



DRUID
society

Paper to be presented at
the DRUID16 20th Anniversary Conference
Copenhagen, June 13-15, 2016

Fitness landscapes in organizational theory: Research challenges and future directions

oana Vuculescu
Aarhus University
BADM
omvuc@asb.dk

Carsten Bergenholtz
Aarhus University
BADM
cabe@mgmt.au.dk

Mads Kock
Aarhus University
Department of Physics and Astronomy
madscock@gmail.com

Jacob Sherson
Aarhus University
Department of Physic and Astronomy
sherson@phys.au.dk

Abstract

Using insights from computer science and biology, we argue that current computational modelling research needs to address several fundamental issues in order to generate more meaningful and falsifiable contributions. Based on comparative simulations and a new type of visualization, we address two key elements that the traditional NK framework has relied on: a) how the NK captures the complexity of organizational problems and b) search behaviours where, despite evidence, local search is often used as the dominant problem solving strategy. We show that these two components are fundamentally intertwined and outline implications for how to simulate organizational problems.

Jelcodes:M10,-

Introduction

In the current literature on search two central issues are still left unresolved: 1) in the continuum from mindless particles to perfect rationality, how to model bounded rational search behaviour and 2) how to conceptualize the space of solutions (i.e. the interdependence structure) in which search takes place (Sorenson 2002, Chang and Harrington 2006, Todd, Hills et al. 2012, Baumann 2015). In line with work from computer science (Jones 1995, Pitzer and Affenzeller 2012), as well as the original fitness model formulation (Wright 1932) we show how the two are fundamentally intertwined and in order to be able to create better models of problem solving, they must be addressed together.

Organizational theory has a long tradition of studying organizations' search for solutions to 'hard' problems, i.e. problems where it is computationally impossible or merely too expensive to list and test all possible solutions (Simon 1956; Cohen et al. 1972). The prevalent way of addressing individual or organizational search behaviour and how to conceptualize the space of solutions stems from early work on population genetics, namely the fitness landscape model (Wright 1932). Within biology, by focusing on fitness interactions between genes, Wright's framework allows for a link between low-level properties of genes and the high-level patterns of the dynamics of evolution (Altenberg 1997). The model's most famous extension, the NK model (Kauffman 1993), explicitly models adaptive evolution as a "search in protein space" (Kauffman 1993: p. 37) which tries to find a maximum point for a chosen fitness function. This approach has grown outside the boundaries of population genetics literature and inspired a series of scholars from computer science (e.g. Pitzer and Affenzeller 2012), organizational theory (e.g. Baumann and Siggelkow 2013) and physics (e.g. Sørensen et al. 2015).

How can problem-solving be addressed in this framework? Imagine trying to solve an innovation problem, for instance designing a new educational app. As any software developer would tell you, there is no need to start from zero: nowadays there are a number of pre-defined libraries you can use, which you can think of as interconnected modules. But how are these modules interconnected and how will this affect your chances of finding a good design? (i.e. what does the task environment look like). Once you have a working prototype, should you then just go through each module, one at a time, and try to make minor improvements? (i.e.

what organizational theory calls ‘local search’). What if you get stuck - should you discard everything and start again from scratch? (i.e. ‘random long-jumps’).

Levinthal (1997) introduces the NK model in order to facilitate the formal modelling and simulation of how the level of interdependence in an organization’s routines affects its long-term chances of finding the optimal configuration of such routines and thus survive in a competitive environment. Typically, the organizational literature is primarily focused on how organizations search over the space of routines for combinations leading to increased performance (e.g. Levinthal 1997), but the underlying intuition is the same. By making explicit assumptions about individual or organizational behaviour and the environment in which the agent evolves, researchers could now simulate how such agents adapt over time. Based on these assumptions one can then map the complex dynamics of organizations being embedded *in* and adapting *to* the competitive environment (Levinthal 1997). This has enabled the field to go beyond static explanations and model possible future trajectories of the current competitive situation: e.g. organizations that are highly coupled (have a high interdependence between routines), Levinthal (1997) argues, have a higher likelihood of failure in the face of changing environments. Later papers have developed this approach, addressing either different interdependence structures (e.g. Rivkin and Siggelkow 2003) or how agents (be it individuals or organizations) search or adapt (Gavetti and Levinthal 2000; Gavetti and Levinthal 2001; Winter, Cattani et al. 2007; Baumann and Siggelkow 2013, Martignoni et al. 2015). Almost all of this work, however, looks at either interdependence structures (e.g. Ethiraj and Levinthal 2004) or search behaviours in isolation (e.g. Winter et al. 2007).

This paper attempts to address these issues to shed light on the challenges of a more systematic approach to modelling organizational and individual problem solving and offers a technical and theoretical comparative analysis of a number of assumptions about search strategies and fitness landscapes. This analysis is substantiated by a novel (to the organizational literature) type of visualization that maps how different search strategies actually ‘generate’ different landscapes, rather than just searching in an a priori given space. We identify and explore two main limitations with modelling problem solving via the NK framework. First, recent research (He et al. 2007) has shown that unless certain complexity-theoretical assumptions are wrong¹, for hard problems, a predictive measure of problem hardness cannot exist. This result mirrors a number of red flags already raised in the organizational literature (e.g. Frenken et al. 1999) regarding

¹ P=NP respectively BPP = NP. See He et al. (2007) for an elaboration.

whether the NK is indeed a ‘tuneable complexity landscape’ and calls for a more elaborated and systematic discussion regarding the kind of problems that can be modelled via the NK framework. Second, we argue that the NK framework has several limitations with respect to modelling more plausible search behaviours, in particular due to particular assumptions regarding the fitness function.

In the following we begin, for historical reasons, to review the use of fitness landscapes in section 2 and search behaviours in section 3. Based on these insights from organizational theory and biology we discuss the implications for organizational theory of the presented visualizations and analysis of simulations in section 4.

2. Fitness landscapes

There are two main elements in the fitness landscape model that need to be specified for problem solving processes to be captured: the task structure (i.e. the problem that is to be solved) and the search behaviour (i.e. how problem solving unfolds). In order to discuss the difference between the “objectively defined task” (Simon and Newell 1971: p. 148) and the fitness landscape, which is the backbone of the NK simulations, we start by providing a short formal definition of fitness landscapes. The fitness landscape is in effect what the solver subjectively perceives or ‘the problem space’ (Simon and Newell 1971). We start by providing an overview of the main elements that describe an NK landscape. We caution that although we attempt to discuss the two issues in sequence, a certain amount of overlap is inevitable.

2.1. NK landscapes

In an optimization problem, a solution to a given problem is represented as a size N vector of (traditionally) binary variables². The quality of a solution, in keeping with the biological inspiration of the model, is given by a ‘fitness function’ (i.e. an objective function).

Let X be the space of all possible solutions to a problem. In a maximizing optimization problem³(i.e. a problem where the goal is to find the solution for which the function f has the highest value), for a function f , a solution $x^* \in X$ is a global optimum if

² The model can be easily and without loss of generality extended to larger alphabets.

$$f(x^*) \geq f(x), \text{ for any } x \in X. \quad (1)$$

where all binary variables are in their optimal position.

The landscape is a mapping between solutions and fitness values that takes into account the connectivity between solutions. In order to define connectedness (or pre-defined similarity) between solutions we need to specify a distance metric. Without a definition of a metric one cannot define a fitness landscape as Wright (1932) conceptualized it. As such, for any type of problem, a fitness landscape $L(X, f, d)$ consists of all $x \in X$, an objective function f that measures the quality of each solution and a distance measure d . Following Pitzer and Affenzeller (2012) we can define a *distance metric* as:

$$d: X \times X \rightarrow \mathbb{R}, \text{ such that } d(s,t) \geq 0, \\ d(s,t) = 0 \iff s=t, d(s,t) \leq d(s,u) + d(u,t), \forall s,t,u \in X. \quad (2)$$

Subsequently, the structure of the search space (Rothlauf 2011; Pitzer and Affenzeller 2012) is defined via the ϵ -neighbourhood of x :

$$M(x) = \{n \in X; n \neq x; d(x;n) \leq \epsilon\}. \quad (3)$$

The neighbourhood function thus defines the set of all solutions that are different from the focal solution and within an ϵ radius of it, *as measured by the chosen distance metric* (as given by Equation 2). Thus, the fitness landscape's shape (or topology) is given only in conjunction with *a given neighbourhood function* (Jones 1995): the fitness landscape (L) is not the same as the fitness function (f) and it is likely that for any function f there can be a number of landscapes L with vastly different properties (Maier et al. 2014).

To illustrate this point, Figure 1 shows an example of the same function mapped onto three different landscapes using three different expressions for the neighbourhood function.

[Insert figure 1 around here]

We used a dimensionality reduction method that transforms high dimensional data to low-dimensional representations while preserving pair-wise similarities (Van der Maaten and Hinton 2008) to create 3D visualizations of the multi-dimensional landscape⁴. The illustration on the left depicts one NK landscape generated by relying on a one-bit-flip (i.e. any two

³ Conversely, in a minimizing optimization problem, $x^* \in X$ is a global optimum, if $f(x^*) \leq f(x)$ for any $x \in X$.

⁴ Given the fact that the t-SNE (t-distributed stochastic neighbour embedding) algorithm is stochastic, as is the NK fitness function, it should be noted that this is one possible illustration of one possible NK landscape with $N=8$, $K=3$. The illustration is not a general result for all NK landscapes of $N=8$ and $K=3$.

solutions are considered similar if the Hamming distance⁵ between them is exactly one), while the ones on the right depicts an NK landscape generated by relying on a two-bit flip neighbourhood structure (i.e. any two solutions are considered similar if the Hamming distance between them is exactly two). Both landscapes have the same underlying fitness function, generated for an NK with $N=8$ and $K=3$. Finally, the lower part of the graphic shows the same NK function mapped by converting the bit of strings to the decimal system (cf. Østman and Adami 2014) and solutions are considered to be similar if their decimal representation transformations are consecutive (i.e. 10011001 is transformed into 153 and its natural neighbours are 151 (10011000) and 154 (10011010)). The decimal representation allows for a up-front intuition regarding the distribution of high fitness peaks in the in the solution space (Østman and Adami 2014).

