



Where are Bottlenecks in NK Fitness Landscapes?

Sébastien Verel, Philippe Collard, Manuel Clergue

► To cite this version:

Sébastien Verel, Philippe Collard, Manuel Clergue. Where are Bottlenecks in NK Fitness Landscapes?. Evolutionary Computation, 2003. CEC'03, Dec 2003, Canberra, Australia. pp.273–280, 10.1109/CEC.2003.1299585 . hal-00159994

HAL Id: hal-00159994

<https://hal.science/hal-00159994>

Submitted on 4 Jul 2007

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.

Where are Bottlenecks in NK Fitness Landscapes?

Sébastien Verel, Philippe Collard and Manuel Clergue

Université de Nice-Sophia Antipolis,
06640, Laboratoire I3S, France,
{verel,pc,clerguem}@i3s.unice.fr

Abstract- Usually the offspring-parent fitness correlation is used to visualize and analyze some characteristics of fitness landscapes such as evolvability. In this paper, we introduce a more general representation of this correlation, the *Fitness Cloud (FC)*. We use the bottleneck metaphor to emphasise fitness levels in landscape that cause local search process to slow down. For a local search heuristic such as hill-climbing or simulated annealing, *FC* allows to visualize bottleneck and neutrality of landscapes. To confirm the relevance of the *FC* representation we show where the bottlenecks are in the well-know NK fitness landscape and also how to use neutrality information from the *FC* to combine some neutral operator with local search heuristic.

Introduction

The fitness correlation between parent and offspring is often used to analyse search space. In this paper, we present the *Fitness Cloud (FC)* which is the scatterplot of points parent-fitness/offspring-fitness. The *FC* allowed us to visualize and analyse the dynamic of local search heuristic at fitness level. *FC* shows evolvability as well neutrality and fitness bottleneck. The bottleneck value is the fitness level that cause local search to slow down and stop. In other word, bottleneck is fitness value where heuristic converge. In this paper, we focus on the *NK* fitness landscapes. First, we present *NK* fitness landscapes and the definition of Fitness Cloud. Section 2 reveals how *FC* exhibit bottleneck in *NK* landscapes for two well-know search heuristics hill climber and simulated annealing. The *FC* represents the neutrality of landscape, section 3 proposes to use this information to design and analyse performances of strategy using neutrality.

1 NK Landscapes and Fitness Cloud

1.1 The Tunable NK-Fitness Landscapes

In this section the basic features of the family of NK fitness landscapes are reviewed. The notion of fitness landscapes [12] as search space is defined as follows: a set of potential solutions (genotypes), a fitness function that evaluates the genotypes and a topology that represents relations between genotypes. NK model proposed by Kauffman [4] is designed to capture the structure of rugged multi-peaked fitness landscapes. These random landscapes are defined on binary strings of length N . The parameter K represents the

number of *epistatic* links¹. By tuning K , landscapes can be generated with varying degrees of ruggedness. In order to compute the overall fitness of one string, one considers that each bit contributes a component to the total fitness based on its own value and the values of K other genes. The random model is used, where fitness contribution of one bit depends on its own value and K other randomly chosen bits². Fitness contributions come from a uniform distribution ranging from 0.0 to 1.0. Fitness of a string is computed as the sum of bits contribution at all N loci divided by N for normalization to the range $[0; 1]$. The case $K = 0$ corresponds to problem without epistasis: fitnesses of neighbourhood points are correlated. There exists a single optimum. A hill climbing search allows to reach this optimum and adaptive walks are thus relatively long ($N/2$). The case $K = N - 1$ corresponds to the maximum number of interaction of the parts. The fitness of any point is random. There exists an enormous number of local optima. Adaptive walks are relatively short ($\log(2N)$) and are likely to end up in local optimum.

1.2 The Offspring-Parent Fitness Correlation Cloud

Plotting fitness against some features is not a new idea. B. Manderick et al. [7] study the correlation coefficient of genetic operators: they compute the correlation between the fitnesses of a number of parents and the fitnesses of their offspring. J. Grefenstette [3] uses fitness distribution of genetic operators to predict GA behaviour. H. Rosé et al. [8] develop the *density of states* approach by plotting the number of genotypes with a same fitness value. Smith et al. [9] focus on notions of *evolvability* and *neutrality*; they plot the average fitness of offspring over fitness according to Hamming neighbourhood. *Evolvability* refers to the efficiency of evolutionary search. It is defined by Altenberg as "the ability of an operator/representation scheme to produce offspring that are fitter than their parents" ([1], [2]).

