



HAL
open science

Deceptiveness and Neutrality - the ND family of fitness landscapes

William Beaudoin, Sébastien Verel, Philippe Collard, Cathy Escazut

► **To cite this version:**

William Beaudoin, Sébastien Verel, Philippe Collard, Cathy Escazut. Deceptiveness and Neutrality - the ND family of fitness landscapes. Genetic And Evolutionary Computation Conference, Jul 2006, Seattle, United States. pp.507 - 514, 10.1145/1143997.1144091 . hal-00164694

HAL Id: hal-00164694

<https://hal.science/hal-00164694>

Submitted on 23 Jan 2009

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.

Deceptiveness and Neutrality

The ND Family of Fitness Landscapes

William Beaudoin, Sébastien Verel, Philippe Collard, Cathy Escazut
University of Nice-Sophia Antipolis
I3S Laboratory
Sophia Antipolis, France
{beaudoin,verel,pc,escazut}@i3s.unice.fr

ABSTRACT

When a considerable number of mutations have no effects on fitness values, the fitness landscape is said neutral. In order to study the interplay between neutrality, which exists in many real-world applications, and performances of metaheuristics, it is useful to design landscapes which make it possible to tune precisely neutral degree distribution. Even though many neutral landscape models have already been designed, none of them are general enough to create landscapes with specific neutral degree distributions. We propose three steps to design such landscapes: first using an algorithm we construct a landscape whose distribution roughly fits the target one, then we use a simulated annealing heuristic to bring closer the two distributions and finally we affect fitness values to each neutral network. Then using this new family of fitness landscapes we are able to highlight the interplay between deceptiveness and neutrality.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search.

General Terms

Algorithms, performance, design, experimentation.

Keywords

Fitness landscapes, genetic algorithms, search, benchmark.

1. INTRODUCTION

The Adaptive Landscape metaphor introduced by S. Wright [1] has dominated the view of adaptive evolution: an uphill walk of a population on a mountainous fitness landscape in which it can get stuck on suboptimal peaks. Results from molecular evolution has changed this picture: Kimura's model [2] assumes that the overwhelming majority of mutations are either effectively neutral or lethal and in the latter case purged by negative selection. This assumption is called the neutral hypothesis. Under this hypothesis,

dynamics of populations evolving on such neutral landscapes are different from those on adaptive landscapes: they are characterized by long periods of fitness stasis (population stashed on a 'neutral network') punctuated by shorter periods of innovation with rapid fitness increases [3]. In the field of evolutionary computation, neutrality plays an important role in real-world problems: in design of digital circuits [4] [5] [6], in evolutionary robotics [7] [8]. In those problems, neutrality is implicitly embedded in the genotype to phenotype mapping.

1.1 Neutrality

We recall a few fundamental concepts about fitness landscapes and neutrality (see [9] for a more detailed treatment). A landscape is a triplet (S, V, f) where S is a set of *potential solutions* i.e. a search space, $V : S \rightarrow 2^S$, a *neighbourhood* structure, is a function that assigns to every $s \in S$ a set of neighbours $V(s)$, and $f : S \rightarrow \mathbb{R}$ is a fitness function that can be pictured as the "height" of the corresponding potential solutions. The neighbourhood is often defined by an operator like bitflip mutation. A *neutral neighbour* of s is a neighbour with the same fitness $f(s)$. The *neutral degree* of a solution is the number of its neutral neighbours. A fitness landscape is neutral if there are many solutions with high neutral degree. A *neutral network*, denoted NN , is a connected graph where vertices are solutions with the same fitness value and two vertices are connected if they are neutral neighbours.

1.2 Fitness Landscapes with Neutrality

In order to study the relationship between neutrality, dynamics of Evolutionary Algorithms (EA) and search difficulty, some benchmarks of neutral landscapes have been proposed. More often neutrality is either an add-on feature, as in NK-landscapes, or an incidental property, as in Royal-Road functions. In most cases the design acts upon the amount of solutions with the same fitness. Royal-Road functions [10] are defined on binary strings of length $N = n.k$ where n is the number of blocks and k the size of one block. The fitness function corresponds to the number of blocks which are set with k bits value 1 and neutrality increases with k . Numerous landscapes are variant of NK-Landscapes [11]. The fitness function of an NK-landscape is a function $f : \{0, 1\}^N \rightarrow [0, 1)$ defined on binary strings with N bits. An 'atom' with fixed epistasis level is represented by a fitness component $f_i : \{0, 1\}^{K+1} \rightarrow [0, 1)$ associated to each bit i . It depends on the value at bit i and also on the values at K other epistatic bits. The fitness f is the average of the values of the N fitness components f_i . Several variants of NK-landscapes try to reduce the number of fitness values in order to add some neutrality. In NKp -landscapes [12], $f_i(x)$ has a probability p to be equal to 0; in NKq -landscapes [13], $f_i(x)$ is uniformly distributed in the interval $[0, q-1] \cap \mathbb{N}$; in *Technological*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

Landscapes [14], tuned by a natural number M , $f(x)$ is rounded so that it can only take M different values. For all those problems, neutrality is tuned by one parameter only: neutrality increases according to p and decreases with q or M (see for example Figure 2).