This visualization is telling in two ways. First, note that the two-bit-flip generates two different landscapes. Depending on the starting point, a subset of solutions is not connected in the graph. Similarly if one attempts to traverse a sequence of consecutive numbers with increments of 2, one generates two distinct and unconnected subsets: odd and even numbers. Thus, the definition of the neighbourhood function can effectively reduce (relative to the entire search space) the size of the landscape. Second, the heat-map preserves information regarding the distribution of fitness scores. One can assess (qualitatively in this case) the relative ease or difficulty of navigating towards the lighter coloured areas, in a traditional NK fashion. The three neighbourhood representations yield three different landscape topologies, i.e. smoother gradients such as the left-hand side of the two-bit flip mean that it would be easy for an agent to find the global optimum, while ‘patchier’ surfaces translate into a lower likelihood of success, such as the one-bit and the right-hand side landscape generated by the two-bit flip. Equivalently, in the 2D decimal representation, one can assess the difficulty of finding the global optimum (the highest fitness value), by looking at the shape of the generated curve. In the 2D case, since the decimal representation is arbitrary, the ‘decimal’ landscape is very “rugged”, thus a solver will likely get stuck in a suboptimal solution.

K/N ratios or epistatic interactions

While in the previous section we have shown how landscapes are defined and how the topology of the landscape changes as a result of the definition of the neighbourhood function, we now

⁵ The number of variables which have different values.

turn to how scholars have attempted to describe or more formally capture the ruggedness of a landscape, within the NK framework.

In the NK model, a solution to a given problem is represented as a size N vector of binary variables and the fitness function f is the average across all contributions in the genome. The details of how these contributions are computed are explained below.

$$f(x) = \frac{1}{N} \sum_{i=1}^N x_i \quad (4)$$

The parameter K gives the number of interactions between the N alleles (i.e. variables). For $K=0$, each contribution can take only two possible values. The single allele contributions are independent and identically distributed random variables. Each allele has one state that is preferable to the other, independently of the values of the other variables. The global optimum is the state where all alleles are in the individual optimal position. The global optimum can be reached from any initial configuration.

For $K>0$ the contribution of each allele depends on the position of a number of other K alleles. The choice of which alleles are interdependent is at the latitude of the modeller, but a common assumption is that neighbouring solutions influence each other. For example, for $N=4$, $K=1$, a solver cannot determine the optimal position for the first site in the solution, but has to compute all four possible combinations for the first and second site: {00, 01, 11, 10} and then choose the maximizing sequence. For $K=N-1$ the entire sequence appears in the argument of each single gene contribution and each step replaces the fitness with a different random number. The interdependence between alleles (or solution components) is known as *epistasis*, a term borrowed from biology where it denotes the fact that the expression of a gene is altered by the presence of another.

Early organizational studies relying on the NK model follow on the path proposed by Kauffman (1993) and study how the attributes of the search space influence the propensity of finding the optimal solution by an one-bit-flip hill climber (e.g. Levinthal 1997; Ethiraj and Levinthal 2004). Indeed, part of the NK model's popularity in organizational literature is due to the fact that it allows the investigation of different problem difficulties (Afuah and Tucci 2012), via the K/N ratio, or the level of epistatic interactions (Weise et al. 2009). Epistasis is equivalent to the non-linearity of a problem or how well a problem can be decomposed into sub-problems (Rothlauf 2011; Pitzer and Affenzeller 2012). In other words, epistasis gives a

measure of signal to noise in what concerns the evaluations of the fitness function. Consider the $N=4$, $K=1$ example used before. A solver with $\{x=(0,0,0,1), f(x)=0.56\}$, might move to $\{x^*=(0,0,0,0), f(x^*)=0.58\}$, even though $\{y=(0,0,1,1), f(y)=0.72\}$. The fact that the optimal setting for the fourth allele is $\{1\}$ is obscured by the epistatic interaction with its neighbour on the third position. In the light switch example, it is naturally easier to detect the optimal configurations for the light switches provided they control distinct parts of the lighting in the room. That is, only for $K=0$, the contribution of each allele is independent and a solver can easily detect the optimal configurations and epistasis is 0.

2.2.Landscape features

Despite the fact that the notion of epistatic interactions, as outlined above, is extremely important within organizational theory, its use in quantifying the hardness of a problem has often been criticized (Mason 1995; Naudts and Verschoren 1999) in particular due to the difficulty of identifying measures of epistasis that have enough predictive power (Pitzer and Affenzeller 2012). There are several known limitations to using epistasis measures as proxies for problem complexity. First, epistatic interactions can be both positive and negative. Whether an interaction effect between two alleles is positive or negative has a significant impact on the difficulty of a problem, but epistasis measures (e.g. epistasis variance or correlation) cannot capture this distinction (Naudts and Kallel 2000). Second, empirical evidence suggests that epistatic interactions can occur at several levels (i.e. there are hierarchical interdependence structures) and this has consequences for the long-term dynamics of the system (Szendro et al. 2013). Thus, what is important is not acknowledging that a problem *has* epistatic interactions, but rather identifying the nature of those interactions.

Paralleling this trend, recent years have seen considerable development when it comes to the study of fitness landscapes (Pitzer and Affenzeller 2012; Malan and Engelbrecht 2013; McClymont 2013; Østman and Adami 2014) with a focus shift from characterizing *problem hardness* (via ruggedness measures) to characterizing *fitness landscapes* in order to determine the appropriate algorithm (McClymont 2013). The shift is due to the fact that fitness landscape analysis allows for “a deeper understanding of a whole problem class” (Pitzer and Affenzeller 2012: p. 3) rather than a specific problem instance. Current research thus aims at identifying relevant features that can describe a fitness landscape and that have known properties with respect to problem solving difficulty (Malan and Engelbrecht 2014). Malan and Engelbrecht

(2014) identify three such features as potentially predictive of performance: ruggedness, neutrality and deceptiveness (Figure 2).

In keeping with the organizational theory approach, where the one-bit hill-climbing algorithm is the dominant search behaviour, in the following we describe how these features can affect the likelihood of finding the optimal solution for a classic one-bit hill-climbing algorithm, but they are not limited to this search heuristic.

[Insert figure 2 around here]

2.2.1. Landscape ruggedness: modality and locality measures

In the previous section we have addressed how prior research has used K or K/N ratios as measures of landscape ruggedness or problem complexity. Indeed, Kauffman (1993) has shown that given the one-bit flip assumption, the ruggedness of an NK landscape is captured by the K parameter. However, although it is clear that highly epistatic landscapes are hard to search, it is not clear how much epistasis “is needed to make a problem difficult” (Jones 1995: p. 134). Thus, in the following, we present a number of alternative measures to capture landscape ruggedness.

In computer science, a frequently used measure of landscape ruggedness is the number of local maxima, or the modality of a landscape. The *modality* of a given landscape is often computed relative to the size of the fitness landscape: the higher the density of such deceiving optima, the more complex the problem, i.e. the higher the likelihood that a solver will be stuck and unable to find the optimal solution. Note that the definition of a distance metric (and implicitly the neighbourhood function) affects the number of local optima, since, by definition, for a problem (X, f) and a *neighbourhood function* M , a solution x^* is called locally optimal *with respect to* M , if

$$f(x) \leq f(x^*) \text{ for all } x \in M(x). \quad (5)$$

The *locality* of a landscape is given by how closely together (with respect to the distance d) solutions with similar fitness values are located (Rothlauf 2011). In general, the lower the distance, the higher the locality and the easier it is to find a global optimum, since better solutions are located closer together (Pitzer and Affenzeller 2012).

Another measure of locality was proposed by Jones and Forrest in 1995 where they proposed a *fitness distance correlation* coefficient.

$$\rho_{FDC} = \frac{C_{fd}}{\sigma(f)\sigma(d_{opt})} \quad (6)$$

Where

$$C_{fd} = \frac{1}{m} \sum_{i=1}^m (f_i - \bar{f})(d_{i,opt} - \bar{d}_{opt}) \quad (7)$$

with $\sigma(f)$ and $\sigma(d_{opt})$ as the standard deviations for the fitness values, respectively the distances to the optimal solution, \bar{f} is the mean value for the fitness function, \bar{d}_{opt} is the mean value for the distance to the optimal solution, f_i the fitness value for solution i and finally $d_{i,opt}$ is the distance of solution i , to the optimal solution x^* .

The fitness-distance correlation coefficient, allows Jones and Forrest (1995) to distinguish between three classes of landscapes:

- a. Straightforward, for $\rho_{FDC} \leq -0.15$. This is the ideal case where the closer a solver gets to the global optimum, the higher the fitness and are roughly correspondent to “smooth” landscapes. NK problems where $K \leq 3$, fall in this category.
- b. Difficult $-0.15 < \rho_{FDC} < 0.15$. There is limited correlation between the fitness difference and the distance to the optimal solution. This makes such optimization problems very hard to solve and renders the search heuristics to random search. According to Jones and Forrest (1995) as K increases over 3, NK landscapes quickly become uncorrelated and ρ_{FDC} approaches 0. These are “rugged” landscapes, with limited or uncorrelated ruggedness.
- c. Misleading $\rho_{FDC} \geq 0.15$. There is an inverse correlation between the fitness difference and the distance to the optimal solution. Thus, the solver is “drawn” *away* from the global optimum. According to Malan and Engelbrecht’s (2014) classifications, these would be “deceptive landscapes”.

2.2.2. Deceptiveness

Recent advances in biology point to the existence of higher-order epistatic interactions which generate multidimensional landscapes (Segre et al. 2005; Kondrashov et al. 2015). These interactions seem to be organized hierarchically in functional modules that interact with each

other (Segre et al. 2005, Jaimovich et al. 2010). This type of interaction structure is reminiscent of the hierarchical structure which has been argued to be an essential feature of organizational problems, at least when it comes to innovation problems (Pelikan et al. 2000; Gavetti 2005). In this context, hierarchy, is seen as the composition of systems out of subsystems with each subsystem in turn having its own hierarchy (Yu et al. 2009), until a certain level of fine grained modularity is achieved. This is a qualitatively different kind of ‘problem complexity’ (as compared to landscape ‘ruggedness’) and the one most likely to be encountered in real-life design problems (Pelikan et al. 2000; Martin 2001; Yu et al. 2009). Note however that hierarchical decomposition and hierarchical interdependence are different from the one-level interdependence, which *is* captured by NK-like landscapes - see also Marengo et al. (2000) for a more detailed account. The latter assumes that the task of solving a problem can be reduced to several low order modules that have intertwined contributions to the overall fitness. In hierarchical problems the interdependence (or interactions) *between* levels is also present and this obstructs single-level decomposition (Pelikan 2005). This description is also in line with Simon’s description of complexity (Simon 1962; 1996).

