
Fitness Landscapes and Evolvability

Tom Smith

toms@cogs.susx.ac.uk

Centre for Computational Neuroscience and Robotics, School of Biological Sciences,
University of Sussex, Brighton, UK

Phil Husbands

philh@cogs.susx.ac.uk

Centre for Computational Neuroscience and Robotics, School of Cognitive and
Computing Sciences, University of Sussex, Brighton, UK

Paul Layzell

paul_layzell@hp.com

Hewlett-Packard Laboratories, Bristol, UK

Michael O'Shea

m.o-shea@sussex.ac.uk

Centre for Computational Neuroscience and Robotics, School of Biological Sciences,
University of Sussex, Brighton, UK

Abstract

In this paper, we develop techniques based on evolvability statistics of the fitness landscape surrounding sampled solutions. Averaging the measures over a sample of equal fitness solutions allows us to build up *fitness evolvability portraits* of the fitness landscape, which we show can be used to compare both the ruggedness and neutrality in a set of tunably rugged and tunably neutral landscapes. We further show that the techniques can be used with solution samples collected through both random sampling of the landscapes and online sampling during optimization. Finally, we apply the techniques to two real evolutionary electronics search spaces and highlight differences between the two search spaces, comparing with the time taken to find good solutions through search.

Keywords

Evolvability, fitness landscape, search space, neutral evolution, NK system, evolutionary electronics.

1 Introduction

In this paper, we develop novel techniques based on local characteristics of the fitness landscape surrounding a solution. Averaging over a sample of equal fitness solutions allows us to build up *fitness evolvability portraits* of the fitness landscape, which we show can be used to compare both the ruggedness and neutrality in a set of tunably rugged and tunably neutral landscapes.

A feature of most fitness landscape descriptions is that a single global metric, e.g., correlation lengths, is used to describe the entire fitness landscape. The techniques presented in this paper develop a set of continuous metrics that vary with solution fitness. This approach allows fitness landscape features to be investigated at different fitness levels, leading to a fuller description of the space.

Many problems to which stochastic search techniques such as evolutionary computation are typically applied, present such highly skewed distributions of solution fitnesses that random sampling (even when some imposed distribution is applied to the

sample) is unlikely to represent fitnesses above a given level, even when such fitnesses are easily found through direct search optimization. In such spaces, we must develop descriptions that work with samples collected using *online* sampling techniques (in the remainder of the paper, we will use the term *online sample* to refer to samples collected during some search process, as opposed to samples collected through random sampling). We show that the fitness evolvability portraits presented work with samples of solutions collected both through random sampling techniques and through online sampling of the best solution so far found during simple hill-climbing optimization.

Finally, we investigate the application of the fitness evolvability portraits to a real evolutionary electronics problem, namely optimization of digital inverter circuits. We show that the portraits can be used to compare two different solution mappings, highlighting differences between the two search spaces and comparing the time taken to find good solutions through search.

The paper proceeds as follows: Section 2 outlines the concepts of fitness landscapes and neutrality and describes the relationship between problem difficulty and fitness landscape structure. Section 3 introduces the notion of solution evolvability as defined by local characteristics of the fitness landscape surrounding the solution and derives and applies the fitness evolvability portraits used in the remainder of the paper. Section 4 describes the tunably rugged and tunably neutral terraced NK landscapes used as test problems in this work. Sections 5 and 6 use the portraits derived in Section 3 to describe the test landscapes and show that they can be used to compare the ruggedness and neutrality in the tunably rugged and tunably neutral landscapes. Section 7 investigates the case where the fitness evolvability portraits are derived from solution samples collected during simple hill-climbing, showing that the portraits are robust to such online sampling. Finally, two real evolutionary electronics search spaces are investigated in Section 8, and the paper closes with discussion.

2 Fitness Landscapes and Neutrality

This section introduces two of the main concepts used in the paper. The *fitness landscape* (Section 2.1), first introduced by Wright (1932), describes the search space as a multidimensional landscape defined by the genotype-to-fitness mapping through which evolution moves. The classical idea of searching this landscape for good genotypes focuses on the difficulty of climbing up to the globally optimal fitness solution and avoiding locally optimal solutions. Here we argue that in difficult search problems, much of the time may be spent in nonadaptive *neutral* evolution (Section 2.2). Thus techniques aimed at describing the space in some way, must take account of the neutrality in the space. Section 2.3 describes how the difficulty of finding good solutions is determined by the structure of the fitness landscape, and Section 2.4 outlines different methods for sampling the fitness landscape structure.

2.1 Fitness Landscapes

Wright (1932) introduced the *fitness landscape* as a nonmathematical aid to visualize the action during evolution of selection and variation (in this paper, we will use the term *evolution* to refer to both natural biological evolution and the artificial evolution class of stochastic search processes that operate through some form of “generate-and-test” algorithm, e.g., genetic algorithms (Holland, 1992), genetic programming (Koza, 1992), evolutionary strategies (Rechenberg, 1973), and evolutionary programming (Fogel et al., 1966)). The description views the space in which evolution takes place as a landscape, with one dimension per genotype locus and an extra dimension, or height,




Figure 1: A two-dimensional model fitness landscape with one globally-optimal and one locally-optimal peak. From a starting point chosen at random, the search process tries to find good solutions. The process creates a new set of solutions through the application of genetic operators to the current solution(s), evaluating whether the new set is better than the current solutions. Evolving populations will tend to get stuck at the locally-optimal peak due to its large basin of attraction, and from there will only find the global optimum with difficulty.

representing the phenotype, or fitness, of that particular genotype.¹ The search space defined by a two-locus representation can thus be viewed as a three-dimensional fitness landscape (Figure 1) with each point corresponding to a single genotype and fitness. Applying a mutation operator to a particular genotype A typically produces a cluster of offspring genotypes lying close to A in the landscape, while recombination of two different genotypes A, B typically produces offspring genotypes lying somewhere between A and B in the landscape. Evolution can thus be viewed as the movement of the population, represented by a set of points (genotypes), towards higher (fitter) areas of the landscape.

This view of the search space leads naturally to the identification of the major problems with which evolution will have to cope: ruggedness and modality (Kauffman, 1993; Naudts and Kallel, 2000). Highly *epistatic* problems, where fitness is dependent on multiple inter-gene interactions, will produce a rugged landscape in which the direction to good solutions is obscured. Similarly, a high degree of *modality*, i.e., large numbers of local optima, will be seen as large numbers of hill-tops in the landscape with no neighbors of higher fitness. The majority of fitness landscape descriptions are based around these features of ruggedness and modality (see Weinberger (1990), Hordijk (1996), Jones and Forrest (1995), and Naudts and Kallel (2000)).

A more exact picture, especially when dealing with solutions represented by discrete-valued genotypes, is the connected graph (Stadler, 1996). Solution vertices, or nodes, are connected directly through the action of the genetic operators. The graph may show the space in a very different way than the fitness landscape: mutation operators acting on more than one locus, and other operators such as recombination, may not “see” fitness landscape hill-tops as local optima at all. However, local optima can

¹Wright defined two forms of fitness landscapes. The first version, used in this work, defines each point on the landscape as representing a single genotype with height corresponding to genotype fitness. The second version has each landscape point representing an entire population, with the values along each dimension representing the allele frequency over the population, and the height corresponding to the mean population fitness. The two approaches may show markedly different properties (Coyne et al., 1997).




Figure 2: Three two-dimensional model fitness landscapes showing the possible advantage of neutrality in a simple landscape with one globally optimal and one (nearly) locally optimal peak. (a) shows the two peaks as unconnected; populations evolving to the locally optimal peak will have difficulty moving to the global optimum. (b) shows the two peaks connected by a single neutral pathway; a population on the suboptimal peak may find the pathway. (c) shows the two peaks connected by a broad plateau; the population will move easily from the suboptimal peak to the global optimum.

clearly exist in the graph, occurring as graph nodes from which all connected nodes are of lower fitness. This definition may produce local optima with respect to genetic operators other than mutation; for example, some solutions may be local optima with respect to recombination operators.

The graph definition of the search space highlights the dangers in the simple visualizable picture afforded to us by the fitness landscape description: our intuitive view may not apply in higher-dimensional spaces. Fisher, for example, argued that local optima may not exist in a large class of high-dimensional spaces; the probability that a solution is optimal in every single dimension simultaneously is negligible (Provine, 1986, 274). However, it should be stressed that many problems clearly do show local optimality, e.g., the traveling salesman problem (Lawler et al., 1985). The next section introduces the idea of search space *neutrality*, one possible way in which some high-dimension spaces may differ radically from our intuitive viewpoint.

2.2 Fitness Landscape Neutrality

In the neutral theory, it is argued that evolving populations may spend relatively large periods of time undergoing nonadaptive neutral mutation (Kimura, 1983), staying at a constant height in the fitness landscape. The evolutionary timescale may be dominated by long periods of neutral epochs (van Nimwegen et al., 1999) interspersed with short periods of rapid fitness increase, i.e., *punctuated equilibrium* (Eldredge and Gould, 1972; Gould and Eldredge, 1977; Elena et al., 1996). During these neutral epochs, the population will move in the space through random drift (note that this is a separate process to Wright's idea of genetic drift due to finite population size (Provine, 1986)). Despite the undirected nature of the population movement, neutrality can be of use in escaping from (nearly) locally optimal solutions: Figure 2 shows three model landscapes illustrating the possible advantages of neutrality.