Dynamics of a population on a neutral network are complex, even on flat landscapes as shown by Derrida [15]. The works [16] [17] [18] [19], at the interplay of molecular evolution and optimization, study the convergence of a population on neutral networks. In the case of infinite population under mutation and selection, they show distribution on a NN is only determined by the topology of this network. That is to say, the population converges to the solutions in the NN with high neutral degree. Thus, the neutral degree distribution is an important feature of neutral landscapes.

In order to study more precisely neutrality, for instance link between neutrality and search difficulty, we need for “neutrality-driven design” where neutrality really guides the design process. In this paper we propose to generate a family of landscapes where it is possible to tune accurately the neutral degree of solutions.

2. ND-LANDSCAPES

In this section, we first present an algorithm to create a landscape with a given *neutral degree distribution*. Then we will refine the method to obtain more accurate landscapes and finally we will study time and space complexity of the algorithm.

2.1 An algorithm to design small ND-Landscape

We now introduce a simple model of neutral landscapes called ND-Landscapes where N refers to the number of bits of a solution and D to the neutral degree distribution. In this first step our aim is to provide an exhaustive definition of the landscape assigning one fitness value to each solution. We fix N to 16 bits and so the size of search space is 2^{16} . Building a ND-Landscape is done by splitting the search space into neutral networks. However the fitness value of each neutral network has no influence on the neutrality. This is why these fitness values are randomly chosen.

Let D be an array of size $N+1$ representing a neutral degree distribution. N and D are given as inputs and the algorithm (see algorithm 1) returns a fitness function f from $\{0, 1\}^N$ to \mathbb{R} such that the neutral degree distribution of the fitness landscape is similar to D . For more simplicity, we chose to give a different fitness value to each neutral network. We define *RouletteWheel(D)* as a random variable whose density is given by distribution D . It is directly inspired from the genetic algorithm selection operator. For example: let Δ be the following distribution :

$\Delta[0]=0 \quad \Delta[1]=0.25 \quad \Delta[2]=0.5 \quad \Delta[3]=0.25$. *RouletteWheel(Δ)* will return value **1** in 25 % of the time, **2** in 50 % of the time and **3** in 25 % of the time. Figure 1 shows the neutral networks of an ideal ND-Landscape (size= 2^5) for the distribution Δ .

2.2 A metaheuristic to improve the ND design

Using algorithm 1, exhaustive fitness allocation does not create a landscape with a neutral degree distribution close enough to the input distribution. The reason is the fitness function is completely defined before the neutral degree of every solution has been considered. Hence, we use a simulated annealing metaheuristic to improve the landscape created by algorithm 1. Here, simulated annealing is not used to find a good solution of a ND-Landscape but to adjust the landscape by modifying the fitness of some solutions such as neutral distribution of a ND-Landscape be closer to the input distribution. The local operator is the changement of fitness value of one solution of the landscape, which can alter at most $N+1$

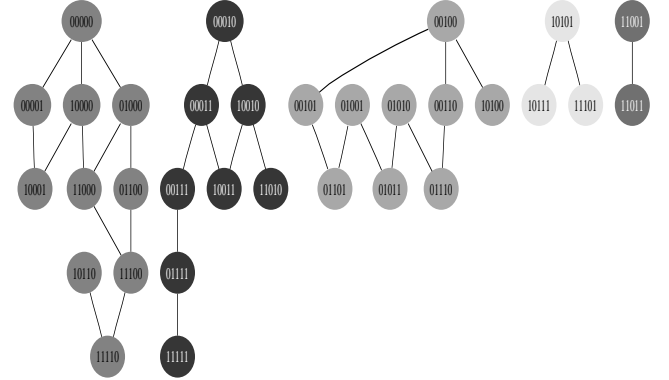


Figure 1: Example of a tiny ND-Landscape. Each node represents a solution and two nodes are connected if they have the same fitness value and are Hamming neighbours. In this example there are five neutral networks.