Such problems are likely to generate *deceptive* landscapes, according to Malan and Engelbrecht’s (2014) classification, since they generate so-called hierarchical traps (Watson and Pollack 1999; Martin 2001; de Jong et al. 2005). The interactions between building blocks make hierarchical problems deceptive (i.e. misleading according to Jones and Forrest 1995) in Hamming space (at lower hierarchical levels), but fully non-deceptive at higher hierarchical levels (Iclanzan and Dumitrescu 2007) – i.e. at higher hierarchical levels (better problem representations), solvers are able to attain better solutions by making incremental changes. In biological terms: the lowest hierarchical level describes “how a mutation in a given gene affects the phenotypic consequence of another mutation and the highest level describes how altered functionality of a given module of genes affects the phenotypic consequence of altered functionality of another module.” (Segre et al. 2005: p. 81)

One example of such function is illustrated in Figure 3.

[Insert figure 3 here]

Figure 3 shows the visualization of a hierarchical problem using a one-bit flip hill-climber (left) and a “chunking” algorithm that was tailored specifically for this problem (see Appendix 1 for

a description). Notice that the chunking landscape is ‘smoother’ than the one-bit-flip landscape. Since the H-XOR⁶ function has $2^{N/2}$ local optima for the one-bit flip hill-climber, the probability that a given point in the one-bit landscape is connected with a path to the global optimum is significantly lower (Figure 4) as compared to the chunking algorithm showing that this problem is ‘deceptive’ for a one-bit-hill-climber but not for an algorithm that can exploit the problem structure.

[Insert figure 4 here]

2.2.3. Neutrality

So far we have only looked at the “smooth vs rugged” distinction and different means of capturing *ruggedness*. A different intuition about how evolutionary dynamics might be influenced by the underlying fitness function comes from models that consider the possibility that some solutions have equal fitness. This was fuelled by developments in molecular biology which have questioned the “rugged landscape” metaphor, in particular its explanation of speciation (Barnett 1997; Gavrilets 1999). This work was largely driven by the neutral theory of molecular evolution and in particular the observation that the majority of mutations at a molecular level do not affect the phenotype (Galván-López et al. 2011). The previous framework assumed that once a population became stuck in a suboptimal peak it could only escape it if the fitness function was changed (e.g. shifting balance theory) or via a long jump. The neutral theory of molecular evolution relies on the conjunction that there must be a series of fitness neutral mutations that would allow even organisms that were currently located in a suboptimal peak to “escape” and undergo further evolution.

In an NKq (Newman and Engelhardt 1998) landscape, for a landscape L we define the *neutral neighbours of x* :

$$M_n(x) = \{x^* \in N(x) \mid f(x) = f(x^*)\}. \quad (8)$$

where fitness contributions are integers drawn from $[0,q)$. The total fitness in this case is given by:

$$f(x) = \frac{1}{N(q-1)} \sum_{i=1}^N x_i \quad (9)$$

⁶ See Appendix for a detailed explanation of HXOR,

Figure 5 shows a 3D reduction for a 1 bit-flip hill-climber as well as the decimal representation used earlier to depict two neutral functions with $N=8$, $K=1$, $q=2$ and $N=8$, $K=1$, $q=3$, respectively.

[Insert figure 5 around here]

The left-hand side of the picture corresponds to the maximally neutral landscape and depicts a very simple, flat fitness landscape, without any local maxima. The right-hand side picture, corresponding to the $q=3$ neutral landscape, has a number of ridges of high fitness states, as well as valleys of low fitness states. Thus, neutral landscapes are not necessarily beneficial for adaptation, since, even for low values of K , a solver is likely to find himself trapped in these “stretches of lethal states” (Franke et al. 2011: p.4).

A number of authors have introduced neutral extensions of the NK landscape and investigate how the new topology might influence the evolutionary processes (e.g. Barnett 1997; Newman and Engelhardt 1998; Lobo et al. 2004). The implementations vary in both details and conclusions regarding the influence of neutrality on the features of the landscape (Geard et al. 2002), but they do conclusively show that neutrality is an important feature that influences search performance and is not captured by traditional measures of ruggedness (Pitzer and Affenzeller 2012), commonly used in NK studies.

Figure 6 provides an illustration of how introducing neutrality can change the dynamics of adaptation. We compare the performance of a one-bit hill climber on an NK landscape ($N=8$, $K=3$) and a NKq landscape ($N=8$, $K=3$, $q=4$ ⁷). Simulations show a higher success ration on the NKq landscape. The success ratio is defined by the ratio between the number of paths and the number of successful paths, and is thus a measure of the likelihood of finding the optimal solution. Simulations were conducted on 1000 different NK (and correspondingly 1000 different NKq landscapes) and the difference was found to be significant ($p=0.04$) with an effect size $r=0.1$.

[Insert figure 6 here]

Thus, if neutrality is a feature that characterizes social science problems, caution should be used when characterizing the fitness landscape by relying on fitness distance correlations or

⁷ We purposely chose a value for q higher than the lowest possible ($Q=2$) which yields maximum neutrality.

K/N ratios (Galván-Lopez and Poli 2006). As Huyen et al. (1996) argue, a small value for the fitness distance correlation (i.e. $-0.15 < \rho_{FDC} < 0.15$) that would normally be connected with a very rugged landscape, is not informative as to the ease/difficulty of finding the global optimum since local optima, when connected, are no longer local (Huynen et al. 1996). This is further explored by Lobo et al. (2004) who conclude that there is an interplay between the ruggedness and neutrality of the landscape. Their simulations suggest that the desirability of neutrality is contingent on the former. For instance, for rugged landscapes, neutrality is beneficial, but for smooth landscapes neutrality just makes adaptation slower.

In consequence, the measures detailed in the previous section do not necessarily capture the relative ease or difficulty an adaptive solver would have on a landscape that does have neutral ridges.

So far we have ignored issues pertaining to the search behaviour, or rather, following the NK literature, we have taken the one-bit flip as a reference. This assumption however isn't as innocuous as it may seem. As such, the features described above (i.e. either metrics such as the number of local optima that a given problem has, or neutrality or deceptiveness) can only be defined with respect to a neighbourhood function M (Pitzer and Affenzeller 2012). We address these concerns in the following section.

3. Search behaviours

Simon (1956) describes agents of increasing intelligence: from the “simple-minded” organism that is driven by a basic stimulus response rule to a more complex, cognitively endowed, actor. Newell and Simon (1976) further introduce the hypothesis that in order for intelligent search to be better than random search, the space of solution has to “exhibit[s] at least some degree of order and pattern” (Newell and Simon 1976: p. 121). Since it is in the interplay between the structure of the problem and the search heuristic that is the focus in these models, it is not just the fitness function that is important, but also the particular search behaviours with which agents are endowed. Furthermore, since search behaviours are not as well defined in the NK model for organizational search as they are in the biological equivalents (i.e. selection, genetic drift, mutation and recombination cf. Huxley 2010) scholars need to build a different empirical foundation in what concerns search behaviours. Still, the most prevalent models in organizational theory seem closer to the ‘mindless particle’ end of the spectrum (cf. Winter et

al. 2007; Csaszar and Levinthal 2015), in what concerns an agent's "ability to store and manipulate symbols" (Newell and Simon 1976: p.115).

Several authors (Jones 1995; Frenken et al. 1999) have cautioned that the NK ruggedness is in fact a property of the landscape and not a property of the task environment, as defined above: the ruggedness is a property of L and not f or its domain. That is, ruggedness, as defined in Kauffman's original model, is assumed to be given by a one bit mutation of the candidate solution (1993). However, as Frenken et al. (1999) point out, the assumption of one-bit flip is of limited relevance in the context of human search behaviours, since such an one-bit conception does not fit human behaviour: human problem solvers are less likely to engage in small, incremental changes. Billinger et al. (2013) e.g. find that the average search distance is above two. As such, the ruggedness of the NK landscape does not allow for intuitions to be formed about problem hardness *in general* (i.e. the likelihood that a solver can find the optimal solution efficiently, provided that local search is not the only or the dominant search heuristic). For example, with a basic hill climbing technique (without random long jumps), a rugged NK landscape ($K > 0$) is reduced to one single peak (likely a local optima) where the solver gets stuck. Thus, even simple heuristics as "hill climbing with long jumps" vastly improve the search process as they are able to cover the entire rugged landscape. In a problem-solving context the heuristic chosen is of outmost importance and the task environment's statistical features (Kauffman 1993) do not preclude the existence of a powerful search heuristic that can in fact resolve to a flat, single-peaked landscape.

3.1. Search in organizational theory

Following the early studies, recent NK model extensions take a more nuanced view on what organizations do in their attempt to find solutions and thus focus on different search behaviours and their performance on landscapes of different complexity (as judged by the $\{N, K\}$ pair).

In the canonical NK model, the search heuristic is inspired by a simple evolutionary mechanism: adaptive mutation. This very simple search algorithm is actually very efficient. Hill-climbing is one of the most powerful domain-general search algorithms (Russel and Norvig 2010). It also provides earlier models with a straight-forward way of implementing bounded rationality assumptions. A solver endowed with such a simple heuristic is clearly light

years away from the all-knowing *homo economicus*, but this simple solver can still solve complex problems, such as identifying the optimal configuration of organizational routines (Levinthal 1997).

Another type of search is where the solver is capable of evaluating all solutions in the one-bit mutation set and chooses the one that maximizes his performance (i.e. ‘offline search’ (Gavetti 2005)). Other scholars have taken into account the fact that the assumption of bounded rationality is not violated if solvers are endowed with more intelligent heuristics (Winter et al. 2007). For example, Gavetti and Levinthal (2000) allow agents to be directed in their search by representations of the search space that are attributed to them a priori, while Winter et al. (2007) assume that agents have exogenously attributed “preferred direction”.

3.2. Alternative search behaviours

Wright (1932) argues that the fundamental mechanism behind speciation must be a non-adaptive one, i.e. it cannot be that hill climbing alone can account for the tremendous variety in species (Gavrilets 1999).