Neutral mutation in a fitness landscape will occur as movement between solutions of equal fitness, i.e., movement along *neutral networks*. Such neutral evolution cannot be distinguished from a population stuck in a local optimum simply through the fitness of the population over time. Instead, the underlying dynamics of the population must be investigated. Two key features have been predicted that distinguish the behavior of the population during these periods of neutral drift from the behavior of a population stuck in a local optimum. The first key feature is *movement*, i.e., is the population moving significantly in genotype space. The second key feature is *constant innovation*, i.e., the number of previously unencountered phenotypes seen over time is constant. Thus any phenotype should be accessible from a large enough neutral network; neutral drift can eventually find a higher fitness genotype and jump up to a higher neutral network.

These key features have been analyzed in many theoretical landscapes. Landscapes produced by RNA secondary structure folding algorithms show that neutral walks (a neutral variant on the random walk, where successive genotypes are generated by application of the mutation operator) have the property of constant innovation (Huynen et al., 1996; Huynen, 1996). Exhaustive analysis of some RNA landscapes has even mapped out the extent of the neutral networks and calculated the transition probabilities between the different networks (Forst et al., 1995; Grüner et al., 1996). Barnett (1998) introduces the *NK p* landscape, a tunably neutral variant on Kauffman's *NK* systems (Kauffman, 1993), and through calculation of population diffusion coefficients, shows highly neutral landscapes produce much more population movement during periods of no fitness change when compared to such periods in landscapes of zero neutrality. An important result from the *NK p* landscape work is that changing the amount of neutrality in the landscape has no effect on the ruggedness correlation function; correlation lengths do not predict the change in evolutionary dynamics seen for *NK p* landscapes of different neutrality (Barnett, 1998). Newman and Engelhardt (1998) investigate a similar tunably neutral variant of the *NK* system, finding that increased neutrality allows high fitness solutions to be found more easily through search. This is supported by Shackleton et al. (2000) who find that adding neutrality through redundancy can improve the level of fitness found through artificial evolution.

Neutrality has also been shown in real-world problem landscapes: In experiments on evolving tone recognition circuits, populations were seen to move in genotype space during periods where fitness did not increase (Harvey and Thompson, 1996). Experiments on evolution of digital circuits under two conditions – allowing neutral mutations and not allowing such changes – have also shown the importance of neutrality to the search process. Experiments where neutral changes were allowed consistently produced two-bit multiplier solutions of higher fitness than experiments without neutral mutation (Vassilev and Miller, 2000). Also, Thompson (2001) reports on the presence of neutral plateaus in evolving speech recognition circuits and shows through exhaustive mutation of genotypes at the start of the plateau that no transitions exist to a higher fitness level; the neutral evolution phase is necessary.

The degree of neutrality in a system is clearly a factor in the population dynamics during evolution. However, many fitness landscape descriptions simply do not identify this factor, e.g., the work on *NK p* landscapes showing that the autocorrelation function does not change with the level of neutrality (Barnett, 1998). In the next section, we describe the relationship between the structure of the fitness landscape and the difficulty of searching in the landscape, and we outline methods by which the fitness landscape is typically described.

2.3 Fitness Landscape Structure and Problem Difficulty

Clearly, the difficulty of searching in a given problem search space is related to the structure of the fitness landscape. However, the exact relationship between different landscape features and the time taken to find good solutions through optimization is not so straightforward. For instance, although landscapes containing large numbers of local optima are often hard to search, the difficulty of finding good solutions through optimization will depend on the fitness, distribution, and accessibility of such local optima. Similarly, landscape neutrality may be useful in both allowing optimization processes to escape from (nearly) local optima and potentially moving to regions of the landscape containing more good solutions but may be positively harmful if search processes spend long periods wandering randomly in flat areas of the space with no gradient information. In this paper, we recognize that the link between fitness landscape structure and the difficulty of finding good solutions through optimization is not yet fully understood and argue that extending techniques for fitness landscape description is an important step along the road to making such a link. Thus the techniques developed in this paper should be seen as complementary to other descriptions of fitness landscapes, some of which are outlined below.

Early work proposed that the correlation structure of the landscape was a key feature corresponding to problem epistasis or landscape ruggedness (Kauffman, 1993; Weinberger, 1990; Manderick et al., 1991; Lipsitch, 1991; Hordijk, 1996). Typically, the correlation structure of the landscape is derived in terms of parent-offspring fitness correlation or the fitness-distance autocorrelation function along a random nonadaptive walk. Similarly, fitness-distance correlation uses the correlation of genotype fitnesses with the fitness and distance of known optima as a metric (Jones and Forrest, 1995).

Related work defines the landscape in terms of the information required to describe the landscape. Fourier analysis decomposes the landscape into its Fourier components; the ruggedness of the system is measured by the number of Fourier components, especially those of high-frequency, required for the decomposition (Stadler and Wagner, 1998). Epistasis variance calculates the level of variance in the system with respect to a set of first order approximations to the system (Davidor, 1991). Similarly, information content analysis relates the ruggedness of the system to the fitness entropy encountered along a random walk (Vassilev, 1997; Vassilev et al., 2000)

However, the majority of such descriptions focus solely on the search space ruggedness measured through such correlation or information analysis as described above. No account is taken of the neutrality in the space; as described in Section 2.2, the NKp landscape correlation functions do not change when the degree of neutrality is varied. In the remainder of this paper, we introduce the technique of *fitness evolvability portraits* based on the characteristics of the fitness landscape surrounding sample solutions and show that such measures can be used to compare both the ruggedness and neutrality in a set of tunably rugged and tunably neutral landscapes.

2.4 Sampling the Fitness Landscape

In many hard problems, the distribution of solution fitnesses may be markedly non-normal, with fitnesses additionally distributed heterogeneously throughout the search space. In such problems, the set of solutions collected through random sampling methods will not accurately represent the set of solutions seen during optimization. All descriptions of the fitness landscape based on random sampling will thus not accurately reflect the fitness landscape as a whole. One extreme example is investigated by Smith et al. (2001a), where analysis of an evolutionary robotics search space shows

that fewer than 0.0001% of solutions in a random sample have fitness above 50% of the maximum in a neural network robot control problem despite this fitness being relatively easy to reach using optimization techniques. Two spaces that differ only in high fitness regions may show markedly different times to find good solutions through optimization, but fitness landscape descriptions based on random sampling will not show these differences (Smith et al., 2001a)

One potential approach is to bias the random sample procedure, keeping only some set percentage of solutions at each fitness. Even this method may fail to collect solutions above some fitness level in reasonable time, and it may be necessary to perform some kind of directed search process to collect the sample. Clearly, there is some point at which the time taken to collect such a sample may well approach a significant fraction of the time taken to solve the problem. For instance, if the sample required to characterize the problem involves collecting solutions at or near the optimum, we will have effectively solved the problem merely in the act of description. A useful analogy could be drawn with Marr's type II systems; the system may not be reducible to a simpler level of description than the system itself (Marr, 1976). By contrast, type I systems can be reduced to a simpler description, e.g., a fitness landscape that can usefully be reduced to a single correlation length description.

In Section 7, we collect samples through simple hill-climbing, and show that the fitness evolvability portraits based on the biased sample set make the same predictions as those based on unbiased random samples. Although the *NK* landscapes used in this paper have approximately normal fitness distributions, verifying that the portraits are reasonably robust to sample bias is important if we are to use them on other problems with highly skewed fitness distributions (Smith et al., 2001b).

In the next section, we introduce the notion of evolvability as the capacity of a solution to evolve, closely tied to the fitness landscape neighboring that solution. We then derive a set of solution and population evolvability metrics using them to build fitness evolvability portraits of sample fitness landscapes.

3 Evolvability and the Transmission Function

Evolvability is loosely defined as the capacity to evolve, alternatively the ability of an individual or population to generate fit variants (Altenberg, 1994; Marrow, 1999; Wagner and Altenberg, 1996). Thus evolvability is more closely allied with the *potential* for fitness than with fitness itself; two equal fitness individuals or populations can have very different evolvabilities (Turney, 1999). Typically, researchers use some definition of evolvability based on the offspring of current individuals or populations: in this paper we follow Cavalli-Sforza and Feldman (1976) and Altenberg (1994) in using the *transmission function* of all possible offspring from a parent to define a set of metrics of evolvability (see Section 3.1 for further details).

It is often argued that there may be long-term trends for evolvability to increase during evolution (see Wilke (2001) and Turney (1999)). However, as evolvability is more directly related to fitness potential than fitness itself, long-term change cannot be due to straight fitness selection. Thus any trend towards change in evolvability can only be understood through some second order selection mechanism by which evolution tends to retain solutions that have a more evolvable genetic system (Dawkins, 1989; Kirschner and Gerhart, 1998).

Researchers in both biology and evolutionary computation typically link evolvability with the local structure of the search space. For example, Burch and Chao (2000) show that RNA virus evolvability can be understood in terms of the mutational neigh-

borhood, while many evolutionary computation researchers (see Ebner et al. (2001) and Marrow (1999)) argue that changing the properties of the search space (through such mechanisms as adding neutrality) can affect evolvability as evidenced by the speed of evolution. The interest in evolvability for evolutionary computation practitioners is thus tied closely to work on the ruggedness and modality of the search space, argued to primarily influence the ease of finding good solutions in the space (Weinberger, 1990; Hordijk, 1996; Jones and Forrest, 1995; Naudts and Kallel, 2000).