Algorithm 1 Generation of ND-Landscapes

```

 $\forall s \in S, f[s] \leftarrow$  unaffected
randomly choose one solution  $s_0$ .
CandidatesList  $\leftarrow$  S sorted by distance from  $s_0$ .
while not empty(CandidatesList) do
   $s \leftarrow$  head(CandidatesList)
  for  $d = 0$  to  $N$  do
    if  $s$  can't have  $d$  neutral neighbours
      then  $D'[d] \leftarrow 0$ 
    else  $D'[d] \leftarrow D[d]$ 
  end for
   $n \leftarrow$  RouletteWheel( $D'[d]$ )
  Give a value to some unaffected neighbours so that  $s$  has exactly  $n$  neutral neighbours and so that the neutral degrees of already chosen solutions ( $\notin$  CandidatesList) are unchanged.
   $D[n] \leftarrow D[n] - \frac{1}{2^N}$ 
  CandidatesList  $\leftarrow$  next(CandidatesList)
end while

```

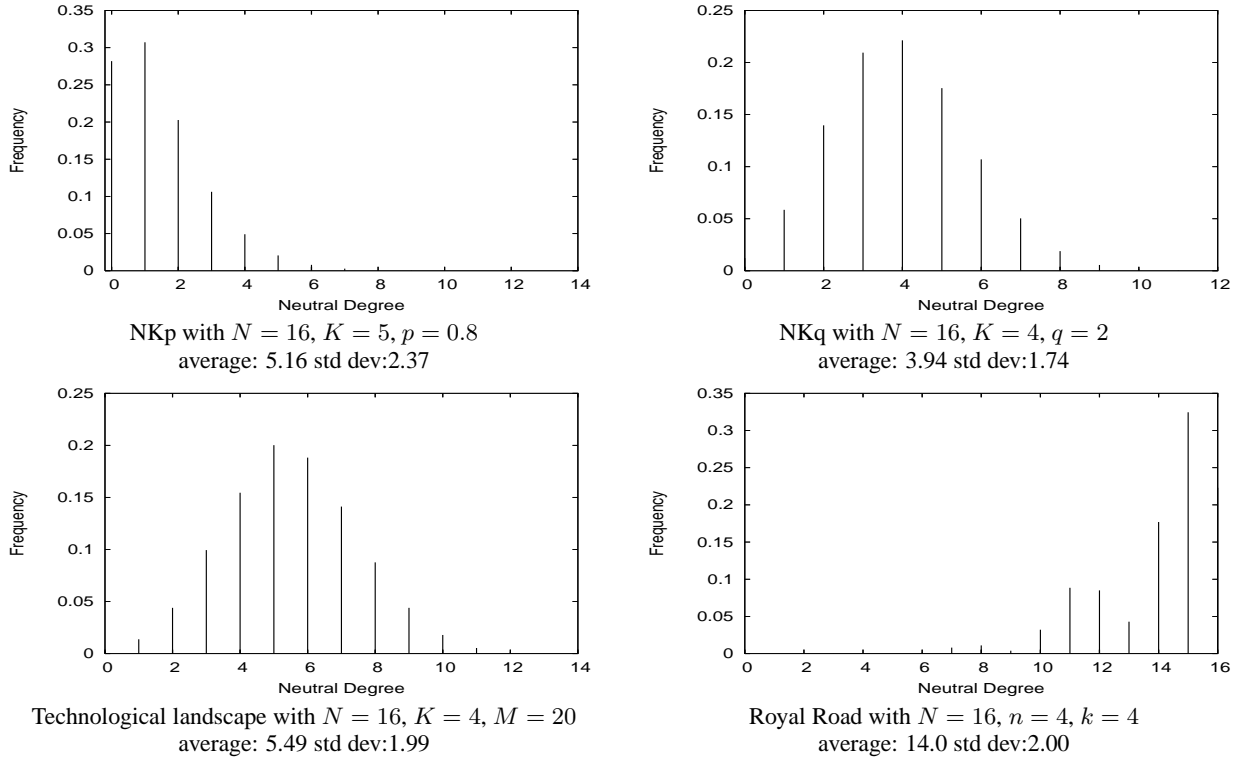


Figure 2: Neutral degree distribution for some neutral landscapes

neutral degrees. The acceptance of a transition is determined by the difference between the distance to the input distribution before and after this transition. The distance we use to compare two distributions is the root mean square :

$$dist(D, D_0) = \sqrt{\sum_{i=0}^N (D[i] - D_0[i])^2}$$

Simulated annealing appeared to be a fast and efficient method for this particular task. Improvements made by simulated annealing are shown in figure 3. Easiest neutral distributions to obtain seemed to be smooth ones like gaussian distributions (which are the most encountered when dealing with real or artificial neutral problems). On the other hand, sharp distributions (like the middle-right one of figure 3) are really hard to approach. In addition, independently of the shape, distributions with higher average neutral degree are harder to approximate.