Natural computing was quick to adopt biological mechanisms and adapt: stochastic hill-climbing, first-choice climbing and random-restart hill-climbing were the first natural successors that already showed a marked improvement over the performance of the canonical hill-climber. Although still extremely simple, these algorithms capture fundamental dynamics of adaptation. For example, Figure 7 illustrates the exploration/exploitation trade-off via the variance in performance for a random restart one-bit hill-climber on an NK landscape ($N=8$, $K=3$). Since such a landscape has a number of local optima, a hill-climber with zero probability of restarting would quickly climb up the nearest peak and the search would stop. The probability of identifying the optimal solution is strictly dependent on the number and size of basins of attraction⁸ for these local optima. For large numbers of local optima and large basins of attraction, the likelihood that the agent finds himself in the vicinity of the global optimum decreases and so does his probability of success. As the probability of restart increases, the solver also increases his chances of “landing” in the right part of the landscape. Evidently, a high probability of restart (in this case $p>0.2$) decreases performance since the agent engages

⁸ The areas around local optima that lead a hill-climbing algorithm directly to the local peak.

excessively in exploration (sampling the landscape) and not enough in exploitation (hill-climbing).

[Insert Figure 7 around here]

Genetic algorithms (Holland 1992) are another class of algorithms that is inspired by evolutionary mechanisms. The major difference between the former and the latter was the use of genetic operators: in addition to selection and mutation, genetic algorithms rely on gene recombination as well (Affenzeller et al. 2009). Genetic algorithms are widely used in practice (e.g. Matthey et al. 2007) as are a number of various other computing tools ranging from fuzzy logic and belief calculus to machine learning like inductive logic programming (Konar 1999).

However, as shown by Wolpert and Macready (1999), an algorithm's average performance is determined by how much knowledge regarding the optimization function is incorporated into the search heuristic (the "No Free Lunch Theorem" for optimization). Thus, computer science has moved away from general purpose evolutionary algorithms with their limited knowledge of the problem space to algorithms that are designed specifically for the problem at hand. To illustrate this point, we compare the performance of three different algorithms on a hierarchical landscape (Figure 8): one-bit flip, and the same "chunking" algorithm that relies on 26 operations that are derived by taking into account the particularities of the hierarchical problem (H-XOR, $N=8$) and random search.

[Insert figure 8 here]

We show how the chunking algorithm significantly outperforms both the hill-climber and a random search. Computational experiments also endorse this: algorithms that embed these principles out-perform traditional recombination (genetic algorithms) or local search, since a solver that uses an inappropriate problem decomposition effectively generates a rugged landscape (Cioffi-Revilla et al. 2012).

Memory in computational agents

Finally, most computational approaches discussed so far rely on agents which do not have memory⁹. The process of problem solving they describe is path dependent and solvers attempt

⁹ Note this is different from Gavetti (2005) who allows agents to have a 'cognitive memory', that is a set of different partial representations of the landscape.

to improve their current best performance, but, for example, in an eight-bit NK problem, at each time-step, a solver still has 255 (2^8-1) possible combinations to choose from. Naturally, the search heuristic implemented restricts the set of available moves, but the assumption is still strong and unjustified. The equivalent claim: ‘draws from an urn with or without replacement yield identical probabilities of success’ highlights the problem with disregarding memory.

[Insert figure 9 around here]

Figure 9 illustrates how the performance of a hill-climber changes as a function of the number of solutions he remembers. Note that memory is not ‘universal a recipe for success’, since remembering past success restricts the following moves that the hill-climber can make, which in turn means that the agent is more easily stuck in a suboptimal solution.

It is beyond the scope of this paper to give a comprehensive overview of the history of natural computing. It should be noted that there is a great variety of increasingly sophisticated algorithms for solving optimization problems: for instance, in addition to algorithms inspired by evolutionary theories, computer science has also developed methods inspired by the human nervous system (artificial neural networks), the collective behaviour of groups of organisms (e.g. particle swarm optimization) or quantum physics (Rozenberg et al. 2011).

Some of these algorithms are directly related to the previous discussion on features of fitness landscapes and rely on different measures of search performance such as e.g. search dispersion to adapt the search as it progresses (Maier et al. 2014). Others rely on more classical mappings of the search space, such as decision trees (Huys et al. 2012) or Bayesian models (Pelikan 2005). These mappings are continuously adapted during the search process, using the information the agent gathers by interacting with the environment.

To sum up, fuelled by the development of Artificial Intelligence, in computer science literature there has been a shift in focus from agents that can be thought of as ‘mindless particles’, to ‘smarter’ agents, which, without being perfectly rational, are capable of observing or partially observing their environment and constructing “beliefs” that allow them to generate subsequent moves (Russel and Norvig 2010).

4. Discussion

Based on the review of organizational theory, computer science and biology perspectives, in the following we try to outline future potential avenues for expanding the models as well as discuss potential implications. Although we acknowledge that models are to be thought of as useful simplifications of reality, in this section we highlight how previously made simplifications and NK legacy elements actually restrict our ability to rely on these models to further our understanding about human/organizational problem-solving dynamics.

Modelling landscapes and complexity

Irrespective of the NK model's biological origins, there are no clear empirical specifications or constraints about how fitness landscapes should be conceptualized. In fact, Kauffman acknowledges that the contributions of individual alleles' are drawn randomly from a uniform distribution, since the exact contribution of a gene to an organism's fitness is not known (Kauffman 1993). The use of the NK canonical function has been justified in a similar manner by the fact that modellers are interested in how the evolutionary process "typically" unfolds (Rivkin 2000: p.828). However, as Jones and Forrest (1995) and Frenken et al. (1999) show, for $K > 3$ NK landscapes become quickly uncorrelated. Given that organizational problems and innovation problems are rarely the real-life correspondent of either completely smooth ($K=0$) or completely uncorrelated landscapes, this questions whether NK functions are meaningful in modelling the dynamics of problem solving systems.

While a mapping between the various components of a given solution and its performance (e.g. the modules of the educational app described in the introduction and its computational performance) is not trivial to make in an organizational setting, it can be even harder to see how it can be argued that such mapping is random and more importantly how the NK captures the interdependence structure of a *typical organizational problem*.

Indeed a number of scholars have attempted more meaningful extensions of the NK model (Siggelkow and Rivkin 2003; Ethiraj and Levinthal 2004; Ethiraj et al. 2008). For example, acknowledging the problematic nature of the random interdependence structure in NK functions, Ethiraj and Levinthal (2004) impose a block near-modular design on top of the NK matrix and then allow solvers to optimize inside the modules as well as resort to recombinative practices. They show that erring on the side of too much decomposition is detrimental to search efficacy. However, Watson and Pollack (2005) caution that according to Simon's definition of modularity, it is only on the *short-term* that modules are quasi-independent and on the contrary,

long-term dynamics should assume stronger inter-module interdependence. This kind of interdependence, they argue, is not captured by structural (as opposed to functional) models of modularity.

One further limitation of these extensions is that the properties of these pseudo-NK landscapes are not as established as results for the canonical NK, so it is not obvious whether assumptions about the structure of the problem or assumptions about the search behaviours are driving the simulation results. For instance, recent work shows that imposing a block structure on the NK interactions qualitatively changes the structure of the landscape by actually *diminishing* the number of evolutionary paths towards the global maximum, under the SSWM¹⁰ condition (Schmiegelt and Krug 2014), while Hebborn et al. (2008) show that imposing a scale free structure on an NK landscape leads to longer adaptive walks and more clustering of optima in the landscape. In addition, all the extensions referenced in this study (e.g. Gavetti and Levinthal 2000; Siggelkow and Rivkin 2003; Ethiraj and Levinthal 2004; Gavetti 2005; Ethiraj et al. 2008) have as a basis the canonical form for the NK fitness function, which averages across individual fitness contributions. This, as Mckelvey et al. (2013) show, inevitably generates the same result: with the increase in N and K, the value for the fitness function converges towards the mean of the uniform distribution (0.5) and this skews the interpretation of the simulation findings. The particular way the fitness function is generated in the NK model is also what Szendro et al. (2013) argue makes such models less amendable to being able to capture different levels of epistatic effects.

In biology, the empirical evidence towards the existence of multi-modal landscapes with numerous epistatic interactions continues to increase (Østman and Adami 2014) with scholars inquiring if it is reasonable to assume that adaption is taking place on a highly uncorrelated landscapes and if it is meaningful to assume there are no ‘neutral ridges’ (Gavrilets 1999) or hierarchical interdependence (Segre et al. 2005). With few exceptions (e.g. Fleming and Sorenson 2001; 2004), a similar, empirically grounded, discussion about how can we create meaningful landscapes for organizational problems seems to be missing.

Paradoxically, we argue that one way forward, as suggested by current developments in computer science, is to revert to what Wright (1932) and Kauffman (1993) originally proposed: relying on a fitness landscape to first acquire a “rough” image of a problem class, instead of

¹⁰ Strong selection, weak mutation.

investigating specific instances (Pitzer and Affenzeller 2012, Malan and Engelbrecht 2014). Investigations into the topology of the fitness landscape would allow for a better understanding of the dynamic processes of adaptation in a similar vein to previous considerations: e.g. rugged landscapes are likely to trap solvers in suboptimal peaks, deceptive landscapes are likely to attract solvers towards suboptimal optima etc. Alongside the few methods introduced in this paper, a number of methods for landscape analysis have been comprehensively developed in this literature.

Modelling search behaviours

While previous paragraphs discuss issues related to how we conceptualize landscapes, we now turn to search behaviours. We have already discussed how several studies have questioned and attempted to expand the human problem solving models of search beyond its biological origins. Their approach is largely driven by theoretical concerns, but recent research is attempting to do the same driven by empirical results. This empirically driven approach allows for a better specification of search behaviours which in turn results in better models. The hill-climber most often used in management science is stochastic-restart hill climbing, rather than the canonical hill-climber. Before Billinger et al. (2013) few if any scholars spend time explaining some of the subsequent (seemingly innocuous) modelling decisions. However, as shown earlier (Figure 7), the performance of a hill-climber with random restarts differs significantly from the performance of a hill-climber without random restarts.

These results mirror a previous study conducted by Mason and Watts (2012). By comparing the performance of actual solvers and computational agents, Mason and Watts (2012) show that heterogeneity in terms of search behaviours has the potential to greatly influence the outcome of the search process. Whether relying on constructs such as attention control (Laureiro-Martínez et al. 2015), intelligence (Steyvers et al. 2009) or cognitive styles (Kirton 1976), there is empirical evidence that there is a great heterogeneity when it comes to human search behaviours, but such heterogeneity is rarely taken into account in modelling approaches (Miller and Page 2009). This is also endorsed by empirical results which suggest that humans are capable of solving hard computational problems (Carruthers and Stege 2013), evidence to the fact that humans have far more sophisticated search strategies.