Recent work has emphasized that in addition to landscape ruggedness and modality, search space neutrality may have impact on the population dynamics of evolution (Section 2.2). This factor may not be identified by many standard measures aimed at the landscape ruggedness and local modality but may be measurable through change in evolvability. For example, recent artificial evolution research has shown that evolvability can change during neutral epochs; populations tend to move to “flatter” areas of the fitness landscape where fewer mutations are deleterious (Wilke et al., 2001; Wilke, 2001). This can clearly have an impact on the speed of search but may not be picked up by the standard landscape ruggedness and modality descriptions.

Other biological research in evolvability is also relevant to evolutionary computation, e.g., the work on adaptation to change in environment through such mechanisms as alleles providing increased mutation rates (Taddei et al., 1997; Sniegowski et al., 1997). However, in this paper we focus on evolvability in terms of the properties of the solutions’ local search space. The next section outlines the offspring transmission function and defines a simple set of evolvability metrics.

3.1 The Transmission Function

In this paper, we follow the definition of evolvability as the ability of individuals and populations to produce fit variants, specifically the ability to both produce fitter variants and to not produce less fit variants. This definition is intimately tied in with research on the *transmission function* T (Altenberg, 1994; Cavalli-Sforza and Feldman, 1976) and the population offspring probability distribution function ϕ from all possible applications of the genetic operators to the parent(s)

$$\phi(g, f) = \int \int \int \int \psi(h, k, h', k') T(g, f : h, k, h', k') dh dk dh' dk' \quad (1)$$

or the probability ϕ (with parental selection function ψ) of obtaining offspring genotype g and phenotype f over all parents of genotypes h, h' and phenotypes k, k' . The transmission function T is the probability density function of obtaining g, f given h, k, h', k' (Cavalli-Sforza and Feldman, 1976).

In the absence of recombination, only a single parent h, k is required to produce offspring through mutation (in Section 9 we discuss the impact of recombination on the techniques developed in this paper):

$$\phi(g, f) = \int_{-\infty}^{\infty} \psi(h, k) T(g, f : h, k) dh dk \quad (2)$$

or the probability of obtaining offspring g, f over all parents h, k with selection ψ . In this paper, we focus on the offspring of a set of single genotypes (saved during the course of evolutionary runs), so do not integrate over the set of all possible parents. Similarly, the selection function can be omitted as we preselect the parent. Since we are interested only in the offspring phenotypes f and not the offspring genotypes g , we can

refer to the transmission function $T(f : h, k)$ as shorthand for the probability density function of offspring fitnesses from a single parent h, k .

The transmission function thus encompasses both the operators and the representation; instead of referring to good and bad genetic operators or good and bad representations, we can talk about the effectiveness of the transmission function. Thus the evolvability of an individual or population, i.e., their ability to generate fit variants, is simply a property of the individual or population transmission function. The next section derives measures for the evolvability of an individual solution in terms of this transmission function for continuous variables.

3.2 Evolvability Metrics: Continuous Variables

The evolvability of a solution genotype h and fitness k is directly tied to the probability of that solution not producing offspring of lower fitness. Thus we derive our first metric of evolvability E_a :

$$E_a = \frac{\int_k^\infty T(f : h, k) df}{\int_{-\infty}^\infty T(f : h, k) df} \tag{3}$$

or the probability that the offspring fitness f is greater or equal to the current fitness k , i.e., the mutation is nondeleterious. Since the transmission function $T(f : h, k)$ is a probability density function, the infinite integral sums to unity, so we have

$$E_a = \int_k^\infty T(f : h, k) df \tag{4}$$

Low fitness solutions may have a larger E_a than high fitness solutions simply due to the increased number of better mutations. The second evolvability metric E_b uses only the offspring fitnesses:

$$E_b = \int_{-\infty}^\infty f T(f : h, k) df \tag{5}$$

or the expected offspring fitness from genotype h . Note, this value is fitness dependent so should not be compared across genotypes without reference to their original fitness. A further problem with both E_a and E_b is their dependence on the entire set of offspring fitnesses; the fraction of offspring that are significantly fitter than the parent may be extremely small. The third measure reflects this dimension of evolvability, looking only at the top C th percentile of the offspring fitnesses

$$E_c = \frac{100}{C} \int_{F_c}^\infty f T(f : h, k) df \tag{6}$$

$$\text{where } F_c \text{ defined by } \int_{F_c}^\infty T(f : h, k) df = \frac{C}{100} \tag{7}$$

or the expected fitness of only the top C th percentile of fitnesses. A similar measure E_d (not shown) calculates the expected fitness of the bottom C th percentile of offspring.

The next section extends the continuous analysis presented above to the discrete set.

3.3 Evolvability Metrics: The Discrete Set

Consider the fitness landscape as a directed graph (V, E) with vertices V (genotypes) connected by edges E (defined by the genetic operators). The set G of offspring from a




Figure 3: Three continuous one-dimensional landscapes $F(x)$ with the parent genotype x_0 shown by the solid circle lying at $x = 0$ (in all cases, $F(0) = 0$). The mutation operator $\mu(x, x_0)$ is a probability distribution function, producing offspring x_1 lying in a uniform distribution around x with range d , shown by the thick bar below each landscape, centered on x_0 . See text for the derived evolvability in each landscape.

parent genotype h, k is thus defined by the vertices connected to the parent vertex:

$$G(h, k) = \{g \in V : E(h, k) = g\} \tag{8}$$

The fitness function F maps each vertex on to a single fitness, so similarly, we define the set of offspring with fitness $F(g)$ equal to or greater than some fitness c :

$$G_c^+(h, k) = \{g \in V : E(h, k) = g, F(g) \geq c\} \tag{9}$$

The probability of the offspring fitness being higher or equal to the parent fitness, or E_a , is simply the fraction of the set with $F(g) \geq k$:

$$E_a = \frac{|G_k^+(h, k)|}{|G(h, k)|} \tag{10}$$

As in the previous section, the mean fitness of the offspring solutions, or E_b , is simply the mean fitness of all members of the set:

$$E_b = \frac{\sum_{g \in G(h, k)} F(g)}{|G(h, k)|} \tag{11}$$

The mean fitness of the set of offspring with fitness in the top C th percentile is similarly defined:

$$E_c = \frac{\sum_{g \in G_{F_c}^+(h, k)} F(g)}{|G_{F_c}^+(h, k)|} \tag{12}$$

$$\text{where } F_c \text{ defined by } |G_{F_c}^+(h, k)| = \frac{C |G(h, k)|}{100} \tag{13}$$

The mean fitness of the set of offspring with fitness in the bottom percentile can be defined through the set $G_{F_d}^-(h, k)$ of offspring with fitness below some fitness F_d .

The next section applies the metrics to a set of simple cases, where the parent genotypes lie at different points in a hypothetical landscape.

Table 1: Derived values for the evolvability of the three sample landscapes shown in Figure 3.

Metric	Plateau $F(x) = 0$	Hill-side $F(x) = kx$	Optima $F(x) = kx^2$
E_a	1	0.5	1
E_b	0	0	$kd^2/3$
$E_{c=25}$	0	$3kd/4$	$37kd^2/48$
$E_{d=25}$	0	$-3kd/4$	$kd^2/48$

3.4 Simple Evolvability Examples

The metrics derived in the previous two sections are here applied to a set of simple cases, showing their ability to discriminate between solutions lying at different points in a hypothetical landscape, including flat plateaus, local optima, and hill-sides. Figure 3 shows three such sample cases.

Through applying the continuous evolvability metrics (Equations 4–7) to the continuous landscapes defined by the functions $F(x)$, with the mutation operator probability distribution around the parent solution $\mu(x, x_0)$, we obtain the following results for the evolvability of parent solutions x_0 (and offspring solutions x_1):

$$E_a \equiv P(F(x_1) \geq F(x_0)) \tag{14}$$

$$= \int_{-\infty}^{\infty} \mu(x, x_0) H(F(x) - F(x_0)) dx \tag{15}$$

where the Heaviside function $H(a) = \begin{cases} 1 & a \geq 0 \\ 0 & \text{else} \end{cases}$ \tag{16}

$$E_b \equiv \langle F(x_1) \rangle = \int_{-\infty}^{\infty} F(x) \mu(x, x_0) dx \tag{17}$$

$$E_c = \frac{100}{C} \int_{-\infty}^{\infty} F(x) \mu(x, x_0) H(F(x) - F_c) dx \tag{18}$$

$$\text{where } \int_{-\infty}^{\infty} \mu(x, x_0) H(F(x) - F_c) dx = \frac{C}{100} \tag{19}$$

Similarly, the mean fitness over the bottom percentile E_d can be defined using the Heaviside function $H(F_d - F(x))$.

Table 1 gives the results for the four evolvability metrics derived on the three landscapes shown in Figure 3 for the uniformly distributed mutation operator $\mu(x) = U(x_0 - d, x_0 + d)$ and a parent genotype $x_0 = 0$. The evolvability data for the first landscape tells us that (1) no mutations are deleterious, (2) the expected offspring fitness is equal to the current fitness, and (3) the expected fitnesses over both the top and bottom quartiles of offspring are equal to the current fitness: we conclude the landscape neighboring the current solution must be a flat plateau. We can similarly identify the hillside and local optima landscapes from the relevant data. No single metric gives enough information to correctly identify the nature of the fitness landscape surrounding the genotype x_0 . For example, the expected offspring fitness (E_b) equals zero for both the plateau and hillside, but the fraction of nondeleterious mutations (E_a) or the tails of the offspring fitness distribution (E_c and E_d) allow the points to be differentiated. Thus the combination of the four metrics allows us to identify each type accurately. The ap-




Figure 4: Three two-dimensional model fitness landscapes showing the possible advantage of neutrality in a simple landscape with one globally optimal and one (nearly) locally-optimal deceptive peak. The two peaks have fitnesses of 1.0 and 0.5, respectively, and the neutral pathway and plateau have fitnesses of 0.5.

proach can also be used on problems with higher dimensional landscapes, although the offspring distributions may need to be approximated through sampled applications of the mutation operator(s).