2.3 Space and Time complexity

To create a landscape with a search space of size 2^N , we use an array of size 2^N containing fitness values and a list of forbidden values for each solution. Thus we need a memory space of size $\mathcal{O}(2^N \times N)$. Consequently the space complexity is : $\mathcal{O}(2^N \times N)$

In order to know what are the possible neutral degrees of an unaffected solution s , we must consider every interesting value for s (the fitnesses of all neighbour solutions and a random value), and for each of these values, we must find out all possible neutral degrees. This can be done in a time $\mathcal{O}(N^2)$. We evaluate the possible neutral degrees once for each solution. Time allowed for simulated annealing is proportional to the time elapsed during construction. Thus, the time complexity of the algorithm is $\mathcal{O}(2^N N^2)$.

Consequently we can only construct ND-Landscapes with a small N (≤ 16) but we will see in section 4 how to create Additive Extended ND-Landscapes with far greater search spaces.

2.4 Sizes of the generated Neutral Networks

Figure 4 shows the diversity of sizes of neutral networks for 4 distributions. For every distribution we created 50 different ND-Landscapes. Graphics on the left show the input and the mean resulting distribution. Graphics on the right show all of the networks of these landscapes sorted by decreasing size with a logarithmic scale. We clearly see that the neutral degree distribution is a really determining parameter for the structure of the generated landscape.

3. TUNING DECEPTIVENESS OF ND-LANDSCAPES

Once we have generated a landscape with a specific neutral degree distribution, we can change the fitness value of all neutral networks without changing the neutral degree distribution (as long as we do not give the same fitness to two adjacent networks). Hence, for a given neutral distribution, we can tune the difficulty of a ND-Landscape. For instance if each NN have a random fitness value from $[0, 1]$ then the landscape is very hard to optimize. Here, we will use the well known *Trap* Functions [20] to affect fitnesses to NN in order to obtain a ND-Landscape with tunable deceptiveness.

The trap functions are defined from the distance to one particular solution. They admit two optima, a global one and a local one. They are parametrized by two values b and r . The first one, b allows to set the width of the attractive basin for each optima, and r sets their relative importance. The function $f_T : \{0, 1\}^N \rightarrow \mathbb{R}$ is so




Figure 3: Neutral Degree Distributions obtained by algorithm 1 (D_1) and then adjusted by simulated annealing (D_2). Neutral degrees are on abscissa. Impulses represent the target distributions (D_0).

defined by:

$$f_T(x) = \begin{cases} 1.0 - \frac{d(x)}{b} & \text{if } d(x) \leq b, \\ \frac{r(d(x)-b)}{1.0-b} & \text{elsewhere} \end{cases}$$

where $d(x)$ is the Hamming distance to the global optimum, divided by N , between x and one particular solution. The problem is most deceptive as r is low and b is high. In our experiment, we will use two kinds of Trap functions with $r = 0.9$, one with $b = 0.25$ and another one with $b = 0.75$ (see figure 5 (a) and (b)).

To affect a fitness value to each neutral network, we first choose the optimum neutral network, denoted NN_{opt} , (for example the one containing the solution 0^N) and set its fitness to the maximal value 1.0. Then, for each neutral network, we compute the distance d between its centroid¹ and the centroid of NN_{opt} ; finally the fit-

¹The centroid of a NN is the string of the frequency of appearance

ness value of the NN is set according to a trap function² and the distance d . In order to ensure that all adjacent networks have different fitness values, it is possible to add a white noise to the fitness values of each NN . In the following experiments, the length of bitstring is $N = 16$. ND-landscapes are constructed with uniform neutral degree distributions. We use the distributions defined by

$$D_{p,w}[i] = \begin{cases} 1/w & \text{if } i \in \{p, p+w-1\}, \\ 0 & \text{elsewhere} \end{cases}$$

where $p \in \{0, 7\}$, $w \in \{3, 4\}$, and the two Trap functions defined in figure 5. For each distribution and each Trap function, 30 landscapes were generated.

of bit value 1 at each position.

