  $\tilde{f}', s'$ )
          Pr[ $s'$ ]  $\leftarrow$  Pr[ $s$ ]/ $n$ 
        end
      else
        Pr[ $s'$ ]  $\leftarrow$  Pr[ $s'$ ] + Pr[ $s$ ]/ $n$ 
      end
    end
  end
  DeleteSolSet (HashTable,  $\tilde{f}$ )
end
return Pr[ $s^*$ ]

```

Algorithm 3: Probability to reach a solution s^* .

Instance	Best	10.00 %	30.00 %	50.00 %	Top 5	#LO	Cov.
1	9.54 %	14.25 %	49.97 %	61.79 %	43.76 %	20	94.37 %
2	2.49 %	10.73 %	33.05 %	59.86 %	17.23 %	36	55.72 %
3	13.43 %	13.43 %	44.98 %	84.58 %	66.86 %	16	90.94 %
4	5.04 %	16.67 %	49.02 %	77.65 %	21.58 %	49	67.82 %
5	7.89 %	17.63 %	55.31 %	75.03 %	17.63 %	54	54.65 %
6	8.14 %	24.76 %	55.45 %	76.76 %	28.17 %	48	74.98 %
7	9.46 %	28.54 %	67.36 %	85.14 %	20.41 %	81	68.91 %
8	12.09 %	12.09 %	45.20 %	73.59 %	51.62 %	16	74.02 %
9	34.89 %	34.89 %	75.04 %	84.90 %	75.04 %	18	96.77 %
10	45.43 %	45.43 %	64.49 %	79.76 %	64.49 %	19	78.66 %
Avg	14.84%	21.84%	53.99%	75.91%	65.70	35.7	75.69
SD	13.93 %	11.41 %	12.35 %	8.96 %	22.55 %	21.93	14.86

Table 2: Probability of achieving different local optima on ($N = 16, K = 2$) instances. Probabilities are obtained from a sampling of 100 initial solutions. The probability of reaching each possible local optimum was calculated. The *Best* column corresponds to the probability of reaching the global optimum. Columns 10%, 30%, 50% refer to the probability of reaching respectively the 10%, 30%, 50% best local optima. Top 5 column corresponds to the probability to reach one of the 5 best local optima. Column #LO indicates the number of local optima of the considered instance. The last column reports the coverage rate of the global optimum on each considered instance.

Instance	Best	10.00 %	30.00 %	50.00 %	Top 5	#LO	Cov.
1	3.69 %	35.41 %	66.62 %	84.05 %	24.17 %	121	91.48
2	4.23 %	27.15 %	61.34 %	82.95 %	13.70 %	129	86.68
3	3.29 %	36.29 %	71.83 %	85.27 %	28.76 %	89	86.89
4	3.64 %	24.92 %	55.68 %	78.67 %	15.09 %	106	92.94
5	6.02 %	32.18 %	70.04 %	85.72 %	20.39 %	99	88.79
6	3.81 %	26.13 %	64.49 %	82.06 %	16.45 %	119	87.76
7	4.56 %	35.03 %	65.15 %	81.66 %	14.74 %	140	87.81
8	3.28 %	29.37 %	61.81 %	80.22 %	14.65 %	139	87.18
9	3.89 %	37.92 %	69.48 %	83.50 %	17.59 %	121	91.06
10	9.28 %	37.87 %	66.56 %	83.12 %	18.45 %	142	87.83
Avg	4.57%	32.23%	65.30%	82.72%	18.40%	120.5	88.85
SD	1.84 %	4.98 %	4.79 %	2.17 %	4.83 %	18.03	2.19

Table 3: Probability of achieving different local optimum on ($N = 16, K = 4$) instances. See Table 2 for columns description.

Instance	Best	10.00 %	30.00 %	50.00 %	Top 5	#LO	Cov.
1	0.54 %	26.93 %	59.95 %	77.52 %	2.92 %	770	68.37
2	2.76 %	31.45 %	59.93 %	79.53 %	5.18 %	770	74.81
3	0.57 %	27.59 %	57.67 %	75.97 %	2.47 %	730	74.26
4	0.61 %	30.21 %	60.84 %	77.41 %	3.46 %	745	66.98
5	0.20 %	27.32 %	59.75 %	77.76 %	1.10 %	794	69.38
6	0.90 %	25.70 %	59.16 %	79.13 %	2.51 %	747	68.51
7	1.11 %	32.36 %	57.71 %	78.28 %	3.20 %	780	69.39
8	0.56 %	30.33 %	60.59 %	77.94 %	3.50 %	737	66.74
9	1.60 %	31.93 %	63.19 %	79.00 %	5.00 %	740	71.33
10	0.59 %	31.89 %	58.65 %	78.15 %	3.06 %	741	71.85
Avg	0.94%	29.57%	59.74%	78.07%	3.24%	755.4	70.17
SD	0.74 %	2.45 %	1.63 %	1.02 %	1.19 %	21.42	2.82

Table 4: Probability of achieving different local optima on ($N = 16, K = 8$) instances. See Table 2 for columns description.