In the next section, we show how these evolvability metrics can be averaged over populations of solutions to produce the fitness evolvability portraits used in the remainder of this paper, and relate these portraits to levels of ruggedness, modality, and neutrality in the landscape.

3.5 Population Fitness Evolvability Portraits


The previous section described how the evolvability metrics could be calculated over the fitness neighborhood for a single solution genotype. We can define the same evolvability metrics over a sampled population of solutions through simply defining the metrics as calculated over the sum of population transmission functions, i.e., we take the distribution of offspring fitnesses from all members of the sample and calculate the evolvability metrics. For the discrete case, this translates to taking the population set of offspring defined over the combined sets of offspring from all members of the population.

Two important ideas emerge from this definition of population evolvability. First, we can compare entire populations simply by comparing their metrics of evolvability. This is not explored further in this paper but has been used by Smith et al. (2001c, 2001d) to investigate the behavior of populations during neutral epochs, in particular whether the populations are moving to more evolvable areas of space during such neutral epochs. Second, we can take samples of equal fitness (in practice, we take samples of nearly equal fitness lying in some range) to build up a fitness evolvability portrait of the landscape. For each equal fitness sample of solutions, we can calculate the population evolvability and plot the evolvability metrics against solution fitness².


²It should be noted that the idea of plotting some measure over fitness was used by Rosé et al. (1996) in their density of states approach. However, in this paper, we focus on the evolvability of solutions at some fitness rather than simply the number of such solutions.




(a) Probability of a nondeleterious mutation E_a



(b) Expected fitness over all mutations E_b



(c) Expected fitness over top quartile of mutations E_c



(d) Expected fitness over bottom quartile of mutations E_d

Figure 5: Fitness evolvability portraits for the three model landscapes shown in Figure 4. The evolvability metrics were calculated from an exhaustive sample set of solutions.

3.5.1 Three Model Landscapes

In this section, we show how the fitness evolvability portraits can be derived for the three model landscapes shown in Figure 4 and illustrate the advantages of the portraits over other available landscape descriptions.

Figure 4 shows the same three model landscapes used in Section 2.2 to illustrate the potential advantages of landscape neutrality. It should be emphasized that the landscapes are used purely to illustrate the potential advantage for searching in landscapes with varying levels of neutrality and are not drawn from real problem spaces. The three landscapes shown here are discrete-valued 100-by-100 grids, and for both the adaptive walks and the evolvability analysis on these landscapes, the same mutation operator was used whereby offspring solutions were created from any one of the eight grid nearest neighbors.




Figure 6: The probability of finding a nondeleterious mutation E_a over the three landscapes shown in Figure 4, shown in detail for fitness levels near 0.5. Over most of the fitness range, the graphs are identical. However, the increase at a fitness of 0.5 (the height of the locally optimal peak) is clear for the neutral pathway (Figure 4(b)) and plateau (Figure 4(c)) landscapes when compared with the unconnected peaks landscape (Figure 4(a)).

3.5.2 Landscape Fitness Evolvability Portraits

Figure 5 shows the four fitness evolvability portraits over the three model landscapes from Figure 4; each is derived in the same fashion. A sample of solutions (in this case we take an exhaustive sample, but in general a random or online sample would be collected) is saved, and a set of offspring fitnesses from each sample solution calculated, and the metrics of evolvability calculated for that solution. To build up the fitness portraits, we average each metric of evolvability across the sample of solutions at equal fitness (in practice, we take solutions lying in some small range of fitness) and plot the average evolvability metric against fitness.

All four graphs in Figure 5 show that the three model landscapes have similar properties over the fitness range, with the exception being seen near fitnesses of 0.5. Figure 6 shows this area in detail for the probability of a neighboring solution being of equal or higher fitness E_a . When the two peaks are unconnected, the probability of finding a fitter neighbor at this fitness drops sharply to nearly 40%, the lowest value seen until we reach fitnesses of nearly 100%. However, when the peaks are connected by a single neutral pathway, this value rises sharply to a value of over 60%, and when the peaks are connected by a broad plateau, the corresponding value is over 90%. This in itself does not allow us to predict the difficulty of finding good solutions (near the top of the fitness range) as we have no way of knowing the importance of the 0.5 fitness level. However, when combined with information on direct search processes getting stuck at fitnesses of 0.5 for the unconnected peaks, we can hypothesize that the increased level of neutrality seen in the fitness evolvability portraits is enabling the search processes to escape from some local optima.

Now we can compare the information gained from the fitness evolvability portraits with other descriptions of the space. Figure 7 shows the fitness distributions over the whole space for the three landscapes (shown with \log_{10} scaling) and the correlation lengths, derived from random walks. As expected, the distribution does show more solutions of fitnesses 0.5 for the neutral pathway and plateau landscapes, but there is no

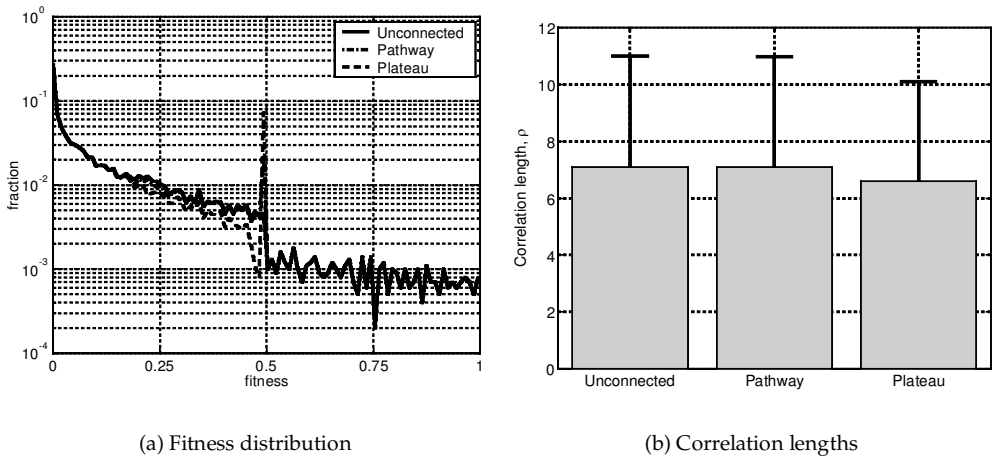


Figure 7: (a) Fitness distributions and (b) correlation lengths over the three model landscapes shown in Figure 4. Note that the fitness distribution is shown with \log_{10} scaling to show the detail at high fitnesses. Correlation lengths were calculated over 100 random walks of 100 steps each.

way to tell how these fitnesses are distributed over the space, i.e., that the spaces contain neutral pathways. The correlation lengths are equal, so give us no way to differentiate the spaces. Thus we see the strength of the fitness evolvability portraits, allowing us to both recognize landscape differences at different fitnesses and giving us a description of the accessibility and distribution of solutions across the space.

3.5.3 Landscape Ruggedness, Modality, and Neutrality

We can also relate the fitness evolvability portraits to general features of the landscape in a direct way.

First, landscape ruggedness is related to the fitness of offspring; the epistasis of a landscape is often measured through the correlation between the fitnesses of solutions a given distance away in the space. This is closely related to the expected fitness of the offspring solutions (E_b shown in figure 5(b)) and, in particular, the gradient of the expected offspring fitness versus parent fitness graph. Barnett (2001) has shown analytically that this gradient is proportional to the autocorrelation function for the mutation operator used. From Figure 5(b), the expected offspring fitness portrait predicts the three landscapes are of equal ruggedness supported by the correlation lengths shown in Figure 7(b). This link is illustrated in the next sections when analyzing NK landscapes of varying ruggedness.

Second, landscape modality is related to the probability of offspring being of lower fitness (E_a shown in Figure 5(a)). There is a direct link between the average number of nondeleterious mutations for a sample of equal fitness solutions and the probability that a given solution at that fitness is a local optima. Thus the fitness evolvability portrait of nondeleterious mutations can, in principle, give us an estimate of the numbers and fitness of local optima in the space. For our three landscapes, the overall probability of nondeleterious mutations does not decrease dramatically as fitness increases (ignoring the behavior near fitnesses of 0.5), thus we might predict the spaces contain




Figure 8: Correlation lengths calculated over random walks on Kauffman’s NK landscapes shown for $N = 25$, $K \in \{0, 1, 2, 6, 12, 18, 24\}$ (mean and standard deviation over 1000 walks shown). The correlation lengths decrease with increasing K , showing the increasing ruggedness in the system.

few, if any, local optima. Again, we see this relationship when analyzing the NK landscapes.

Finally, landscape neutrality can be linked to the probability of obtaining nondeleterious offspring in the same way as landscape modality. Typically, neutrality is defined through the number of neutral neighbors, so increasing levels of neutrality are seen as higher probabilities of nondeleterious mutations, thus lower probabilities of solutions at a given fitness being a local optima. However, we can also identify at what fitness levels this increasing neutrality is expressed, e.g., do the neutral networks permeate the space at all fitness levels? For our model landscapes, it is not simple to disassociate the effects of modality and neutrality, however the lack of increase in nondeleterious mutations at fitnesses just above 0.5 for the neutral pathway and plateau landscapes (Figure 6) leads us to predict that the increase in such mutations at fitness 0.5 is due to the existence of neutral networks at this fitness.