²This trap function is defined for all real numbers between 0 and N

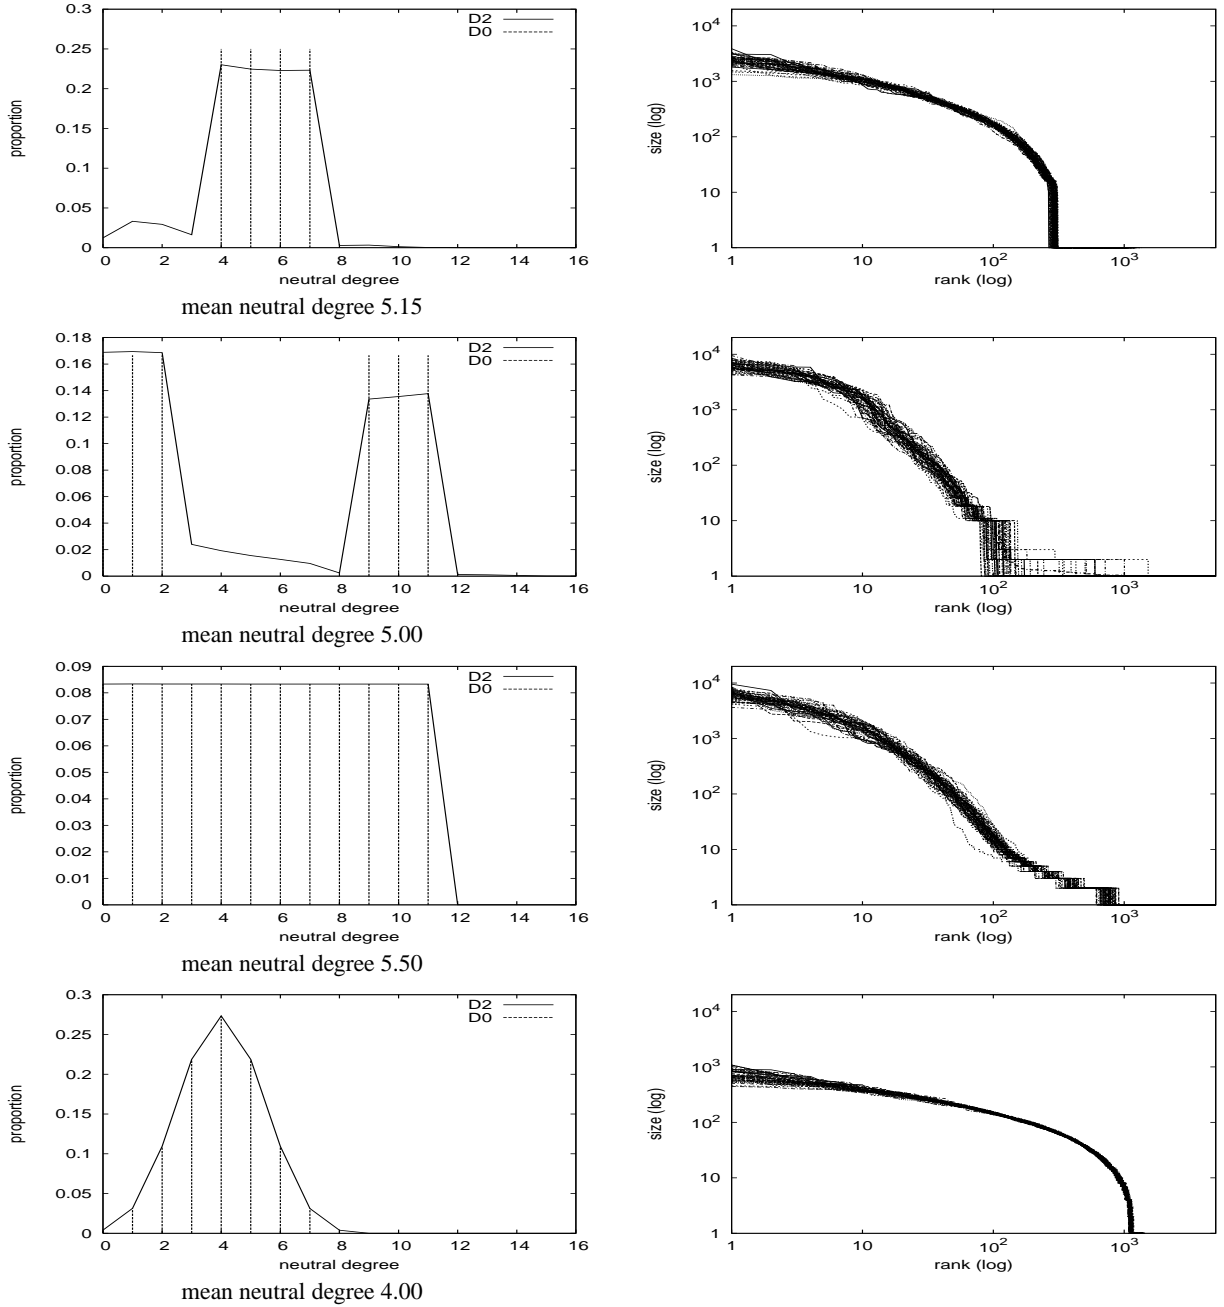


Figure 4: Neutral networks sizes for various ND-landscapes

3.1 Fitness Distance Correlation of ND-Landscapes

To estimate the difficulty to search in these landscapes we will use a measure introduced by Jones [21] called *fitness distance correlation* (FDC). Given a set $F = \{f_1, f_2, \dots, f_m\}$ of m individual fitness values and a corresponding set $D = \{d_1, d_2, \dots, d_m\}$ of the m distances to the global optimum, FDC is defined as:

$$FDC = \frac{C_{FD}}{\sigma_F \sigma_D}$$

where:

$$C_{FD} = \frac{1}{m} \sum_{i=1}^m (f_i - \bar{f})(d_i - \bar{d})$$

is the covariance of F and D and σ_F , σ_D , \bar{f} and \bar{d} are the standard deviations and averages of F and D . Thus, by definition, FDC stands in the range $[-1, 1]$. As we hope that fitness increases as distance to global optimum decreases, we expect that, with an ideal fitness function, FDC will assume the value of -1 . According to Jones [21], problems can be classified in three classes, depending on the value of the FDC coefficient:




Figure 5: Two trap functions: Deceptive one with $b = 0.25$ and $r = 0.9$ (a), Easy one with $b = 0.75$ and $r = 0.9$ (b)

- *Deceptive problems* ($FDC \geq 0.15$), in which fitness increases with distance to optimum.
- *Hard problems* ($-0.15 < FDC < 0.15$) in which there is no correlation between fitness and distance.
- *Easy problems* ($FDC \leq -0.15$) in which fitness increases as the global optimum approaches.

Hard problems are in fact hard to predict, since in this case, the FDC brings little information. The threshold interval $[-0.15, 0.15]$ has been empirically determined by Jones. When FDC does not give a clear indication i.e., in the interval $[-0.15, 0.15]$, examining the scatterplot of fitness versus distance can be useful. The FDC has been criticized on the grounds that counterexamples can be constructed for which the measure gives wrong results [22, 23, 24].

Figure 6 shows the average and standard deviation of FDC over ND-Landscapes for each set of parameters, neutral distribution and deceptiveness. The absolute value of FDC decreases as we generate ND-Landscapes with more and more neutrality. When adding neutrality, the landscapes are increasingly flatter and thus less easy or deceptive. So neutrality smoothes correlation. Adding neutrality to a deceptive landscape makes it easier and adding neutrality to an easy landscape makes it harder.

3.2 Genetic Algorithm Performances on ND-Landscapes

In this section, difficulty is measured by genetic algorithm performances which is the success rate over 10 independent runs. In order to minimize the influence of the random creation of ND-Landscapes, we consider 30 different landscapes for each distribution D and each trap function. For the GA, one-bit mutation and one-point crossover are used with rates of respectively 0.8 and 0.2. The evolution, without elitism and with 3-tournament selection of a population of 50 individuals took place during 50 generations. Figure 8 shows average and standard deviation of GA performances over the 30 ND-landscapes. Until the average neutral




Figure 6: Average and standard deviation of FDC as a function of average neutral degree for ND-landscapes created from easy and deceptive Trap functions.

degree 5, the landscapes are fully deceptive (a) or fully easy (b). Between neutral degree 5 and 7, the deceptiveness and the easiness decrease. After neutral degree 7, the two trap functions have nearly a same good success rate (0.7). Whereas the FDC coefficient gave no information about performances, examining the scatter plot fitness/distance allowed to predict these good performances (figure 7 (c)). When neutrality increases on deceptive trap, performances increase whereas on easy trap they decrease. These results confirm conclusions found in section 3.1 from FDC measures.

4. ADDITIVE EXTENDED ND-LANDSCAPES


Exhaustive fitness allocation allows only to generate landscapes with small search space (in our experiments 2^{16} solutions). Hence we must find a way to construct similar problems on a larger scale. We propose here to concatenate several small ND-Landscapes to create an additive ND-Landscape. Even though we are not able to create an additive ND-Landscape directly from a neutral degree distribution we can choose the mean, the standard deviation and less precisely the shape of his neutral degree distribution. Moreover, this method allows us to know exactly the neutral degree distribution of the resulting landscape. Let be $P_1 = (E_1, V_1, f_1)$ and $P_2 = (E_2, V_2, f_2)$ two fitness landscapes. We define the extended landscape $P = P_1 \oplus P_2 = (E, V, f)$ such as :

- $E = E_1 \times E_2$
- $\forall (x_1, y_1) \in E_1^2, \forall (x_2, y_2) \in E_2^2,$
 $(x_1, x_2) \in V(y_1, y_2) \iff (x_1 \in V_1(y_1) \text{ and } x_2 = y_2) \text{ or } (x_1 = y_1 \text{ and } x_2 \in V_2(y_2))$
- $f(x_1, x_2) = f_1(x_1) + f_2(x_2)$