Fitness Cloud In order to get a visual rendering of evolvability, we proposed a more general representation in the plan parent-fitness / offspring-fitness. We consider that two strings are neighbours if there is a transformation related to

¹Epistasis is defined as the influence of the genotype at one locus on the effect of a mutation at another locus

²Weinberger [11] proved the NK optimization problem with random neighbourhoods is NP complete for $K \geq 3$

a local search heuristics or an operator, which allows "to pass" from one string to the other one. For each string x in the genotype space³, we plot one point which have for abscissa the fitness $f(x)$ of x and for ordinate the fitness $\tilde{f}(x)$ of a peculiar neighbour of x . Thus, we obtain a scatterplot, the so-called *offspring-parent Fitness Cloud (FC)*. The choice of one peculiar neighbour among all the possible ones is a feature of the heuristic. Implicitly the fitness cloud gives some insight on the genotype to phenotype map. The set of genotypes that all have equal fitness is a *neutral set* [5]. Such a set corresponds to one abscissa in the *FC*; according to this abscissa, a vertical slice from the cloud represents the fitness values that could be reached from this set of neutrality. For a given offspring-fitness value \tilde{f} , an horizontal slice represents all the fitness values from which a local operator can reach \tilde{f} . Evolvability against a fitness level can be characterized by the repartition of points over the diagonal line in the *FC*. In this paper we will use the *FC* to track the dynamic and to locate the bottlenecks of local search heuristic. To get a more synthetic view on the *FC*, we define three functions :

$$\tilde{f}_{min}(\varphi) = \min_{x \in G_\varphi} \tilde{f}(x)$$

$$\tilde{f}_{max}(\varphi) = \max_{x \in G_\varphi} \tilde{f}(x)$$

$$\tilde{f}_{mean}(\varphi) = \text{mean}_{x \in G_\varphi} \tilde{f}(x)$$

where G_φ is the neutral set defined by: $\{x \in Gtype \mid f(x) = \varphi\}$. Practically two fitness values are taken as equal if they both stand in the same interval⁴. We call FC_{min} , FC_{max} and FC_{mean} respectively the representative curve of \tilde{f}_{min} , \tilde{f}_{max} and \tilde{f}_{mean} . Plotting the curves FC_{max} and FC_{min} allows to materialize the edge of the cloud (see fig. 1). A peculiar fitness value β is defined as solution of equation $\tilde{f}_{mean}(\beta) = \beta$; it corresponds to the abscissa of intersection between the FC_{mean} curve and the diagonal line (see fig. 1). On average, G_β is invariant by the heuristic, i.e. the heuristic is neutral on the set G_β .

Offspring-Parent Fitness Cloud and Hamming Distance

Before study a specific heuristic, it may be useful to get a view on the scatterplot fitness vs. fitness based on Hamming neighbourhood which is independent from any heuristic. So, we plot a cloud where all of the genotypes that can be produced by a single bit flip are selected (see fig. 1). Hence, the entire neighbourhood is represented without condition induced by some heuristics.




Figure 1: Fitness Cloud from the Hamming neighbourhood: border (min and max) and mean (NK-landscapes with $N = 25$, $K = 20$ and $K = 5$)




Figure 2: Distribution of points on the Fitness Cloud (NK landscape with $N = 25$ and $K = 20$)

Figure 1 shows the FC_{mean} curve is coarsely a line. This illustrates a well-known result: Weinberger [10] established the following relation between the mean offspring-fitness and the length N , the epistatic parameter K and the fitness value f :

$$\tilde{f}_{mean}(f) = \left(1 - \frac{K+1}{N}\right) f + \left(\frac{K+1}{N}\right) \beta \quad (1)$$

where β is constant. Therefore the mean offspring-fitness depends linearly from the parent-fitness whatever the epistatic parameter K is. As reported by Smith [9], let's note that the slop coefficient $1 - \frac{K+1}{N}$ is the offspring-parent fitnesses correlation [10]. The β fitness level is always equal to 0.5. So, when the parameter K varies from 0 to $N-1$, the FC_{mean} line turns around the $(\beta; \beta)$ point (see

³Data are collected from an exhaustive enumeration of the search space

⁴in our experiments the range is 0.002

fig. 1). For $K = 0$ the problem is linear and the FC_{mean} line is near the diagonal; at the opposite when epistasis is upper limit ($K = N - 1$), the FC_{mean} line is close to the horizontal.