We have shown how assumptions about memory influence performance (Figure 8). Memory is not only important when it comes to restricting the search space (e.g. Gavetti 2005), but also as one way of advancing more plausible assumptions into modelling human problem-solving. Human problem solvers are not guided in their search merely by immediate feedback, but also by representations of the problem they form over time, via accumulated experience (Doll et al. 2012). The modelling community has struggled to capture this via ‘preferred direction’ (Winter et al. 2007) or partial representations of the solution (Gavetti and Levinthal 2000, Gavetti 2005), but the attempt is complicated by the fact that one assumes an agent has a problem representation *before* solving a problem. Where could this representation come from?

One answer lies in the fact that the mental models or representations that solvers use to guide their search are themselves adaptive and subject to reinforcement learning (Miller and Page 2009). It is then not only in the generation and evaluation of *solutions* that feedback loops are important (Bonabeau 2009), but the same mechanism can account for the emergence and evolution of problem representations as solvers engage in “imagining future events” (Schacter et al. 2007: p. 659). This approach insures that solvers do not have to start with an ex-ante map of the landscape, but gradually formulate it, or as Simon and Newell describe it (1972): they incorporate knowledge into their search heuristic. The computational solutions involve the implementation of machine learning techniques (Rand 2006) which can range from Bayesian algorithms (Pelikan et al. 2003), decision trees (Huys et al. 2012) and more recently deep learning (Mnih et al. 2015). Irrespective of the details, these methods have the potential to bring forward a middle way in modelling problem solving that is in keeping with the bounded rationality assumption, but at the same time allows more than “ant-like” behaviour (Winter et al. 2007) consistent with some of the theories regarding human cognition (Le 2013). The underlying idea is that while engaging in problem solving, as more information is available human solvers are able to detect and abstract essential features, in the same way visual pattern recognition works (Roland and Gulyás 1995). As a result of that, in structured environments subsequent variations or mutations are not random, but closer to what biologists call “facilitated variation”: generated new solutions are potentially *useful* (Parter et al. 2008: p.2). This however, is not possible in NK landscapes where fitness contributions are drawn randomly from an underlying distribution since, as Watson et al. (2011) argue the environment has to display a certain degree of regularity that the agents can exploit.

Conclusion

While the simulation approach has gained attention in high-status outlets within organizational theory, we acknowledge and address the sceptical concerns still being raised about theoretical assumptions (Fioretti 2013) and the weak empirical grounding of these assumptions (Chang and Harrington 2006; Mason and Watts 2012). Much like in the original biological setup, the organizational literature has had simplistic assumptions about agent behaviour, embedded in a relatively undefined fitness landscape (Ganco and Hoetker 2008, McKelvey et al. 2013). However, unlike micro-biology, where evolutionary forces are well known (Huxley 2010), defining human search behaviours in this conceptual framework turns out to be elusive: how do we model and define what constitutes intelligent boundedly rational behaviour; for instance, do agents have memory and how good are they at interpreting the landscape?

Given our limited understanding about the genotype-phenotype mapping in a technological setting (Solée et al. 2013), we suggest that the focus should not be on the statistical features of the landscape to be searched under the one-bit flip condition, but on how the interplay of search behaviours and the different natures of interdependence structures translates into problem solving performance. Only then we can focus on how the search can be best organized in such a way that solvers effortlessly find themselves in the vicinity of the optimal solution (cf. Felin and Zenger 2014).

Additionally, we argue that any model of organizational learning should allow for more plausible (and if possible, empirically validated) assumptions regarding learning and expertise. For un-informed solvers, the fitness landscape will be extremely large and rugged, as they have to deal with a seemingly unconstrained search space (Zhang and Norman 1994). However, a problem representation works by effectively constraining the search space, generating a different set of possible solutions. Early investigations in the use of problem representations (Kotovsky and Simon 1990) show that knowledge about the landscape changes the structure of the landscape: “<the easy problem> problem” (Winter 2004)¹¹. These ideas are not foreign to organizational literature which has a long tradition of looking at managerial decisions through the lenses of cognitive frames (or schemas) which seem to be the primary source of difficulty for organizations in turbulent environments (Kaplan 2008; Bingham and Kahl 2013). Still, most modelling approaches do not take into account this perspective about how cognitive frames change the search behavior *and* implicitly the landscape.

¹¹ Learning and practice as well as context influence whether a problem is perceived as “easy” by a given solver.

We argue that moving away from “armchair speculations” Simon (1982) regarding human search behaviour and the nature of the problem is essential in these settings as seemingly innocuous assumptions can drastically change the problem solving performance. We further identify two potential avenues for future research: focusing on different landscape features and creating “smarter” agents by relying on the recent developments in computer science, both which, we argue, should be endorsed by empirical calibration and validation.

References

- Affenzeller, M., S. Wagner, et al. (2009). Genetic algorithms and genetic programming: modern concepts and practical applications, Crc Press.
- Altenberg, L. (1997). "NK fitness landscapes." Handbook of evolutionary computation 7: 5-B2.
- Barnett, L. (1997). "TANGLED WEBS--Evolutionary Dynamics on Fitness Landscapes with Neutrality."
- Barnett L. (1998). “Ruggedness and neutrality-the NKp family of fitness landscapes”. In Artificial Life VI: Proceedings of the sixth international conference on Artificial life Aug 28 (pp. 18-27).
- Billinger, S., N. Stieglitz, et al. (2013). "Search on rugged landscapes: An experimental study." Organization Science 25(1): 93-108.
- Bingham, C. B. and S. J. Kahl (2013). "The process of schema emergence: Assimilation, deconstruction, unitization and the plurality of analogies." Academy of Management Journal 56(1): 14-34.
- Bonabeau, E. (2009). "Decisions 2.0: The power of collective intelligence." MIT Sloan Management Review 50(2): 45-52.
- Baumann, O. and N. Siggelkow (2013). "Dealing with Complexity: Integrated vs. Chunky Search Processes." Organization Science 24(1): 116-132.
- Baumann, Oliver. "Models of complex adaptive systems in strategy and organization research." Mind & Society (2015): 1-15.
- Burton, R. M. and B. Obel (1980). "A Computer Simulation of the M-Form Hypothesis." Administrative Science Quarterly 25(3): 457-466.
- Carruthers, S. and U. Stege (2013). "On evaluating human problem solving of computationally hard problems." The Journal of Problem Solving 5(2): 4.

- Chang, M. H., & Harrington, J. E. (2006). Agent-based models of organizations. Handbook of computational economics, 2, 1273-1337.
- Cioffi-revilla, C., K. De Jong, et al. (2012). "Evolutionary computation and agent-based modeling: biologically-inspired approaches for understanding complex social systems." Computational and Mathematical Organization Theory **18**(3): 356-373.
- Cohen, M. D., J. G. March, et al. (1972). "A garbage can model of organizational choice." Administrative science quarterly: 1-25.
- Csaszar, F. A. and D. A. Levinthal (2015). "Mental representation and the discovery of new strategies." Strategic Management Journal **Forthcoming**.
- de Jong, E. D., R. A. Watson, et al. (2005). A generator for hierarchical problems. Proceedings of the 2005 workshops on Genetic and evolutionary computation, ACM.
- Doll, B. B., D. A. Simon, et al. (2012). "The ubiquity of model-based reinforcement learning." Current opinion in neurobiology **22**(6): 1075-1081.
- Ethiraj, S. K. and D. Levinthal (2004). "Modularity and innovation in complex systems." Management Science **50**(2): 159-173.
- Felin, T. and T. Zenger (2014). "Closed or open innovation? Problem solving and the governance choice". Research Policy **43**(5): 914-925.
- Fioretti, G. (2013). "Agent-based simulation models in organization science." Organizational Research Methods **16**(2): 227-242.
- Fleming, L., & Sorenson, O. (2001). "Technology as a complex adaptive system: evidence from patent data". Research Policy, 30(7), 1019-1039.
- Fleming, L. and O. Sorenson (2004). "Science as a map in technological search." Strategic Management Journal **25**(8-9): 909-928.
- Frenken, K., L. Marengo, et al. (1999). Interdependencies, nearly-decomposability and adaptation. Computational techniques for modelling learning in economics, Springer: 145-165.
- Galván-López, E., R. Poli, et al. (2011). "Neutrality in evolutionary algorithms... What do we know?" Evolving Systems **2**(3): 145-163.
- Galván-López E, Poli R. (2006) An empirical investigation of how and why neutrality affects evolutionary search. In Proceedings of the 8th annual conference on Genetic and evolutionary computation 2006 Jul 8: 1149-1156. ACM.

- Ganco, M. and G. Hoetker (2009). "NK modeling methodology in the strategy literature: bounded search on a rugged landscape." Research methodology in strategy and management **5**(2009): 237-268.
- Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly **45**(1): 113-137.
- Gavetti, G. and D. Levinthal (2001). "Bringing Cognition Back In and Moving Forward." Journal of Management and Governance **5**(3): 213-216.
- Gavetti, G. (2005). Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. Organization Science, 16(6), 599-617.
- Gavrilets, S. (1999). "Evolution and speciation in a hyperspace: the roles of neutrality, selection, mutation and random drift." Evolution: 1.
- Geard, N., J. Wiles, et al. (2002). A comparison of neutral landscapes-NK, NKp and NKq. Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on, IEEE.
- Golberg, David E. (1989) Genetic algorithms in search, optimization, and machine learning. Addison Wesley, Reading (MA)
- Goldberg, D. (2002). The Design of Innovation (Genetic Algorithms and Evolutionary Computation), Springer New York.
- Hargadon, A. and R. I. Sutton (1997). "Technology Brokering and Innovation in a Product Development Firm." Administrative Science Quarterly **42**(4): 716-749.
- He, J., C. Reeves, et al. (2007). "A note on problem difficulty measures in black-box optimization: Classification, realizations and predictability." Evolutionary Computation **15**(4): 435-443.
- Hebbron, T., S. Bullock, et al. (2008). NKalpha: Non-uniform epistatic interactions in an extended NK model. Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems, MIT Press, Cambridge, MA: 234-241
- Hedström, P. and G. Manzo (2015). "Recent Trends in Agent-based Computational Research: A Brief Introduction." Sociological Methods & Research **44**(2): 179-185.
- Holland, J. H. (1992). Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence, MIT press.
- Huxley, J.S.: Evolution (2010): The modern synthesis, MIT Press
- Huynen, Martijn A., Peter F. Stadler, and Walter Fontana. "Smoothness within ruggedness: the role of neutrality in adaptation." Proceedings of the National Academy of Sciences **93**.1 (1996): 397-401.