The size of the larger ND-landscape will be the product of the sizes of the small ones. The neutral degree distribution of the resulting landscape is the *convolution product* of the two components distributions. The convolution product of two distributions D_1 and D_2 is the distribution D (see figure 9) such as

$$\forall n \in \mathbb{N}, D(n) = \sum_{i=0}^n (D_1(i) \times D_2(n-i))$$


We have the following properties :



(a)



(b)



(c)

Figure 7: ND-Landscapes : FDC scatter plot from easy trap (cf Figure 5 a) with respective average neutral degree (a) 1.5 (b) 6 (c) 8.5




Figure 8: Average with standard deviation of success rate of a GA on ND-landscapes created from easy and deceptive trap functions

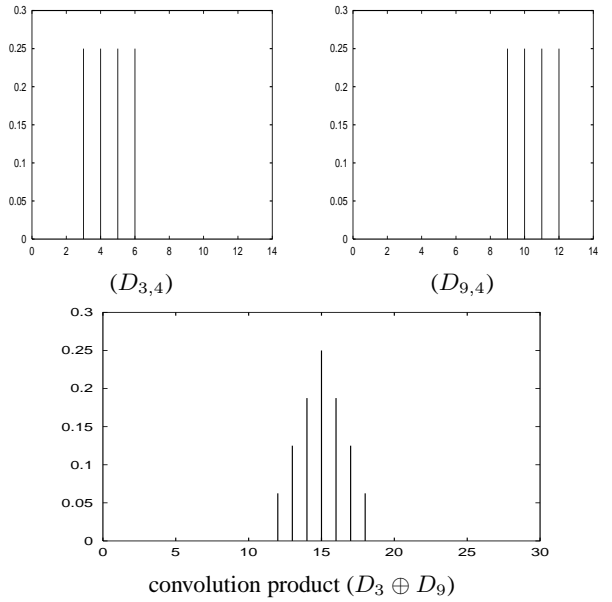


Figure 9: Example of convolution product

$average(D) = average(D_1) + average(D_2)$
and $\sigma(D) = \sqrt{\sigma^2(D_1) + \sigma^2(D_2)}$ where σ is the standard deviation.

The convolution product of two normal (resp. χ^2 , Poisson) distributions is a normal (resp. χ^2 , Poisson) distribution.

Concatenation is commutative, so it is simple to concatenate a number of small landscapes. Hence, we can tack together many small ND-Landscapes generated exhaustively to obtain a bigger one with known neutral degree distribution. Moreover, it has been proven by Jones [21] that the FDC coefficient is unchanged when multiple copies of a problem are concatenated to form a larger problem.

Conclusion

This paper presents the family of ND-Landscapes as a model of neutral fitness landscapes. Most of the academic fitness landscapes

found in the literature deals with neutrality, either as an add-on feature or as an incidental property. Theoretical studies in molecular evolution, as well as in optimization, show that the distribution of neutral degree is a key factor in the evolution of population on a neutral network. Hence to understand and study the influence of neutrality on evolutionary search process it may be helpful to design landscapes with a given neutral distribution. We propose three steps to design such landscapes: first using an algorithm we construct a landscape whose distribution roughly fits the target one, then we use a simulated annealing heuristic to bring closer the two distributions, and finally we affect fitness values to each neutral network. The last step can be used to tune the landscape difficulty according to other criteria than neutrality. So we can study interaction between neutrality and another source of hardness. In this paper we use trap functions to affect fitness values, this allows us to study the interplay between deceptiveness and neutrality; in particular, experimental results show that neutrality smoothes the correlation between fitness and distance. For some problems, introducing neutrality could be benefic as shown in [25] by increasing the evolvability. On the other hand, neutrality could destroy useful information such as correlation as it is shown in our experiments. As these landscapes need an exhaustive enumeration of the search space we propose to concatenate small ND-Landscapes to scale up over 16 bits. Then, using the convolution product of distribution, we are able to design large landscapes with known neutral distribution.

Future work along these lines includes studying the influence of various distributions and different ways to affect fitness value on the dynamics of a population on a neutral network. With ND-trap functions we have focused on the interplay between deceptiveness and neutrality. By replacing trap functions by NK functions[11], we might be able to highlight correlation between epistasis and neutrality.