In the next section, we outline the NK and terraced NK landscapes.

4 The NK and Terraced NK Landscapes

4.1 Kauffman’s Tunably Rugged NK Landscapes

Consider a genotype consisting of $N > 0$ loci, with each locus having some state drawn from an alphabet of size A . In addition, each locus epistatically interacts with $0 \leq K < N$ linked loci. The genotype fitness is defined as the mean fitness over the N loci, with each locus fitness uniquely determined by both its own state and the state of the K loci it is epistatically linked to. Thus each locus has a total number of A^{K+1} assigned fitnesses (one fitness for each of the possible states for the locus and K linked loci) typically drawn from a uniform distribution over $[0, 1]$, and the entire landscape is defined by the table of NA^{K+1} fitnesses. In the remainder of this paper, we take a binary alphabet $A = 2$.

This NK landscape is a generalized version of the spin-glass model found in statistical physics (see Binder and Young (1986)). As K increases from the unimodal $K = 0$ landscape, the *ruggedness* of the landscape increases up to the maximally rugged random $K = N - 1$ landscape, corresponding to the random energy spin-glass model




Figure 9: Correlation lengths calculated over random walks on the neutrally terraced NK landscapes shown for $N = 25$, $K \in \{0, 1, 2, 6, 12, 18, 24\}$, $F = 4$ (mean and standard deviation over 1000 walks shown). The correlation lengths are identical to the corresponding lengths shown in Figure 8 (calculated on the non-neutral NK landscape) and decrease with increasing K . Thus the correlation length measures still predict increasing ruggedness, even in the presence of significant levels of neutrality in the system.

(Derrida, 1981). Early work by Weinberger (1990, 1991) showed that increasing ruggedness of the system is well predicted by the decrease in correlation length for the system, and that the number of locally optimal peaks increases dramatically with K ; this has long been a benchmark result arguing that landscape ruggedness is the key feature for problem difficulty. Figure 8 shows the correlation lengths derived from random walks in the $N = 25$, $K \in \{0, 1, 2, 6, 12, 18, 24\}$ landscapes; increasing K clearly produces more rugged landscapes with shorter correlation lengths.

The next section describes tunably neutral extensions to Kauffman's NK systems developed by Barnett (1998) and Newman and Engelhardt (1998).

4.2 Tunably Neutral NK Landscapes

It has been argued by many researchers that other features of problem difficulty are not adequately represented by measures of landscape ruggedness. One such feature is landscape neutrality (Section 2.2), and two different methods have been proposed to incorporate a tunable level of neutrality into the NK landscapes.

Barnett (1998) has argued that in real systems, many loci will not contribute to fitness, thus some fraction $0 \leq p \leq 1$ of the NA^{K+1} locus fitnesses are taken to equal zero. The neutrality of the system is thus represented by p , which can be tuned from zero (the standard NK landscape) to 1.0 (all genotypes have fitness zero). The resulting NKp landscapes are thus both tunably rugged and tunably neutral. Similarly, Reidys and Stadler (2001) derive the more general additive fitness landscape class of which the NKp landscapes are a special case. Barnett goes on to show that changing the level of neutrality on the system has large impact on the population dynamics of evolutionary search but produces no change in the correlation length (Barnett, 1997; Barnett, 1998).

A second approach taken is to discretize the possible locus fitnesses through only allowing each locus to take one of F fitnesses, or *terraces* (Newman and Engelhardt, 1998). Varying this terrace parameter F from an infinite value (corresponding to the standard NK landscape) down to the minimal value of 2 (where each locus has a fit-

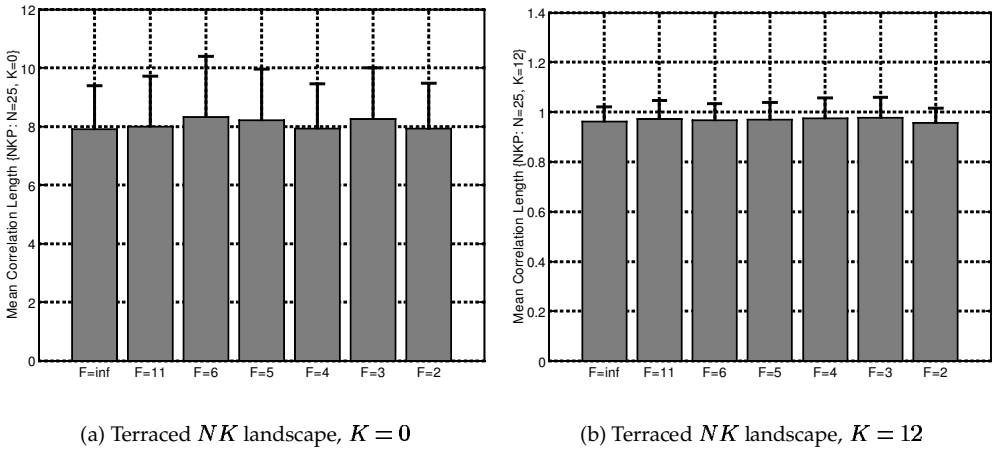


Figure 10: Correlation lengths calculated over random walks on the neutrally terraced NK landscapes shown for $N = 25$, $K \in \{0, 12\}$, $F \in \{\infty, 11, 6, 5, 4, 3, 2\}$ (mean and standard deviation over 1000 walks shown). The correlation lengths are equal for a given K despite the varying amount of neutrality in the system (shown by the varying number of neutral terraces F). Thus correlation lengths do not predict the difference in time required to find high fitness solutions.

ness of 0 or 1) thus decreases the number of possible fitnesses, so increasing the fraction of equal fitness neutral solutions (the actual number of distinct solution fitnesses will be equal to $1 + (N - 1)(F - 1)$). Figure 9 shows that the correlation lengths still predict increasing ruggedness in the presence of significant neutrality. Newman and Engelhardt (1998) go on to show that increasing the neutrality in the system reduces the evolutionary time required to find higher fitness solutions. In Figure 10, we show that this evolutionary speed difference is not predicted by the change in correlation length; the correlation length for given K is constant for all levels of neutrality.

In the following two sections, we investigate the evolvability of the terraced NK landscapes with varying degrees of ruggedness and neutrality.

5 Evolvability of a Tunably Rugged Landscape

For a given genotype length N , the overall distribution of *genotype* fitnesses will be approximately equal to the normal distribution $N(M_N, \sigma_N^2)$ for all K , i.e., the distribution obtained from the mean of N *loci* fitnesses sampled from a uniform $[0, 1]$ distribution (M_N and σ_N^2 are the mean and variance, respectively). As N increases, the mean of this genotype fitness distribution M_N will stay constant at $M_N = 0.5$, but the standard deviation will decrease with the number of samples: $\sigma_N = \sigma/\sqrt{N}$. Thus as our genotypes increase in length, the fraction of extremely fit and unfit genotypes will decrease as the variance of the population decreases (note that some NK models address this issue through scaling the genotype fitnesses by \sqrt{N} (see Sibani and Pederson (1999))).

Although the overall distribution of genotype fitnesses is equal for all K over a given genotype length N , the distribution of these fitnesses over the search space is decidedly not equal for all K . For the unimodal $K = 0$ landscape, all high fitness genotypes are grouped in a small volume of the search space, while for the maximally

multimodal $K = N - 1$ landscape, the fitness distribution over the search space is random. In general, the distribution of fitnesses neighboring a solution of given fitness is approximately normal with mean and deviation dependent on N, K and the current solution fitness (see Weinberger (1990) and Stadler and Schnabl (1992)). From this it is possible to derive the expected fitnesses (and the time taken on both adaptive and random walks) at which local optima are reached for various N and K (again, see Weinberger (1990) and Stadler and Schnabl (1992)). In the next section, we derive analytic and empirical results for the evolvability measures when applied to the NK landscapes.

5.1 Analytically Derived Evolvability for NK Landscapes

In this section, we focus on the probability that an offspring derived from a single bit mutation of the parent has a higher (or equal) fitness than the parent, i.e., the first evolvability metric E_a (Section 3.1), as a function of the parental fitness. The other evolvability metrics derived in Section 3 can be similarly treated. Consider a parent genotype of fitness F_0 , the mean of the N locus fitnesses f_i drawn from a uniform distribution over $[0, 1]$:

$$F_0 = \frac{1}{N} \sum_{i=1}^N f_i \text{ where } f_i \in U[0, 1] \tag{20}$$

Now, the probability that the offspring fitness F_1 is not lower than the parent fitness is simply the probability that the $K + 1$ loci affected by a single bit mutation do not, on average, decrease in fitness:

$$E_a \equiv P(F_1 \geq F_0) = P\left(\left(\frac{1}{K+1} \sum_{i=1}^{K+1} f_i\right) \geq F_0\right) \text{ where } f_i \in U[0, 1] \tag{21}$$

is the probability that the mean of $K + 1$ uniformly distributed samples is not smaller than the current fitness. For the unimodal $K = 0$ we can solve trivially

$$P(F_1 \geq F_0) = P(f_1 \geq F_0) = 1 - F_0 \text{ where } f_1 \in U[0, 1] \tag{22}$$