- Huys, Q., N. Eshel, et al. (2012). "Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees." PLoS Comput. Biol **8**(3): e1002410.
- Iclanzan, D. and D. Dumitrescu (2007). Overcoming hierarchical difficulty by hill-climbing the building block structure. Proceedings of the 9th annual conference on Genetic and evolutionary computation, ACM.
- Jaimovich, A., Rinott, R., Schuldiner, M., Margalit, H., & Friedman, N. (2010). Modularity and directionality in genetic interaction maps. Bioinformatics, 26(12), i228-i236.
- Jones, T. (1995). Evolutionary algorithms, fitness landscapes and search, Citeseer.
- Jones, T. and S. Forrest (1995). Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms. ICGA.
- Kaplan, S. (2008). "Framing Contests: Strategy Making Under Uncertainty." Organization Science **19**(5): 729-752.
- Karaboga, D. (2005). An idea based on honey bee swarm for numerical optimization, Technical report-tr06, Erciyes university, engineering faculty, computer engineering department.
- Katada, Y. and K. Ohkura (2006). Estimating the Degree of Neutrality in Fitness Landscapes by the Nei's Standard Genetic Distance-An Application to Evolutionary Robotics. Evolutionary Computation, 2006. CEC 2006. IEEE Congress on, IEEE.
- Kauffman, S. A. (1993). The origins of order: self-organization and selection in evolution. New York, Oxford University Press.
- Kirton, M. (1976). Adaptors and innovators: A description and measure. Journal of applied psychology, 61(5), 622.
- Knudsen, T., and D. A. Levinthal. "Two faces of search: Alternative generation and alternative evaluation." Organization Science 18.1 (2007): 39-54.
- Konar, A. (1999). Artificial intelligence and soft computing: behavioral and cognitive modeling of the human brain, CRC press.
- Kondrashov, D. A., & Kondrashov, F. A. (2015). Topological features of rugged fitness landscapes in sequence space. Trends in Genetics, 31(1), 24-33.
- Kotovskiy, K. and H. A. Simon (1990). "What makes some problems really hard: Explorations in the problem space of difficulty." Cognitive Psychology **22**(2): 143-183.
- Laureiro-Martínez, D., Brusoni, S., Canessa, N., & Zollo, M. (2015). Understanding the exploration–exploitation dilemma: An fMRI study of attention control and decision-making performance. Strategic Management Journal, 36(3), 319-338.

- Le, Q. V. (2013). Building high-level features using large scale unsupervised learning. Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on, IEEE.
- Levinthal, D. A. (1997). "Adaptation on rugged landscapes." Management Science **43**(7): 934-950.
- Lobo, J., J. H. Miller, et al. (2004). "Neutrality in technological landscapes." Sante Fe Institute Working Paper.
- Maier, H., Z. Kapelan, et al. (2014). "Evolutionary algorithms and other metaheuristics in water resources: current status, research challenges and future directions." Environmental Modelling & Software **62**: 271-299.
- Malan, K. M. and A. P. Engelbrecht (2013). "A survey of techniques for characterising fitness landscapes and some possible ways forward." Information Sciences **241**: 148-163.
- Malan, K. M. and A. P. Engelbrecht (2014). Fitness landscape analysis for metaheuristic performance prediction. Recent advances in the theory and application of fitness landscapes, Springer: 103-132.
- Marengo, L. and G. Dosi (2005). "Division of labor, organizational coordination and market mechanisms in collective problem-solving." Journal of Economic Behavior & Organization **58**(2): 303-326.
- Marengo, L., G. Dosi, et al. (2000). "The structure of problem-solving knowledge and the structure of organizations." Industrial and Corporate Change **9**(4): 757-788.
- Martignoni, D., Menon, A., & Siggelkow, N. (2015). "Consequences of Misspecified Mental Models: Contrasting Effects and the Role of Cognitive Fit". Strategic Management Journal.
- Martin, W. (2001). Foundations of Genetic Algorithms 2001 (Foga 6), Morgan Kaufmann.
- Mason, A. (1995). Non-linearity measure of a problem's crossover suitability. Proceedings of the IEEE Conference on Evolutionary Computation, IEEE.
- Mason, W. and D. J. Watts (2012). "Collective problem solving in networks." Proceedings of the National Academy of Sciences.
- Mata, R. and B. von Helversen (2015). "Search and the Aging Mind: The Promise and Limits of the Cognitive Control Hypothesis of Age Differences in Search." Topics in Cognitive Science **7**(3): 416-427.
- Matthey, D., J. Wang, et al. (2007). "Enhanced bonding of gold nanoparticles on oxidized TiO₂ (110)." Science **315**(5819): 1692-1696.

- McClymont, K. (2013). Recent advances in problem understanding: Changes in the landscape a year on. Proceedings of the 15th annual conference companion on Genetic and evolutionary computation, ACM.
- Mckelvey, B., Li, M., Xu, H., & Vidgen, R. (2013). Re-thinking Kauffman's NK fitness landscape: From artifact and groupthink to weak-tie effects. Human Systems Management, 32(1), 17-42.
- Miller, J. H. and S. E. Page (2009). Complex adaptive systems: an introduction to computational models of social life: an introduction to computational models of social life, Princeton university press.
- Miner, A. S., P. Bassof, et al. (2001). "Organizational improvisation and learning: A field study." Administrative science quarterly 46(2): 304-337.
- Mnih, V., K. Kavukcuoglu, et al. (2015). "Human-level control through deep reinforcement learning." Nature 518(7540): 529-533.
- Naudts, B. and A. Verschoren (1999). "Epistasis and deceptivity." Simon Stevin-Bulletin of the Belgian Mathematical Society 6: 147-154.
- Naudts, B., & Kallel, L. (2000). "Comparison of summary statistics of fitness landscapes." IEEE Trans. Evol. Comp, 4, 1-15.
- Newell, A. and H. A. Simon (1972). Human problem solving. Englewood Cliffs, N.J., Prentice-Hall.
- Newell, A. and H. A. Simon (1976). "Computer science as empirical inquiry: Symbols and search." Communications of the ACM 19(3): 113-126.
- Newman, M. and R. Engelhardt (1998). "Effects of selective neutrality on the evolution of molecular species." Proceedings of the Royal Society of London B: Biological Sciences 265(1403): 1333-1338.
- Nickerson, J. A. and T. R. Zenger (2004). "A knowledge-based theory of the firm—The problem-solving perspective." Organization science 15(6): 617-632.
- Østman, B. and C. Adami (2014). Predicting evolution and visualizing high-dimensional fitness landscapes. Recent Advances in the Theory and Application of Fitness Landscapes, Springer: 509-526.
- Østman, B., A. Hintze, et al. (2010). "Critical properties of complex fitness landscapes." arXiv preprint arXiv:1006.2908.

- Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated variation: how evolution learns from past environments to generalize to new environments. PLoS Comput Biol, *4*(11), e1000206.
- Pelikan, M. (2005). Bayesian Optimization Algorithm. Hierarchical Bayesian Optimization Algorithm, Springer: 31-48.
- Pelikan, M., D. E. Goldberg, et al. (2000). Hierarchical Problem Solving and the Bayesian Optimization Algorithm. GECCO.
- Pelikan, M., D. E. Goldberg, et al. (2003). Hierarchical Bayesian optimization algorithm: Toward a new generation of evolutionary algorithms. New York, Ieee.
- Pitzer, E. and M. Affenzeller (2012). A comprehensive survey on fitness landscape analysis. Recent Advances in Intelligent Engineering Systems, Springer: 161-191.
- Rand, W. (2006). Machine learning meets agent-based modeling: When not to go to a bar. the Proceedings of Agent.
- Reidys, C. M. and P. F. Stadler (2002). "Combinatorial landscapes." SIAM review *44*(1): 3-54.
- Rivkin, J. W. (2000). "Imitation of complex strategies." Management Science *46*(6): 824-844.
- Rivkin, J. W. and N. Siggelkow (2003). "Balancing search and stability: Interdependencies among elements of organizational design." Management Science *49*(3): 290-311.
- Roland, P. and B. Gulyás (1995). "Visual memory, visual imagery, and visual recognition of large field patterns by the human brain: functional anatomy by positron emission tomography." Cerebral Cortex *5*(1): 79-93.
- Romero, Philip A., and Frances H. Arnold. "Exploring protein fitness landscapes by directed evolution." *Nature Reviews Molecular Cell Biology* 10.12 (2009): 866-876.
- Rothlauf, F. (2011). Design of modern heuristics: principles and application, Springer Science & Business Media.
- Rozenberg, G., T. Bäck, et al. (2011). Handbook of natural computing, Springer Publishing Company, Incorporated.
- Russel, S. and P. Norvig (2010). "Artificial Intelligence: A Modern Approach (3d edition)." EUA: Prentice Hall.
- Schacter, D. L., D. R. Addis, et al. (2007). "Remembering the past to imagine the future: the prospective brain." Nature Reviews Neuroscience *8*(9): 657-661.
- Schmiegelt, B. and J. Krug (2014). "Evolutionary accessibility of modular fitness landscapes." Journal of Statistical Physics *154*(1-2): 334-355.

- Segre, D., DeLuna, A., Church, G.M. and Kishony, R., 2005. Modular epistasis in yeast metabolism. *Nature genetics*, 37(1), pp.77-83.
- Siggelkow, N. and D. A. Levinthal (2003). "Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation." *Organization Science* 14(6): 650-669.
- Shoham, Y., R. Powers, et al. (2007). "If multi-agent learning is the answer, what is the question?" *Artificial Intelligence* 171(7): 365-377.
- Simon, H. A. (1956). "Rational choice and the structure of the environment." *Psychological review* 63(2): 129.
- Simon, H. A. (1962). "The architecture of complexity." *Proceedings of the American philosophical society* 106(6): 467-482.
- Simon, H. A. (1974). "The structure of ill structured problems." *Artificial intelligence* 4(3): 181-201.
- Simon, H. A. (1996). *The sciences of the artificial*. Cambridge, Mass., MIT Press.
- Simon, H. A. and A. Newell (1971). "Human problem solving: The state of the theory in 1970." *American Psychologist* 26(2): 145.
- Simon, H. A., M. Augier, and J. G. March. *Models of a man: Essays in memory of Herbert A. Simon*. MIT Press, 2004.
- Smith, T., P. Husbands and M. O'Shea. (2002). "Fitness landscapes and evolvability." *Evolutionary computation* 10.1 : 1-34.
- Solée, R. V., S. Valverde, et al. (2013). "The evolutionary ecology of technological innovations." *Complexity* 18(4): 15-27.
- Sorenson, O. (2002). Interorganizational complexity and computation. *Companion to organizations*, 664-685.
- Sørensen, J. J. W., M. K. Pedersen, et al. (2015). "Exploring the Quantum Speed Limit with Computer Games." *arXiv preprint arXiv:1506.09091*.
- Szendro, I. G., Schenk, M. F., Franke, J., Krug, J., & De Visser, J. A. G. (2013). Quantitative analyses of empirical fitness landscapes. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01), P01005.
- Steyvers, M., Lee, M. D., & Wagenmakers, E. J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3), 168-179.