2 Fitness Bottleneck and Limit Fitness Cloud

In this section we show that the fitness cloud is useful in identifying a bottleneck's location at fitness level. In a first step we plot the FC according to a local search heuristic. So we are able to locate the bottleneck. Then modeling dynamic by the way of the *limit fitness cloud* allows to confirm the realness of bottleneck value. In the following two computational search techniques are used: *myopic hill climbing* (mHC) and *simulated annealing* (SA). They implement adaptive local search; the neighbourhood is defined in terms of applying Hamming mutation.

2.1 Modeling Dynamics at Fitness Level

Let H an heuristic, we assume there is a function $H(X)$ which allows to model the average dynamics of H at fitness level. Given an initial genotype of fitness f_1 , applying the heuristic generates a sequence f_1, f_2, \dots by the iteration $f_{k+1} = H(f_k)$. Our hope is to gain knowledge from function H in order to help us to predict the behaviour of the heuristic. To illustrate this approach, let us consider the heuristic (noted H_{ham}) corresponding to a *random walk*: starting from a initial genotype, at each step the next genotype is chosen at random in the Hamming neighbourhood. Equation 1 may be reformulated in

$$H_{ham}(f) = \left(1 - \frac{K+1}{N}\right) f + \left(\frac{K+1}{N}\right) \beta \quad (2)$$

From an initial fitness value f_1 less than β , the sequence f_1, f_2, \dots increases to β . On average offspring-fitness is higher than parent-fitness; thus the heuristic is selectively advantageous. If fitness is greater than β , the mean offspring-fitness is lower than fitness: on average the heuristic is deleterious. Property $\{\beta = 0.5\}$ means that on average H_{ham} is selectively neutral⁵ on NK-landscapes whatever epistasis is. Starting from $f_1 = 0.5$ ⁶ the heuristic generates the sequence $f_1 = 0.5, f_2 = 0.5, \dots$. In order to get a visual rendering of the long term behavior of an heuristic, for each string in the genotype space a point is plotted; the abscissa of which is the fitness f and the ordinate the fitness f^* of a genotype reached after applying the heuristic a given number of times. Thus, a new scatterplot, the so-called *Limit Fitness Cloud* (noted FC^*), is drawn. We

define the following function:

$$f_{mean}^*(\varphi) = \text{mean}_{x \in G_\varphi} f^*(x)$$

and the fitness value β^* as $f_{mean}^*(\beta)$ ⁷. We call FC_{mean}^* the representative curve of f_{mean}^* .

2.2 Myopic Hill Climber

A myopic hill climbing heuristic (so-called mHC) is used. At each step, the fittest of all of the genotypes that can be produced by a single bit flip is selected. Entire neighbourhood is searched and selection occurs in all cases, even when the best of the one-mutant neighbours of a genotype is less fit than it. Figure 3 (a) shows the FC_{mean} curve is coarsely a line too: the mean offspring-fitness is in proportion to fitness, whatever the epistatic parameter K is. Relation between \tilde{f}_{mean} and the length N , the epistatic parameter K and fitness f verifies the following equation:

$$\tilde{f}_{mean}(f) = \left(1 - \frac{K+1}{N}\right) f + \left(\frac{K+1}{N}\right) E(X) \quad (3)$$