For $K \gg 0$, the mean of affected loci fitnesses tends to a normal distribution with mean $M_{K+1} = 0.5$ and deviation $\sigma_{K+1} = \sigma/\sqrt{K+1}$ (where σ is the deviation of loci fitnesses as $K \rightarrow \infty$, assumed to be non-zero and finite). For a normal distribution $N(M, \sigma^2)$, the probability density function n is given by


$$n = \frac{1}{\sigma} \phi\left(\frac{x - M}{\sigma}\right) \tag{23}$$

$$\text{where } \phi(z) = \frac{1}{\sqrt{2\pi}} \exp(-0.5z^2) \tag{24}$$


The probability that a random number drawn from this distribution is greater than some value F_0 is given by the integral of the probability density function over the relevant limits (with mean $M = 0.5$ and deviation $\sigma_{K+1} = \sigma/\sqrt{K+1}$):

$$P(F_1 \geq F_0) = \frac{\sqrt{K+1}}{\sigma} \int_{F_0}^{\infty} \phi\left(\frac{\sqrt{K+1}(x - 0.5)}{\sigma}\right) dx \tag{25}$$

$$= \sqrt{\frac{K+1}{2\pi\sigma^2}} \int_{F_0}^{\infty} \exp\left(\frac{-(K+1)}{2\sigma^2}(x - 0.5)^2\right) dx \tag{26}$$



(a) Analytically derived probability of a non-deleterious mutation E_a



(b) Empirically derived probability of a non-deleterious mutation E_a

Figure 11: Analytically and empirically derived probabilities of a nondeleterious mutation E_a . The derived probability is plotted as $50\text{erfc}(C(F_0 - 0.5))$. The empirical probability is calculated on the NK landscape with $N = 25$, $K \in \{0, 6, 12, 24\}$ for a random sample set of solutions.

which is simply the complementary error function $\text{erfc}(x)$:

$$\text{erfc}(x) \equiv \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-z^2) dz \tag{27}$$


$$\text{and } P(F_1 \geq F_0) = 0.5 \text{erfc} \left(\sqrt{\frac{K+1}{2\sigma^2}} (F_0 - 0.5) \right) \tag{28}$$

Note that an equivalent result to Equation 28 is derived by Stadler and Schnabl (1992) in order to calculate the probability of solutions of given fitness being local optima; Section 3.5.3 describes the link between the fitness evolvability portraits and general landscape features.


5.2 Empirically Derived Evolvability for NK Landscapes

Figure 11 shows data generated from Equation 28 compared to the fitness evolvability portraits derived from empirical random sampling of simulated NK landscapes ($N = 25$, $K \in \{0, 1, 2, 6, 12, 24\}$), showing good agreement between the analytically and empirically derived data. All NK random sample sets used in this paper consist of 1000 individuals sampled from each of 100 generated landscapes – a total of 100,000 sampled solutions. Note that for each set of N, K , there is an arbitrarily large number of fitness lookup tables that can be generated, and thus an arbitrarily large number of possible landscapes. For this reason, we sample both a set of individuals and a set of landscapes for each value of N, K .


Both sets of data predict that as K increases, the probability of finding a fitter mutant *increases* for parent fitnesses below the population mean of 0.5. Only for parent fitnesses above this mean value of 0.5 does the probability of reaching a fitter mutant favor the lower K landscapes. This can be understood by considering that a single bit flip mutation can affect the fitness F_0 by a fraction of order $O(\frac{K+1}{N})$. Low K landscapes




(a) Probability of a nondeleterious mutation E_a



(b) Expected fitness over all mutations E_b



(c) Expected fitness over top quartile of mutations E_c




(d) Expected fitness over bottom quartile of mutations E_d

Figure 12: Fitness evolvability portraits for the NK landscapes with $N = 25$, $K \in \{0, 6, 12, 24\}$. The evolvability metrics were calculated from a random sample set of solutions.


are thus highly correlated³, and offspring fitnesses are close to parent fitnesses. For high K landscapes, the offspring-parent fitnesses are less correlated, thus offspring fitnesses *on average* are close to the population mean of 0.5, and the distribution of genotype fitnesses is essentially random in space.

The other evolvability measures can be derived in similar fashion and give good agreement with the fitness evolvability portraits derived from empirical simulation of the NK landscapes. Figure 12 shows empirical data for the evolvability metrics E_a, E_b, E_c, E_d . First, we correctly identify the increasing *modality* of the spaces at higher fitnesses with increasing K ; as fitness increases, the probability of nondeleterious mutations E_a tails off faster for high K than for low K , showing that the number of local optima increase extremely rapidly above fitnesses of 0.5 for $K = N - 1$. Second, the


³The offspring-parent correlation is simply $\rho = 1 - \frac{K+1}{N}$ (Weinberger, 1990) with the correlation length $\tau = -1/\ln(\rho)$.




(a) Probability of a nondeleterious mutation E_a



(b) Expected fitness over all mutations E_b



(c) Expected fitness over top quartile of mutations E_c



(d) Expected fitness over bottom quartile of mutations E_d


Figure 13: Fitness evolvability portraits for the terraced NK landscapes with $N = 25$, $K \in \{0, 6, 12, 24\}$, $F = 4$. The evolvability metrics were calculated from a random sample set of solutions.

portraits correctly predict the increasing *ruggedness* of the spaces with increasing K ; the expected offspring fitness graphs (E_b shown in Figure 12(c)) show decreasing gradient with increased K (remember from Section 3.5.3 that this gradient is proportional to the autocorrelation function).


The next section applies the evolvability analysis to the tunably neutral terraced NK landscapes.

6 Evolvability of a Tunably Neutral Landscape


In the previous section, we saw how we can discriminate between landscapes of varying ruggedness using the fitness evolvability portraits derived in Section 3.1. In this section, we derive portraits for the tunably neutral terraced NK landscape (Section 4) in order to discriminate between landscapes of varying ruggedness in the presence of




(a) Probability of a nondeleterious mutation E_a for $K = 0$



(b) Expected fitness over all mutations E_b for $K = 0$



(c) Probability of a nondeleterious mutation E_a for $K = 12$



(d) Expected fitness over all mutations E_b for $K = 12$

Figure 14: Fitness evolvability portraits for the terraced NK landscapes with $N = 25$, $K \in \{0, 12\}$, $F \in \{2, 3, 6, \infty\}$. The evolvability metrics were calculated from a random sample set of solutions.

neutrality and landscapes of varying neutrality.

Figure 13 shows the four fitness evolvability portraits for the terraced NK landscape with the number of possible locus fitness terraces fixed at $F = 4$ ($N = 25$ and $K \in \{0, 6, 12, 24\}$). We see that the presence of a large degree of neutrality in the system does not affect the ability of the portraits to distinguish between differing levels of ruggedness and modality in the system. In particular, the ruggedness of the landscapes are not affected by the increased neutrality (compare Figure 13(b) to 12(b)), but as expected, the probability of reaching local optima at a given fitness is much lower for all landscapes (compare Figure 13(a) to 12(a)).

Figure 14 shows portraits for the terraced NK landscapes with varying neutrality and fixed ruggedness (the probability of obtaining a nondeleterious mutation and

the expected fitness of all mutations for $N = 25$, $K \in \{0, 12\}$, $F \in \{2, 3, 6, \infty\}$). We see that the expected mutation fitness (Figures 14(b) and 14(d)) does not change with differing levels of neutrality in the system; tallying with the results that the autocorrelation and ruggedness do not change with neutrality in the NK landscapes. However, the probability of obtaining a nondeleterious mutation (Figures 14(a) and 14(c)) does show such change. As $F \rightarrow 2$, neutrality increases and the probability of obtaining a nondeleterious mutation increases. For $K = 0$, even at high fitnesses there are still on average $1/F$ neutral mutations. For high K , this probability tends to zero at high fitnesses as all $K + 1$ loci fitness affected by the mutation need to show a neutral mutation. However, this decrease in the probability of finding nondeleterious mutations is slower for landscapes with more neutrality. The difference is significant: at a fitness of 0.6, roughly 37% of mutations in the $F = 2$ landscape are nondeleterious compared with roughly 13% of such mutations for the non-neutral $F = \infty$ landscape. At a fitness of 0.7, the corresponding percentages are roughly 18% and 0%. Thus, in the highly neutral $F = 2$ landscape, the probability of the search process reaching a local optimum is significantly smaller than the probability of reaching a local optimum in the non-neutral $F = \infty$. Rather than sticking in local optima, the search process can explore more of the space along neutral networks, eventually reaching higher fitness solutions.

Thus the fitness evolvability portraits do indeed differentiate between landscapes of both varying ruggedness (with constant neutrality) and varying neutrality (with constant ruggedness). In particular, two general features are seen for the terraced NK landscapes with the portrait descriptions. First, as neutrality increases, the landscape ruggedness does not change, as evidenced by the expected offspring fitness portrait (this result is also shown with the correlation lengths shown in Figure 10). Second, as neutrality increases, the number of nondeleterious mutations increases *at all levels of fitness* and for all K . Thus as expected, the number of local optima falls with increasing neutrality, but also the number of local optima decrease at all fitness levels in the space.

In the next section, we show that fitness evolvability portraits based on samples collected during simple hill-climbing optimization show the same features as when based on the random samples used in the previous two sections. This is crucial for problems with extremely skewed solution fitness distributions for which random sampling is inappropriate and biased sampling techniques must be used. If we are to describe the landscape structure of such problems, the portraits must be robust when based on such biased samples.

7 Online Sampling Evolvability

In the previous sections, we have investigated empirically derived evolvabilities for the tunably rugged and tunably neutral terraced NK landscapes through random sampling of the space of all solutions. This random sampling technique works well with the NK landscapes where solution fitnesses are defined as the linear sum of all loci fitnesses; due to the central limit theorem, the solution fitnesses will be approximately normally distributed. However, in many problems, such normally distributed solution fitnesses will not be encountered, and measures based on random sampling of the space may in general be less successful in describing the landscape.