- Todd, P. M., T. T. Hills, et al. (2012). Cognitive search: Evolution, algorithms, and the brain, MIT press.
- Thrun, S. (1995). "Learning to play the game of chess." Advances in neural information processing systems **7**.
- Valente, M. (2008). Pseudo-NK: An enhanced model of complexity, LEM Working Paper Series.
- Van der Maaten, L. and G. Hinton (2008). "Visualizing data using t-SNE." Journal of Machine Learning Research **9**(2579-2605): 85.
- Vuculescu, O. (2015). Searching far away from the lamp-post: an agent based model, presented at Organizing Crowds and Innovation Conference, Oxford 2015
- Watson, R. A. and J. B. Pollack (1999). Hierarchically consistent test problems for genetic algorithms. Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on, IEEE.
- Watson, R. A., & Pollack, J. B. (2005). Modular interdependency in complex dynamical systems. Artificial Life, 11(4), 445-457.
- Watson, R. A., Buckley, C. L., & Mills, R. (2011). Optimization in "self-modeling" complex adaptive systems. Complexity, 16(5), 17-26.
- Weinberger, E. D. (1996). Np completeness of kauffman's nk model, a tuneable rugged fitness landscape.
- Weise, T., M. Zapf, et al. (2009). Why is optimization difficult? Nature-Inspired Algorithms for Optimisation, Springer: 1-50.
- Winter, S, G. "The easy problem problem" in Simon et al 2004 *Models of a man: Essays in memory of Herbert A. Simon*. MITPress, 2004
- Winter, S. G., G. Cattani, et al. (2007). "The value of moderate obsession: Insights from a new model of organizational search." Organization Science **18**(3): 403-419.
- Wolpert, David H., and William G. Macready. "No free lunch theorems for optimization." Evolutionary Computation, IEEE Transactions on 1.1 (1997): 67-82.
- Wood, S., J. Busemeyer, et al. (2005). "Older adults as adaptive decision makers: evidence from the Iowa Gambling Task." Psychology and aging **20**(2): 220.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution, na.
- Yu, T. L., D. E. Goldberg, et al. (2009). "Dependency Structure Matrix, Genetic Algorithms, and Effective Recombination." Evolutionary Computation **17**(4): 595-U593.

Zhang, J. and D. A. Norman (1994). "Representations in distributed cognitive tasks." Cognitive science **18**(1): 87-122.

Zott, C. (2002). "When Adaptation Fails An Agent-based Explanation of Inefficient Bargaining Under Private Information." Journal of Conflict Resolution **46**(6): 727-753.

Figures and tables

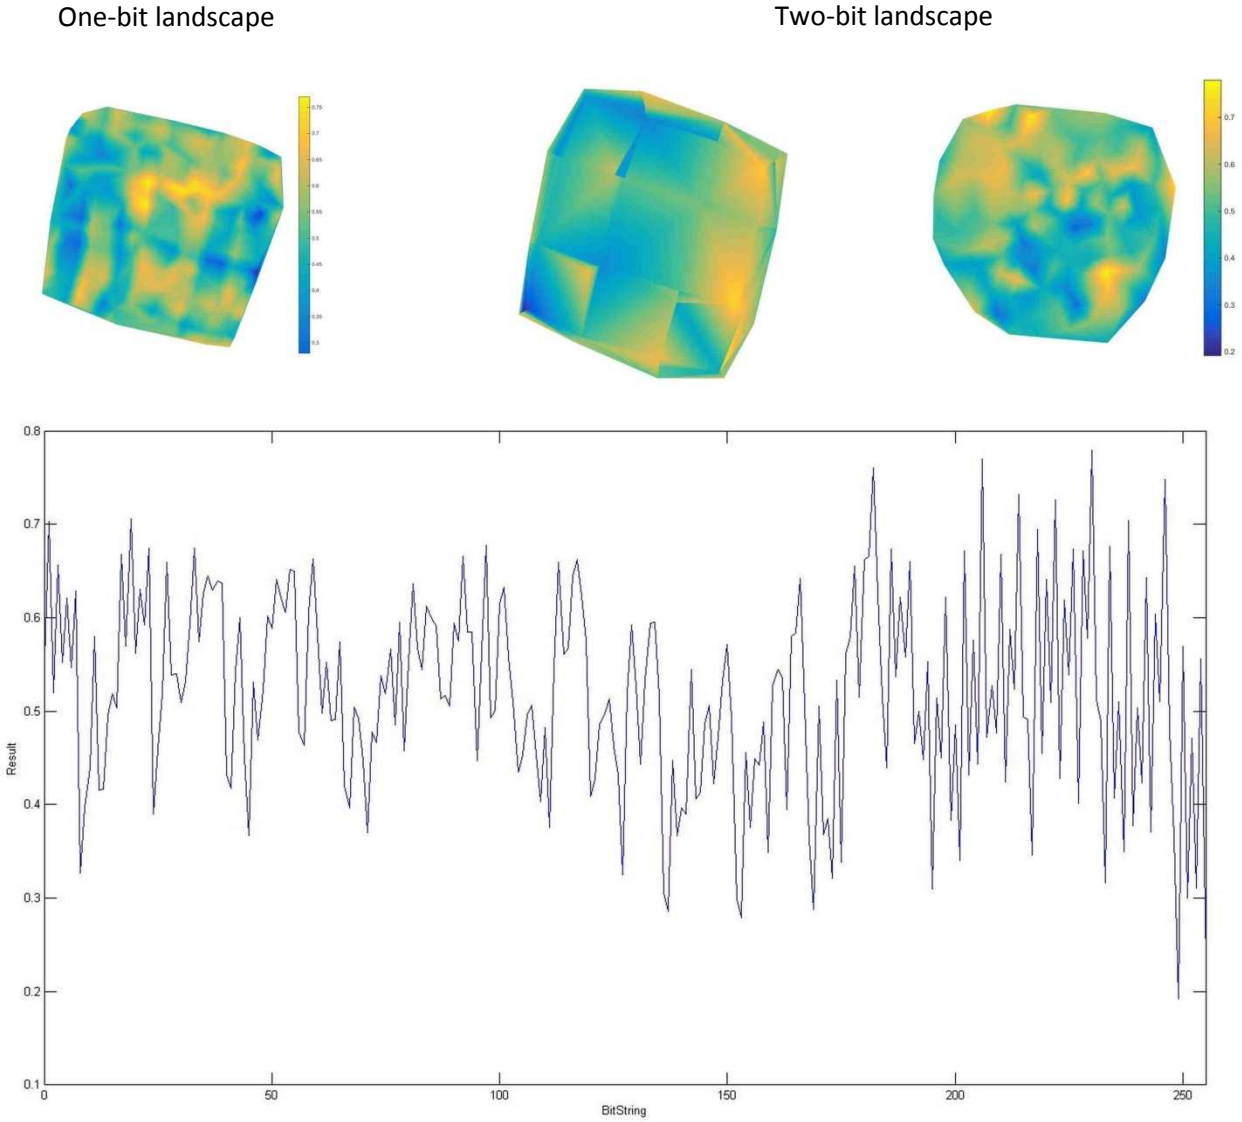


Figure 1. The same function (N=8, K=3) mapped with three different definitions of neighbourhood.