The mean term $E(X)$ is equal to $E(X(N, K))$ where $X(N, K) = \max(X_1, \dots, X_N)$ and X_i follows normal law $\mathcal{N}(0.5, \frac{1}{\sqrt{12(K+1)}})$. Then, correlation between parent-fitness and the mean offspring-fitness is linear. Therefore $\beta = E(X(N, K))$ is not any more constant but depends from N and K . When N is fixed, β grows as the amount of epistasis decreases. A least squares regression is computed from the FC_{mean} set for $N = 20$ and $K = 15$, we find:

$$\tilde{f}_{mean} = 0.200f + 0.516 \quad (4)$$

which agrees to equation 3

The FC allows to predict whether in the fitness space a bottleneck is likely to arise. We conjecture there is such a bottleneck in the vicinity of the β fitness level. This means that, applying mHC heuristic from a point which fitness is below β , on average the search process breaks off around β . In particular from a random initial point, dynamic pushes fitnesses toward β . In order to validate this hypothesis, from each genotype as initial point, mHC is ran over 50 generations to collect the fitness of the last point encountered. All these informations are got together to build up the corresponding *limit fitness cloud* see figure 3 (b). According to the initial fitness value (abscissa) we can observe two main cases about the limit behavior of the mHC heuristic. Applying the heuristic a given number of times results to an equilibrium state where the mean fitness f^* no longer changes from generation to generation. First, for an initial fitness below β , on average mHC "converges" to the fitness value β^* .

⁵ H_{ham} induces no effect on fitness level

⁶ Fitness of a random initial genotype is on average closed to the mean fitness over the search space ($\bar{f} = 0.5$)

⁷ of course, in this definition we assume that β exists

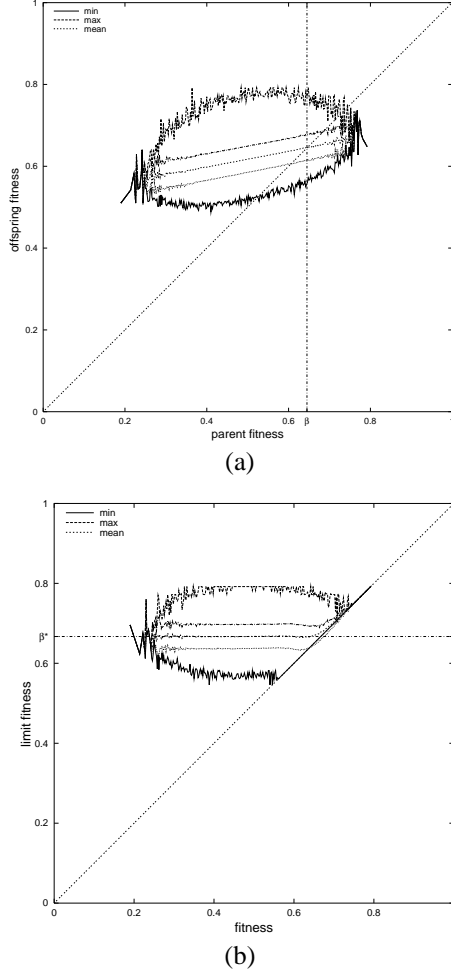


Figure 3: Fitness Cloud under mHC: border (min and max) and mean (with standard deviation) (FC (a), FC^* (b)) (NK-landscape with $N = 20$ and $K = 15$)

Let's note that β is smaller than β^* but have the same magnitude (see tab. 1). Second, fitnesses f above β are fixed points in the fitness space ($f^* = f$). Let's note that there is a transition range around β where f^* depends not linearly on f . These experiments support the *bottleneck conjecture*: it is difficult to bypass the β^* fitness level for the mHC heuristic. Of course, breaking the bottleneck may occur for a particular initial genotype as in this study we look for the average behavior only.

2.3 Simulated Annealing

Simulated annealing (SA) can be seen as a way of trying to allow solution to get away from local optima and move toward fitter point. The SA algorithm employs a random search that not only accepts changes that increase the fitness function, but also some changes that decrease the fit-

metaheuristic	β	β^*
mHC	0.645	0.667
SA ($T = 0.10$)	0.524	0.559
SA ($T = 0.05$)	0.548	0.590
SA ($T = 0.01$)	[0.604, 0.792]	0.656
SA (Generation 50)	-	0.560
SA (Generation 1000)	-	0.613
SA (Generation 1900)	-	0.682
SA (Generation 2450)	-	0.701
nHC	[0.686, 0.792]	0.746