With such skewed solution fitness distributions, it may be necessary to bias the collected sample through keeping only a percentage of solutions found at each fitness and define the landscape description over this biased sample. With even more extremely skewed distributions, it may be necessary to collect a biased sample through some direct search optimization procedure such as a simple hill-climber. For example, Smith

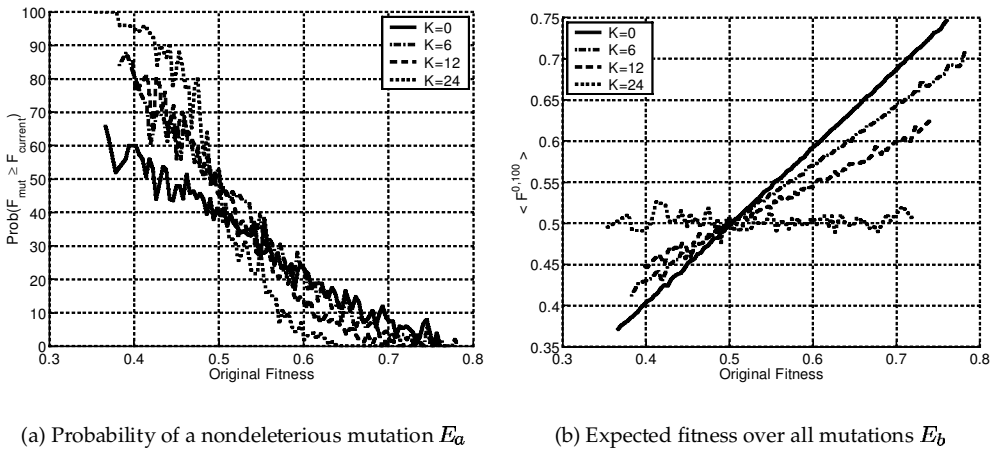


Figure 15: Fitness evolvability portraits for the NK landscapes with $N = 25$, $K \in \{0, 6, 12, 24\}$. The evolvability metrics were calculated from a sample set of solutions collected during hill-climbing.

et al. (2001a) find only 0.0001% of randomly generated solutions have fitness above 50% of the maximum in a neural network robot control problem despite this fitness being relatively easy to reach using optimization techniques.

In this section, we show that the fitness evolvability portraits presented in the previous sections still describe the general features of the terraced NK landscapes when based on a biased sample collected using a (1+1) evolutionary strategy hill-climber (Rechenberg, 1973). 100 runs of the hill-climber were performed for each parameter setting (generating a new landscape for each run). From an initial randomly generated solution, random mutations were applied (using both single bit mutation and mutation probability per bit gives similar results) with nondeleterious mutations accepted and deleterious mutations rejected. All new encountered genotypes were saved for analysis, and the hill-climber stopped after 1,000 mutations had been tried. The following analysis uses the saved samples over each parameter setting.


Figure 15 shows the probability of a nondeleterious mutation and the expected mutation fitness over the NK landscape with $N = 25$, $K \in \{0, 6, 12, 24\}$ for the biased hill-climber sample. As seen with the fitness evolvability portraits based on random sampling, the ruggedness of the landscape, as measured through the gradient of the expected mutation fitness against fitnesses, increases with K . Also, the number of local optima increases with both K and the level of fitness.

Figure 16 shows the probability of a nondeleterious mutation for the terraced NK landscapes with $N = 25$, $F \in \{2, 3, 6, \infty\}$, $K \in \{0, 12\}$ for the biased hill-climber sample. Again we see the same results as for the unbiased random sample (Section 6); as expected, the more neutral landscapes with lower numbers of terraces F show greater probability of nondeleterious mutations over the entire fitness range.

The results for the biased hill-climber sample over both the tunably rugged (Figure 15) and the tunably neutral landscapes (Figure 16) are equivalent to the results seen for the unbiased random sample used in Sections 5 and 6. Thus the fitness evolvability portraits are indeed robust over the use of biased and unbiased samples. In the next



(a) Probability of a nondeleterious mutation E_a for $K = 0$



(b) Probability of a nondeleterious mutation E_a for $K = 12$

Figure 16: Fitness evolvability portraits for the terraced NK landscapes with $N = 25$, $K \in \{0, 12\}$, $F \in \{2, 3, 6, \infty\}$. The evolvability metrics were calculated from a sample set of solutions collected during hill-climbing.

section, we derive the fitness evolvability portraits for a real search space from the evolutionary hardware domain, evolution of a digital inverter circuit, and compare with results from optimization runs.

8 An Evolutionary Hardware Problem

In this section, we apply the same evolvability analysis to search spaces corresponding to real engineering applications and compare with results from optimization runs. It should be emphasized that the specific implementation details outlined below are unimportant; what should be stressed is that the two different search spaces correspond to two solution representations, “multiplex” and “direct”, for the same evolutionary electronics problem.

In Layzell (1999, 2001), various electronic circuits were evolved directly in hardware using an Evolvable Motherboard (EM); a purpose-built research platform consisting of bipolar transistors connected to a triangular matrix of reconfigurable analog switches. The triangular switch matrix allows all possible combinations of interconnection between the transistors, power supply, and I/O, but by altering the genotype-to-phenotype mapping, various more restrictive interconnection architectures can be investigated. This research explored two different architectures: direct mapping, for which each genotype bit corresponds to the on/off state of an EM switch, and multiplex mapping, for which the genotype specified the matrix position of a limited number of EM switches as well as their state. Figure 17 shows how the two mappings differ in the way that the state of the EM switches are specified.

Circuit evolution was carried out both in the noisy environment of physical hardware and a noise-free environment attained by simulating the EM with proprietary electronics design software. In general, this research required search-spaces of the order 2^{1000} in size, with the multiplex mapping consistently proving more conducive to evolutionary search. Multiplex mapping was originally designed to improve the evo-

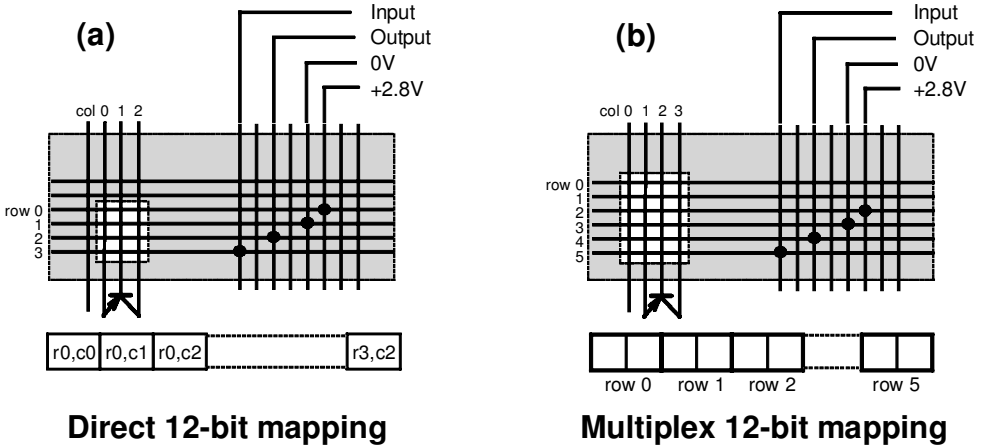


Figure 17: The two 12-bit evolutionary electronics mappings. Direct mapping, for which each genotype bit corresponds to the on/off state of an EM switch, and multiplex mapping, for which the genotype specified the matrix position of a limited number of EM switches as well as their state.






Figure 18: \log_{10} fitness distribution for the 12-bit direct and multiplex evolutionary electronics spaces. Note that the fitness function for the digital inverter circuit assigns negative fitness to solutions that produce output increasing with input, zero fitness to solutions with constant output, and positive fitness to solutions with output decreasing with increasing input. The majority of solutions have near zero fitness, while the multiplex mapping space has many more positive fitness solutions than the direct mapping space.

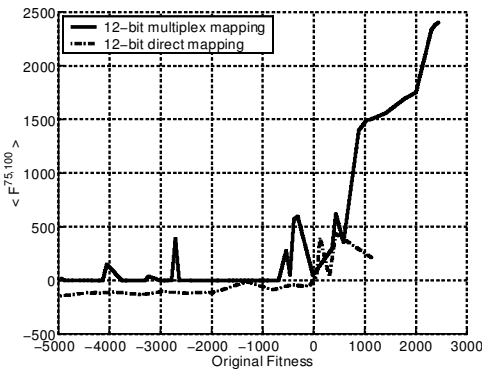
lutionary search from an electronics perspective (by reducing the likelihood of short-circuits on the transistor pins), but a study was later undertaken to investigate the relative merits of the two mappings from an evolutionary theoretic viewpoint (Layzell, 2001). This study used a minimal set of components to evolve a digital inverter with search-space sizes between 2^8 and 2^{16} – small enough to conduct exhaustive searches. Although small, they are real electronic spaces and exhibit all the characteristics (such as configuration-dependent noise levels and power-supply related epistasis) of the




(a) Probability of a nondeleterious mutation E_a



(b) Expected fitness over all mutations E_b



(c) Expected fitness over top quartile of mutations E_c



(d) Expected fitness over bottom quartile of mutations E_d

Figure 19: Fitness evolvability portraits for the evolutionary hardware multiplex and direct mappings. The metrics were calculated from exhaustive sampling of solutions.

much larger spaces also explored by Layzell (1999, 2001).