Results

Tables 2, 3 and 4 report the probability of reaching local optima on NK landscapes, for $N = 16$. The average probability to reach the global optimum (best) is relatively low when $K = 4$ and $K = 8$ (respectively 4.5% and 1%). As expected, global optima’s high coverage rates do not imply high reaching probabilities. For $K = 2$ landscapes, this probability is relatively high (more than 10%). Globally, the probability of reaching the global optimum decreases while the ruggedness of the problem increases, which is natural since the number of local optima is greater. Nevertheless, let us recall that this phenomenon is not present for the covering rate. Since the probability of reaching the global optimum is a numerical indicator of the landscape difficulty, this experimentation confirms that the ruggedness level of a landscape directly affect its difficulty.

The column 50% provides us an additional piece of information. Indeed, for all ruggedness levels, the probability to reach one of the top 50% local optima is over 70%. This means that a first improvement hill-climbing globally tends to converge to high optima of the landscape. Similar conclusion were obtained by Ochoa *et al.* (Ochoa *et al.* 2008) which were focusing on best improvement. They showed that high local optima benefited from wider basins of attraction than lower local optima.

Correlation coverage rate vs. neighborhood distance

This section aims at evaluating the coverage rate of a configuration according to its neighborhood distance to other configurations. To achieve this, we distinguish different distance-based coverage rate values of a global optimum x_{opt} , depending of the neighborhood distance of configurations to x_{opt} . While a correlation between coverage rate and neighborhood distance is expected, it will be interesting to determine if coverage rates are still significant from distant configurations.

Here we will consider the distance-based coverage rates of global optima only. Let us notice that considering NK landscapes, neighborhood distances between solutions are distances, varying from 1 to N . The number of configurations being at a Hamming distance of d of any particular configuration is $\binom{N}{d}$. To compute distance-based coverage rates, we first split the basin of attraction of x_{opt} into N subsets with respect to the Hamming distance $h : X^2 \rightarrow \mathbb{N}$ between configurations of $B_{x_{opt}}$ and x_{opt} , and then compute the rates separately. More precisely, the coverage rate of x_{opt} at distance d is:

$$\frac{\#\{x \in B_{x_{opt}}, h(x, x_{opt}) = d\}}{\binom{N}{d}}$$

Let us precise that the computation of distance-based coverage rates can be included in algorithm 2 by considering N basin cardinalities rather than a single one.

Figure 2 presents an interesting view of the basin of attraction of an NK landscape global optimum x_{opt} ($N = 16$,

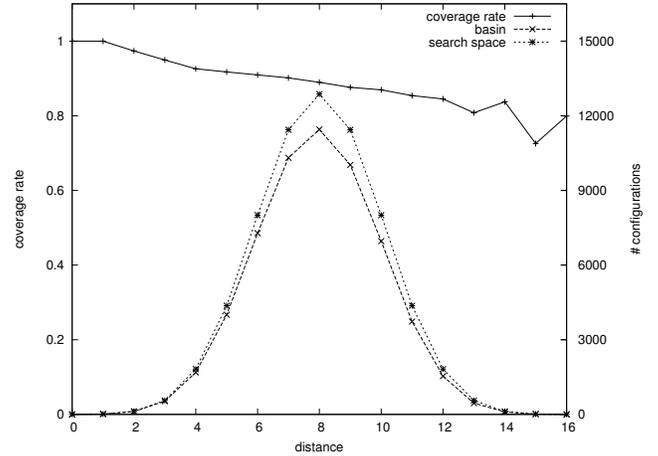


Figure 2: Average coverage rates of NK landscape global optima ($N = 16$, $K = 4$, 10 instances, global coverage rate: 88.85%) with respect to distances from optimum. We indicate, for all distance values, the average number of configurations belonging to the basin of attraction of each instance global optimum (dashed line), and the total number of configurations of a search space (dotted line).

$K = 4$). We illustrate here the proportion of configurations belonging to the basin of attraction of x_{opt} with respect to their neighborhood (Hamming) distance of x_{opt} . One can observe that although the distance-based coverage rate tends to decrease linearly while the neighborhood distance increases, it remains relatively high even for farrest solutions.

Correlation neighborhood distance / search length

In this section, we aim at comparing the neighborhood (Hamming) distance between starting points and global optimum with the minimal length of a hill-climbing path linking these two configurations. A strong correlation could allow to reduce the computational effort to solve large instances by focusing on direct (or nearly direct) paths — ie. without flipping the same bit several times.