Table 1: Experimental values of β and β^* for $N = 20$ and $K = 15$ with mHC, SA and nHC. The maximum fitness value for this fitness landscape is 0.792

ness value, thus allowing SA to jump out of local maxima. SA search technique takes its inspiration from the models of the annealing physical process [6]. It is search process based on using a parameter which can play the role of temperature. The ability to avoid to get stuck in local optima depends on the choice of initial temperature, the number of iterations performed at each temperature, and the way the temperature is decremented. At each step, one genotype from all of the genotypes that can be produced by a single bit flip is selected and the resulting change, $\Delta f = \tilde{f} - f$, in fitness is computed. If $\Delta f > 0$, the new point is accepted; else, it is accepted with probability $e^{\Delta f/T}$, where T is the temperature control parameter. One major problem with SA is to control the *cooling process*. Often the cooling schedule is developed by trial and error for each particular landscape. First, with regard to given temperature values, the fitness cloud is analyzed; then a cooling process is implemented. For a given temperature T , relation between \tilde{f}_{mean} , f , N , K and T can be derived from equation 3. In this case,


$$\begin{aligned}
 E(X) &= E(X(f, N, K, T)) \\
 &= 1 - \phi\left(\frac{f - 0.5}{\sigma_K}\right) + \int_{-\infty}^f \varphi\left(\frac{x - f}{\sigma_K}\right) e^{\frac{x-f}{T}} dx
 \end{aligned}$$

where φ and Φ are respectively the density and repartition function of the reduced centered normal law. So $E(X)$ depends, not only on N and K , but on the fitness value too. As a consequence, the set FC_{mean} is no more represented by a line but by a curved shape.

High temperature As predicted by our analytical study, at high temperature ($T = 0.10$), the FC_{mean} set is represented by a curved shape (see fig. 4 (a)). This curve crosses the diagonal at point $(\beta; \beta)$ distinctly; so it is easy to estimate β (see tab. 1). Plotting the limit fitness cloud FC^* shows that it is difficult to bypass the β^* bottleneck level for the SA at high temperature (see fig. 4 (b) and tab. 1).



(a)




(b)

Figure 4: Fitness Cloud under SA at **high temperature**: border (min and max) and mean (with standard deviation) (FC (a), FC^* (b)) (NK-landscape with $N = 20$ and $K = 15$)


Let's note that, β^* is reached whatever the initial fitness is (except for extreme fitness values).

Medium temperature At medium temperature ($T = 0.05$), the fitness cloud is roughly shared by the diagonal line as the probability to accept deleterious mutation remains significant (see fig. 5 (a)). It is easy to estimate both β and β^* (see tab. 1). The β^* fitness level appears to be attractive on the limit fitness cloud (see fig. 5 (b) and tab. 1). Once again, β^* is reached whatever the initial fitness is.

Low temperature At low temperature ($T = 0.01$), the greatest part of the fitness cloud is above the diagonal line as the probability to accept deleterious mutation is small (see fig. 6 (a)). The FC_{mean} curve is a curved shape; as fitness increases, it glides slope toward the diagonal. So it is diffi-



(a)



(b)

Figure 5: Fitness Cloud under SA at **medium temperature**: border (min and max) and mean (with standard deviation) (FC (a), FC^* (b)) (NK-landscape with $N = 20$ and $K = 15$)

cult to visualize point which abscissa is β . Examining data, we can find an interval where the FC_{mean} curve is close to the diagonal line (with a accuracy of 0.002) (see tab. 1). The set FC_{mean}^* is roughly represented by an horizontal line except for high fitnesses where it follows the diagonal line (see fig. 6 (b) and tab. 1). As the constant value of f_{mean}^* corresponds to a bottleneck, β^* stands for this value, although β is not discerned.

Cooling process The notion of system temperature is intrinsic to the SA process. By slowly lowering the temperature of an initially random system, we encourage the elements of the system to assume an orderly, minimal energy arrangement. In search process terms, a slow cooling can thus lead to an optimal state. Our annealing schedule is defined as follows: temperature starts at 0.10, and drops to




Figure 7: Limit fitness cloud under SA with **cooling process**: generation 50 (a), generation 1000 (b), generation 1900 (c), generation 2450 (d) (NK-landscape with $N = 20$ and $K = 15$)

0.01 over a geometric decrements where $T = 0.95T^8$. 2450 generations are performed; changes occurs each 50 generations. The cooling process is sufficiently slow such that for each temperature an equilibrium state, where the mean fitness no longer changes from generation to generation, is reached. Figure 7 shows snapshots of the FC^* cloud at generations 50, 100, 1900 and 2450. As for low temperature, β^* stands for the constant value of f_{mean}^* since β is not significant. During the cooling process β^* increases with generations to finally, reaches its greatest value (see tab. 1). Let's note that the final β^* value for SA (0.701) is greater than the one for mHC (0.667).