5. REFERENCES

- [1] S. Wright. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the Sixth International Congress of Genetics 1*, pages 356–366, 1932.
- [2] M. Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK, 1983.
- [3] N. Eldredge and S. J. Gould. Punctuated equilibria: An alternative to phyletic gradualism. In T. J. M. Schopf, editor, *Models in Paleobiology*, pages 82–115. Freeman Cooper, San Francisco, 1972.
- [4] Inman Harvey and Adrian Thompson. Through the labyrinth evolution finds a way: A silicon ridge. In *ICES*, pages 406–422, 1996.
- [5] Vesselin K. Vassilev and Julian F. Miller. The advantages of landscape neutrality in digital circuit evolution. In *ICES*, pages 252–263, 2000.
- [6] Tina Yu and Julian F. Miller. Finding needles in haystacks is not hard with neutrality. In *EuroGP*, pages 13–25, 2002.
- [7] T.M.C. Smith, A. Philippides, P. Husbands, and M. O’Shea. Neutrality and ruggedness in robot landscapes. In *Congress on Evolutionary Computation: CEC2002*, pages 1348–1353. IEEE Press, 2002.
- [8] Chad W. Seys and Randall D. Beer. Evolving walking: The anatomy of an evolutionary search. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, editors, *Animals to Animals, Proceedings of the Eight international Conference on Simulation of Adaptive Behavior*, pages 357–363. The MIT Press, 2004.
- [9] P. F. Stadler. Fitness landscapes. In M. Lässig and Valleriani, editors, *Biological Evolution and Statistical Physics*, volume 585 of *Lecture Notes Physics*, pages 187–207, Heidelberg, 2002. Springer-Verlag.
- [10] M. Mitchell, S. Forrest, and J. H. Holland. The royal road for genetic algorithms: Fitness landscape and GA performance. In F.J. Varela and P. Bourguine, editors, *Proceedings of the First European Conference on Artificial Life*, pages 245–254, Cambridge, MA, 1992. MIT Press/Bradford Books.
- [11] S. A. Kauffman. *The Origins of Order*. Oxford University Press, New York, 1993.
- [12] L. Barnett. Ruggedness and neutrality - the NKp family of fitness landscapes. In C. Adami, R. K. Belew, H. Kitano, and C. Taylor, editors, *ALIFE VI, Proceedings of the Sixth International Conference on Artificial Life*, pages 18–27. ALIFE, The MIT Press, 1998.
- [13] M. Newman and R. Engelhardt. Effect of neutral selection on the evolution of molecular species. In *Proc. R. Soc. London B.*, volume 256, pages 1333–1338, 1998.
- [14] J. Lobo, J. H. Miller, and W. Fontana. Neutrality in technological landscapes. *Santa Fe Working Paper*, 2004.
- [15] B. Derrida and L. Peliti. Evolution in a flat fitness landscape. *Bull. Math. Biol.*, 53:355–382, 1991.
- [16] E. Bornberg-Bauer and H.S. Chan. Modeling evolutionary landscapes: Mutational stability, topology and superfunnels in sequence space. *Proc. Natl. Acad. Sci. USA*, track 2, 96(19):10689–10694, 1999.
- [17] Erik Van Nimwegen, James P. Crutchfield, and Martijn Huynen. Neutral evolution of mutational robustness. In *Proc. Nat. Acad. Sci. USA* 96, pages 9716–9720, 1999.
- [18] Erik Van Nimwegen and James P. Crutchfield. Optimizing epochal evolutionary search: Population-size dependent theory. *Mach. Learn.*, 45(1):77–114, 2001.
- [19] C. O. Wilke. Adaptive evolution on neutral networks. *Bull. Math. Biol.*, 63:715–730, 2001.
- [20] K. Deb and D. E. Goldberg. Analyzing deception in trap functions. In D. Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 93–108. Morgan Kaufmann, 1993.
- [21] T. Jones. *Evolutionary Algorithms, Fitness Landscapes and Search*. PhD thesis, University of New Mexico, Albuquerque, 1995.
- [22] L. Altenberg. Fitness distance correlation: an instructive counterexample. In T. Back, editor, *Seventh International Conference on Genetic Algorithms*, pages 57–64. Morgan Kaufmann, 1997.
- [23] R.J. Quick, V.J. Rayward-Smith, and G.D. Smith. Fitness distance correlation and ridge functions. In A. E. Eiben et al., editor, *Fifth Conference on Parallel Problems Solving from Nature (PPSN’98)*, volume 1498 of *Lecture Notes in Computer Science*, pages 77–86. Springer-Verlag, Heidelberg, 1998.
- [24] Manuel Clergue and Philippe Collard. Ga-hard functions built by combination of trap functions. In David B. Fogel, Mohamed A. El-Sharkawi, Xin Yao, Garry Greenwood, Hitoshi Iba, Paul Marrow, and Mark Shackleton, editors, *Proceedings of the 2002 Congress on Evolutionary Computation CEC2002*, pages 249–254. IEEE Press, 2002.
- [25] Marc Ebner, Mark Shackleton, and Rob Shipman. How neutral networks influence evolvability. *Complex Systems*, 7(2):19–33, 2001.