The correlation between Hamming distance and search length is performed on the whole search space. We focus here on computing this correlation with the global optimum as target configuration. The principle is similar to the coverage computation provided in section . The set of configurations belonging in the global optimum basin of attraction is computed thanks to algorithm 2, except that we compute minimal paths leading to the global optimum. In most cases, if there is a hill-climbing path to reach the global optimum, then numerous different search paths can generally be followed.

Results on NK instances (with $N \in \{16, 24\}$) are reported in table 5. From most configurations belonging to the global optimum basin of attraction, there exists a direct

N	K	Dist.	d	$d + 2$	$d + 4$	$d + 6$	$d + 8$
16	2	Avg	91.73%	6.27%	1.39%	0.40%	0.19%
		SD	0.042	0.029	0.008	0.004	0.002
	4	Avg	89.96%	7.40%	1.81%	0.57%	0.18%
		SD	0.029	0.016	0.008	0.004	0.001
	8	Avg	81.99%	12.76%	3.63%	1.15%	0.35%
		SD	0.023	0.011	0.007	0.004	0.002
24	2	Avg	88.16%	7.63%	2.38%	0.99%	0.46%
		SD	0.059	0.033	0.015	0.007	0.004
	4	Avg	91.53%	6.26%	1.41%	0.49%	0.19%
		SD	0.015	0.010	0.003	0.001	0.001
	8	Avg	87.73%	8.21%	2.46%	0.95%	0.40%
		SD	0.013	0.007	0.003	0.002	0.001
Cumulative avg.			88.52%	96.61%	98.79%	99.55%	99.85%

Table 5: Minimal hill-climbing path length with respect to the initial Hamming distance starting points / global optimum (d). Considered starting points are configurations belonging to the global optimum basin of attraction. Rates are averaged from the 10 considered instances of each (N, K) parameterization.

hill-climbing path reaching the global optimum in a number of moves corresponding to their neighborhood distance from the optimum (d). Cumulative results indicate that only 3.4% of configurations of the global optimum basin cannot reach the optimum in hill-climbing paths of $d + 2$ moves — ie. only one bit is flipped back to its initial value during the search. These observations indicate that if we focus on local searches which avoid or limit the application of a move twice during the search (e.g., using prohibition mechanisms), this do not influence significantly the possibility to reach the global optimum.

Conclusion

In this paper we aimed at extracting information from NK landscapes concerning the ability of hill-climbers to reach global optima. To achieve this, we performed an analysis through exhaustive explorations of NK landscapes. The different experiments realized provide us some interesting conclusions. A first experiment, which aims at evaluating the attainability of global optima, shows that they are reachable from most configurations of the search space through a basic hill-climbing process.

While focusing on the probability to reach the different local optima of the search space, results showed that a first improvement reaches more likely higher local optima than others.

Additional studies extended attainability results by adding neighborhood distance information. In particular, we pointed out that even furthest configurations are often able to reach the global optimum thanks to a strict hill-climbing. Moreover, by reducing the number of allowed moves, results indicate that when a solution belongs in the global optimum basin, in most cases there exists a direct path to attain it. It means that there often exists a hill-climbing path from an initial solution to the global optimum where the number of steps is equal or close to the initial neighborhood distance between both configurations.

It would be interesting to find a way to estimate the pro-

posed metrics for large of highly rugged landscapes where we are not able to enumerate all paths exhaustively. Nevertheless, this study lead us to believe that there is possibilities to improve hill-climbing search efficiency by defining alternative pivoting rules (Basseur and Goëffon 2014), which aim to increase the probability to reach the best local optima and then the average expected final fitness achieved.

References

- Basseur, M., and Goëffon, A. 2013. Hill-climbing strategies on various landscapes: an empirical comparison. In Blum, C., and Alba, E., eds., *GECCO*, 479–486. ACM.
- Basseur, M., and Goëffon, A. 2014. On the efficiency of worst improvement for climbing NK landscapes. In *GECCO*. ACM.
- Battiti, R.; Brunato, M.; and Mascia, F. 2008. *Reactive Search and Intelligent Optimization*, volume 45 of *Operations research/Computer Science Interfaces*. Springer Verlag.
- Hoos, H. H., and Stützle, T. 2004. *Stochastic Local Search: Foundations & Applications*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Kauffman, S. A., and 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(2):211–245.
- Ochoa, G.; Tomassini, M.; Vérel, S.; and Darabos, C. 2008. A study of NK landscapes’ basins and local optima networks. In Ryan, C., and Keijzer, M., eds., *GECCO*, 555–562. ACM.
- Ochoa, G.; Vérel, S.; and Tomassini, M. 2010. First-improvement vs. best-improvement local optima networks of nk landscapes. In Schaefer, R.; Cotta, C.; Kolodziej, J.; and Rudolph, G., eds., *PPSN (I)*, volume 6238 of *Lecture Notes in Computer Science*, 104–113. Springer.
- Selman, B., and Gomes, C. 2002. Hill climbing search. Nature Encyclopedia of Cognition, Nature Publ.