⁸Experience has shown that the decrement coefficient should be between 0.8 and 0.99, with better results being found in the higher end of the range

3 Fitness Cloud and neutrality : Neutral Hill Climber

In the Fitness Cloud, a vertical slice represents the set of fitnesses that could be reached from this set of neutrality. Consequently the FC shows the potential interest in using neutral operator. To implement such an operator noted nOP, first the entire search space is partitioned according to fitness, then we are able to choose at random a genotype with a given fitness value⁹. So each genotype with the same fitness are connected by an elementary neutral move. To combine neutral exploration with local search technique, we define the *neutral Hill Climbing* heuristic (so-called nHC). First mHC is applied on the current genotype g and the resulting change, $\Delta f = \tilde{f} - f$, in fitness is computed. If $\Delta f > 0$, the new point is accepted; else, we obtain off-

⁹with an uniform law

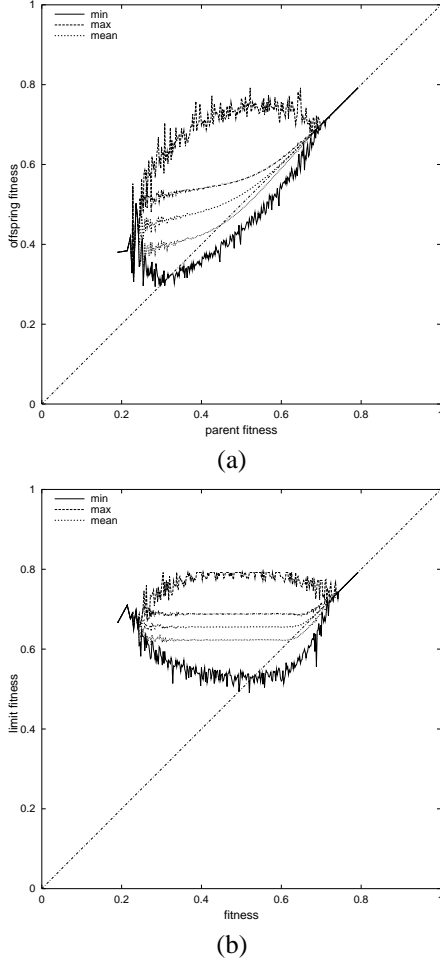


Figure 6: Fitness Cloud under SA at **low temperature**: border (min and max) and mean (with standard deviation) (FC (a), FC^* (b)) (NK-landscape with $N = 20$ and $K = 15$)

spring by applying nOP to the genotype g . Within a neutral set there is no productive gradient information, then the gain fitness comes from Hamming based mutation only. The relation between \tilde{f}_{mean} , f , N and K can be derived from equation 3. In this case, $E(X) = E(X(f, N, K))$ where $X(f, N, K)$ is $\max(f, X_1, \dots, X_N)$ with X_i follows normal law $\mathcal{N}(0.5, \frac{1}{\sqrt{12(K+1)}})$. As a consequence, the set FC_{mean} is represented by a curved shape. Figure 8 shows the FC_{mean} curve is coarsely a line except for high fitness values where it glides slope toward the diagonal. So it is difficult to visualize point which abscissa is β . Examining data, we can find an interval where the FC_{mean} curve is close to the diagonal line with a accuracy of 0.002 (see tab. 1). To support the bottleneck conjecture the *limit fitness cloud* is plotted (see fig. 8 (b)). According to the initial fitness value we can observe two main cases about the limit

behavior of the nHC heuristic. First, for an initial fitness below $\beta^* = 0.746$ (as for SA with low temperature, β^* is the constant value of f_{mean}^*), on average nHC converges to a fitness value close to β^* . Second, fitnesses above β^* are fixed points in the fitness space. Therefore the β^* fitness level is a bottleneck for the nHC heuristic. In nHC, leaving the genotype invariant instead of applying nOP, the fitness cloud remains identical. The FC^* allows to show the influence of nOP: neutral exploration allows to find better fitnesses. Let's note that the fitness bottleneck for nHC (0.746) is greater than the one for both mHC (0.667) and SA (0.701). These experiments show the potential interest in using neutral operator when each plateau is a graph connected by a neutral operator. Of course this is an ideal case, in more realistic situations we must consider the topology graph of neutral sets induced by hamming mutation or a specific neutral operator. In real word problems one must take into account the computational cost in using neutral operator as well the availability of such operator. However in many problems neutral operator may derive from specific knowledge as symmetry properties or redundancies.