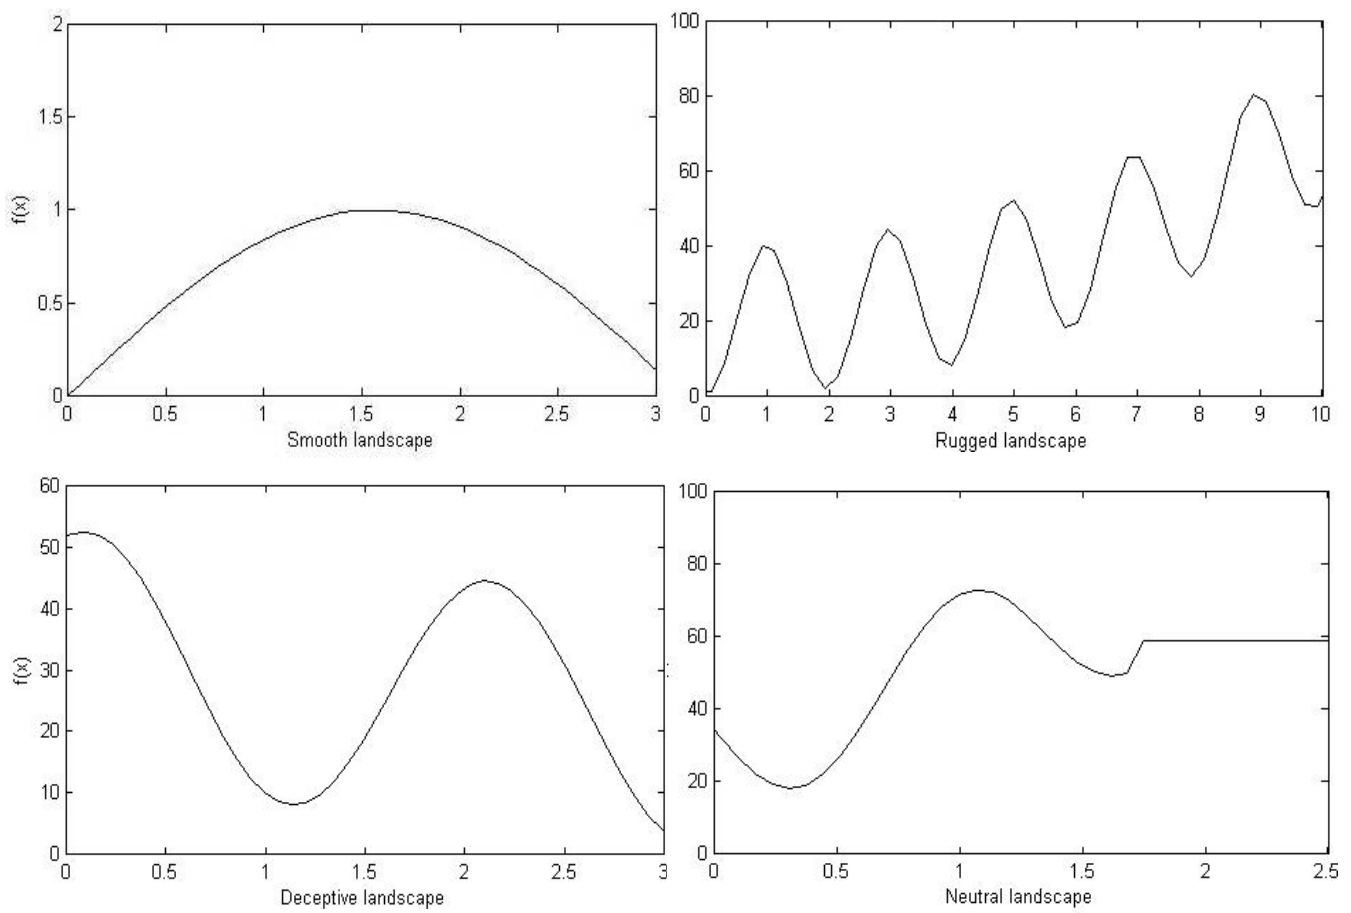


Figure 2. Landscape features adapted from Malan and Engelbrecht 2013.

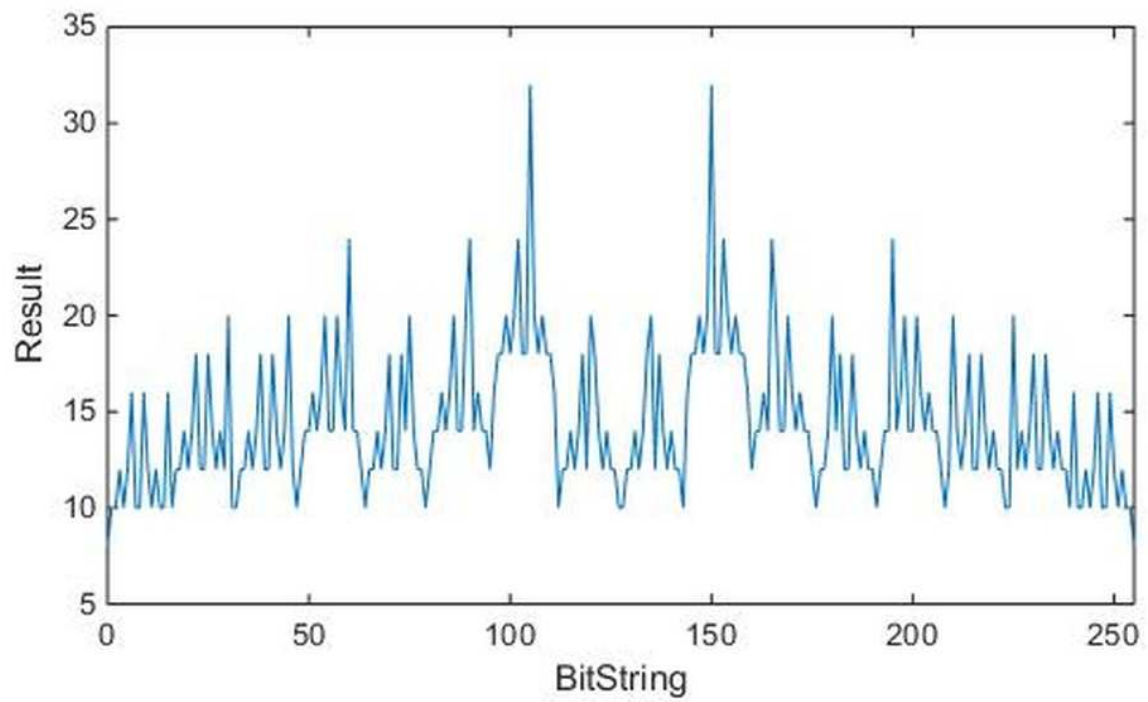



Figure 3. Same function (HXOR N=8) mapped with three different definitions of neighbourhood.

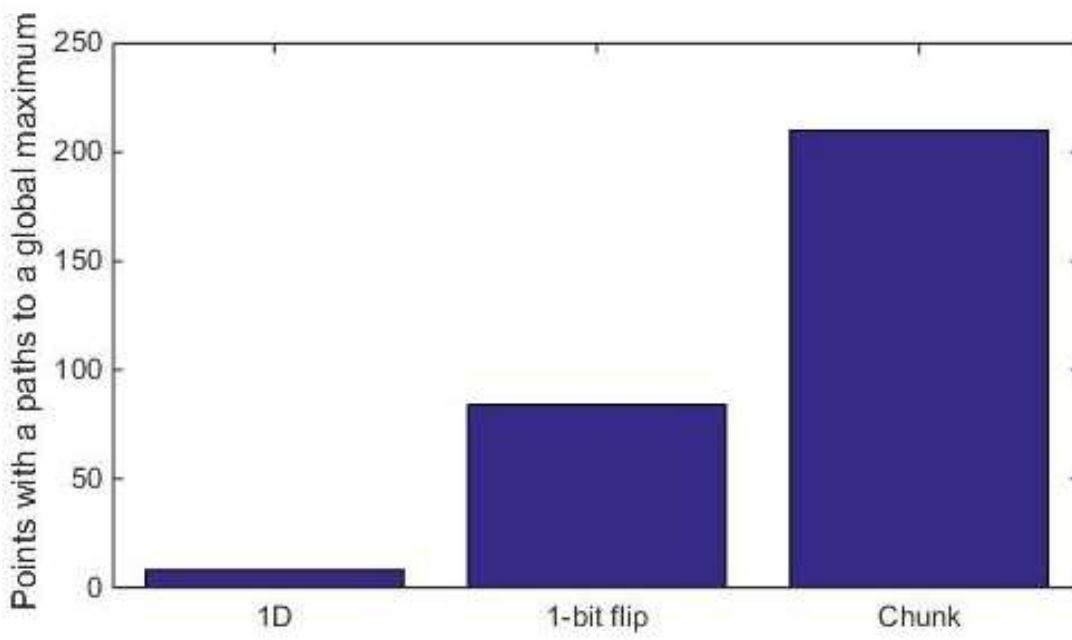



Figure 4. Comparative performance of 1D, 1-bit flip and chunking.

Maximum neutrality landscape (q=2)



Neutral landscape (q=4)






Figure 5. The same function ($N=8$, $K=3$, $q=4$) mapped with two different definitions of neighbourhood.




Figure 6. Comparative performance of a 1-bit hill climber on two landscapes: $N=8$, $K=3$ and its neutral counterpart $N=8$, $K=3$, $Q=4$.




Figure 7. Performance varying with restart probability for a stochastic hill-climber.




Figure 8. Performance of three different search algorithms on an HXOR landscape (N=8): 1-bit flip, chunking and random.




Figure 9. Performance of a 1-bit flip hill climber as a function of the number of remembered previous solutions.

Appendix 1

H-XOR function

The H-XOR function (Watson and Pollack 1999) is given by applying recursively an ‘exclusive or’ transformation onto the solution string where adjacent positions are considered starting with the leftmost. For instance, a {1010 0010} string becomes first {11 -1} and then {- -}. Once the transformation is completed, the payoff function rewards each non-null position in the hierarchy. Thus, a solution which contains an alternating pattern {1010 1010} would give a better score than a {1111 1111} since it will generate payoffs at lower levels of the hierarchy as well. The second level transformation for the first solution is {11 11} while for the second it is {- -}. The maximum score is given by {1001 0110} or symmetrically by {0110 1001} (see a more extensive description in paper 1).

Operations for the ‘chunking algorithm’

Chunks 8

Inverse all	e.g. 01111111->10000000
Mirror all	e.g. 01111111->11111110

Chunks 4 4

Inverse the 1 st chunk	e.g. 0111 1111->1000 1111
Inverse the 2 nd chunk	e.g. 0111 1111->0111 0000
Mirror the 1 st chunk	e.g. 0111 1111->1110 1111
Mirror the 2 nd chunk	e.g. 0111 1111->0111 0111
Permute the 1 st and 2 nd chunk	e.g. 0111 1111->1111 0111

Chunks 3 2 3

Inverse the 1 st chunk	e.g. 011 11 111->100 11 111
Inverse the 2 nd chunk	e.g. 011 11 111->011 00 111
Inverse the 3 rd chunk	e.g. 011 11 111->011 11 000
Mirror the 1 st chunk	e.g. 011 11 111->110 11 000
Mirror the 2 nd chunk	e.g. 011 10 111->011 01 000
Mirror the 3 rd chunk	e.g. 011 11 011->011 11 110

Permute the 1st and 3rd chunk e.g. 011 11 111->111 11 011

Chunks 2 2 2 2

Inverse the 1st chunk e.g. 01 11 11 11->10 11 11 11

Inverse the 2nd chunk e.g. 01 11 11 11->01 00 11 11

Inverse the 3rd chunk e.g. 01 11 11 11->01 11 00 11

Inverse the 4th chunk e.g. 01 11 11 11->01 11 11 00

Mirror the 1st chunk e.g. 01 01 01 01->10 01 01 01

Mirror the 2nd chunk e.g. 01 01 01 01->01 10 01 01

Mirror the 3rd chunk e.g. 01 01 01 01->01 01 10 01

Mirror the 4th chunk e.g. 01 01 01 01->01 01 01 10

Permute the 1st and 2nd chunk e.g. 01 11 11 11->11 01 11 11

Permute the 1st and 4th chunk e.g. 01 11 11 11->11 11 11 01

Permute the 2nd and 3rd chunk e.g. 11 01 11 11->11 11 01 11

Permute the 3rd and 4th chunk e.g. 11 11 01 11->11 11 11 01

Watson, R. A. and J. B. Pollack (1999). Hierarchically consistent test problems for genetic algorithms. Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on, IEEE.