Conclusion

In this paper we have presented the Offspring-Parents Fitness Cloud. The Fitness Cloud shows evolvability as well neutrality and dynamic at fitness level of local search heuristic. We have used FC to formulate the *bottleneck conjecture*: there is a fitness level in the NK-landscape that causes local search process to slow down and stop. Moreover the FC is useful in identifying the bottleneck's location in the neighbourhood of a fitness value β . This conjecture deals with the average behavior of local search heuristics only. To confirm this conjecture, the Limit Fitness Cloud is introduced; it gives a visual rendering of the long term behavior. The experiments performed on NK-landscape with Hill Climbing, Simulated Annealing and Neutral Hill Climbing support our conjecture. Indeed there is a fitness level β^* close to β which is difficult to bypass. FC can also represent the neutrality of landscape. So, we have designed the Neutral Hill Climbing (nHC) which shows how to exploit the information about neutrality given by FC. We have found experimentally that the bottleneck for nHC is greater than the one for both mHC and SA.

In this work exhaustive enumeration of the search space is used; future works should address the question of how to get FC from data collected through random sampling or during the search process. To track the dynamics of population based heuristics, as *genetic algorithms*, notion of population must be taken into account. The approach can be extended to others operators than local variations, in particular we project to draw the FC for crossover. Another extension is to study the effect of choosing one representation: for in-




Figure 8: Fitness Cloud under nHC: border (min and max) and mean (with standard deviation) (FC (a), FC^* (b)) (NK-landscape with $N = 20$ and $K = 15$)

stance, what happen to the FC when switching from Integer coding to Gray coding ? The NK-landscapes can be useful for initial investigations, but results gained on them cannot be guaranteed to transfer to real word problems in general. So we obviously have to study fitness cloud and the fitness bottleneck hypothesis on other problems than NK-landscapes.

Bibliography

- [1] Lee Altenberg. The evolution of evolvability in genetic programming. In *In Kinneer, Kim (editor). Advances in Genetic Programming*. Cambrige, MA, pages 47–74. The MIT Press, 1994.
- [2] L. Altenberg G. P. Wagner. Complexes adaptations and the evolution of evolvability. In *Evolution*, pages 967–976, 1996.
- [3] J. J. Grefenstette. Predictive models using fitness distributions of genetic operators. In D. Whitley, editor, *Foundations of Genetic Algorithms*, San Mateo, CA, 1995. Morgan Kaufmann Publishers.
- [4] S. A. Kauffman. “*The origins of order*”. *Self-organization and selection in evolution*. Oxford University Press, New-York, 1993.
- [5] M. Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK, 1983.
- [6] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. In *Science*, pages 671–680, 1983.
- [7] B. Manderick, M. de Weger, and P. Spiessens. The genetic algorithm and the structure of the fitness landscape. *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 143–150, 1991.
- [8] Helge Rosé, Werner Ebeling, and Torsten Asselmeyer. The density of states - a measure of the difficulty of optimisation problems. In *Parallel Problem Solving from Nature*, pages 208–217, 1996.
- [9] Smith, Husbands, Layzell, and O’Shea. Fitness landscapes and evolvability. *Evolutionary Computation*, 1(10):1–34, 2001.
- [10] E. D. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. In *Biological Cybernetics*, pages 63:325–336, 1990.
- [11] E. D. Weinberger. NP completeness of kauffman’s NK model, a tuneable rugged fitness landscape. *Santa Fe Institute: Working Papers* 96-02-003, 1996.
- [12] S. Wright. The theory of gene frequency. In *Evolution and the genetics of Population*, volume 2, pages 120–143. University of Chicago Press, 1969.