The fitness function for the digital inverter circuits assigns negative fitness to solutions that produce output increasing with input, zero fitness to solutions with constant output, and positive fitness to solutions with output decreasing with increasing input. It should be emphasized that the distribution of fitnesses will not necessarily be equal for the multiplex and direct mapping spaces and is markedly non-normal. In particular, the majority of the solutions have zero fitness due to constant output, and the number of positive fitness solutions is small. Figure 18 shows the logged fitness distribution for the multiplex and direct mapping spaces; note that the majority of solutions have near zero fitness, while the multiplex mapping space has many more positive fitness solutions.

Figure 19 shows the four fitness evolvability portraits for the spaces. Several points can be made. First, the portraits are much noisier than the corresponding NK por-

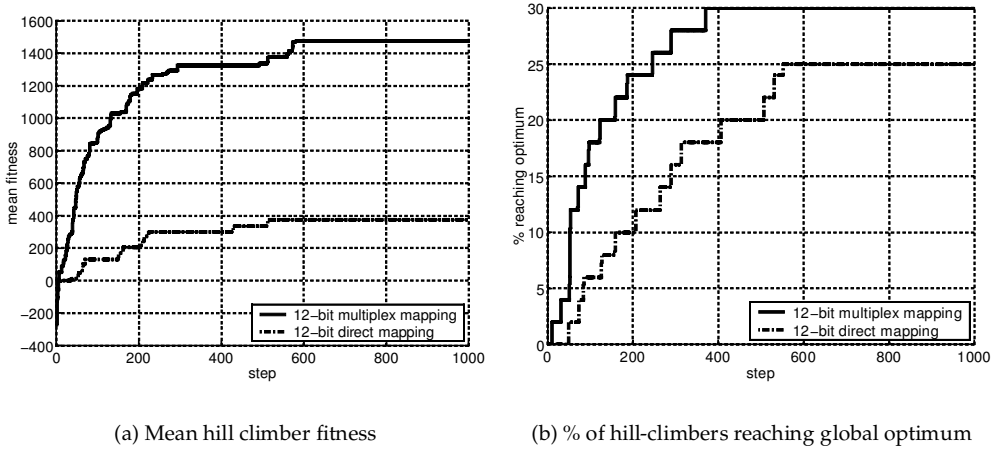


Figure 20: (a) Mean fitness reached through 1+1 evolutionary strategy hill-climbing, and (b) percentage of hill-climbers that reach the global optimum. Data calculated over 50 runs of a non-greedy hill-climber (accepting the first nondeleterious mutation).

traits; the fitness distributions for these evolutionary electronics spaces are markedly non-normal, and many fitness levels contain only a few solutions, producing large variation in the calculated evolvability. Second, as evidenced by the gradients of the expected offspring fitness portrait (Figure 19(b)), the spaces are of similar ruggedness to each other, which is supported by autocorrelation analysis (not shown). Third, the distribution and density of local optima varies markedly over the space (see the probability of nondeleterious mutations, Figure 19(a)). At fitnesses at and below zero, the majority of mutations are nondeleterious, and there are few local optima. Once above zero fitness however, the character of the space changes dramatically; there are now far fewer neutral or fitter mutations and many local optima.

The general features of the two spaces are thus similar in terms of ruggedness and modality, but we can identify differences at and above fitnesses of zero. In particular, the probabilities of nondeleterious mutations are higher for the multiplex spaces than for the direct spaces (Figure 19(a)), thus we might expect the hill-climbers to reach high fitnesses faster for the multiplex mapping. This is seen in Figure 20(a); the mean hill-climb fitness for the multiplex space is significantly larger than for the direct space. However, this is heavily weighted by the higher global optimum in the multiplex space, so Figure 20(b) shows the percentage of hill-climbers reaching the global optimum; roughly equal numbers of hill-climbers in the two spaces reach the global optimum (30% for the multiplex, 25% for the direct mapping), but the hill-climbers in the multiplex space reach the optimum quicker than the direct space hill-climbers. This tallies with the suggestion from the fitness evolvability portraits, that the multiplex spaces contain fewer deleterious mutations at higher fitnesses.

In the final section, we discuss the results, relevance, and limitations of the fitness evolvability portraits applied to real spaces.

9 Discussion

In this paper, we have developed the notion of evolvability based on the fitness distribution of the offspring of sampled solutions. Averaging evolvability over a population of equal fitness solutions allows us to build up fitness evolvability portraits of a landscape, which we have shown can be used to describe general features of ruggedness, modality, and neutrality in a set of tunably rugged and tunably neutral landscapes. We have also shown that such portraits can be based on samples of solutions collected both through random sampling and through online biased sampling. Finally, we have derived the fitness evolvability portraits for two real evolutionary electronics search spaces and suggested that some characteristics of the two search spaces can be linked to the ease of finding good solutions.

The primary aim of this work is to develop techniques that can be used to describe features of fitness landscapes for real problems. In previous work, we have shown that standard descriptions based on random sampling of the space, such as the autocorrelation function, fail when applied to spaces with highly skewed fitness distributions (Smith et al., 2001a). Although such techniques can be used with biased sample sets, it is unclear to what extent their underlying assumptions are violated by the heterogeneous anisotropic nature of certain problem spaces, e.g., evolutionary electronics landscapes. This is also an issue with the techniques described here; we are using methods based on *local* space features to predict *global* space properties. This implicit assumption holds in homogeneous problem spaces such as the *NK* landscapes, but the evolutionary electronics spaces investigated show no such homogeneity. In such cases of extremely heterogeneous spaces, it may be unrealistic to expect to identify general features; variation in the populations of equal fitness solutions may well be swamped through the averaging procedures used. Instead, the techniques should be used as part of a larger investigation for description, and hopefully characterization, of some fitness landscape.

A second issue that arises when applying the techniques to heterogeneous fitness landscapes is that of sparse sampling. In the evolutionary electronics case, many levels of fitness were extremely poorly represented. Although in this case, exhaustive sampling was used, the argument applies to much larger spaces where similar distributions of solution fitnesses may be obtained through random or biased sampling. Sparsely sampled fitness levels will produce large variation in the evolvability portraits (seen in the electronics spaces in Section 8), making it difficult to correctly interpret the portraits. Further, fitness levels that are not represented at all in the collected sample, e.g., extremely high fitness levels, will clearly have no impact on the evolvability portraits. One future area of work is to develop techniques for predicting the *accessibility* of such high fitness levels; one such measure could potentially be based on the increase in probability of deleterious mutations as fitness increases.

In the work reported in this paper, we have concentrated on deriving fitness evolvability portraits based on the fitness of offspring in a small volume surrounding the sampled solutions, with the volume typically defined through single bit mutation. There is no reason in principle that the techniques could not be extended to incorporate crossover and other operators; only the fitness distribution of offspring is required to derive the portraits (although there are clearly issues to be explored regarding the other solution(s) involved in recombination). In the same vein, the portraits could in principle be used to investigate the suitability of operators for a particular problem, comparing features of the landscapes defined by two or more different sets of operators.

Finally, it should be emphasized that the techniques of evolvability and fitness portraits developed in this paper should be seen as complementary to other methods for describing and characterizing fitness landscapes. No single measure or description can possibly characterize any high-dimensional heterogeneous search space; the fitness evolvability portraits are presented as a useful method to investigate the properties of the fitness landscape over the entire fitness range.

Acknowledgments

The authors would like to thank the three anonymous reviewers. We also thank Andy Philippides, Lionel Barnett, Inman Harvey, and the members of the Centre for Computational Neuroscience and Robotics (<http://www.cogs.sussex.ac.uk/ccnr/>) for constructive comments and discussion. We would also like to thank the Sussex High Performance Computing Initiative (<http://www.hpc.sussex.ac.uk/>) for computing support. TS is funded by a British Telecom Future Technologies Group (<http://innovate.bt.com/projects/ftg.htm>) sponsored Biotechnology and Biology Science Research Council Case Award.

References

- Altenberg, L. (1994). The evolution of evolvability in genetic programming. In Kinnear Jr, K., editor, *Advances in Genetic Programming*, Chapter 3, pages 47–74, MIT Press, Cambridge, Massachusetts.
- Barnett, L. (1997). *Tangled webs—evolutionary dynamics on fitness landscapes with neutrality*. Master's thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex, UK.
- Barnett, L. (1998). Ruggedness and neutrality: The NKp family of fitness landscapes. In Adami, C. et al., editors, *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pages 18–27, MIT Press, Cambridge, Massachusetts.
- Barnett, L. (2001). Statistical properties of the NKp family of fitness landscapes. In press.
- Binder, K. and Young, A. (1986). Spin Glasses: Experimental Facts, Theoretical Concepts and Open Questions. *Reviews of Modern Physics*, 58(4):801–976.
- Burch, C. and Chao, L. (2000). Evolvability of an RNA Virus is Determined by its Mutational Neighbourhood. *Nature*, 406:625–628.
- Cavalli-Sforza, L. and Feldman, M. (1976). Evolution of continuous variation: Direct approach through joint distribution of genotypes and phenotypes. *Proceedings of the National Academy of Sciences*, 73:1689–1692.
- Coyne, J., Barton, N., and Turelli, M. (1997). Perspective: A critique of Sewall Wright's shifting balance theory of evolution. *Evolution*, 51(3):643–671.
- Davidor, Y. (1991). Epistasis variance: A viewpoint on GA-hardness. In Rawlins, G., editor, *Foundations of Genetic Algorithms (FOGA'1)*, pages 22–35, Morgan Kaufmann, San Mateo, California.
- Dawkins, R. (1989). The evolution of evolvability. In Langton, C., editor, *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Volume VI of Santa Fe Institute Studies in the Sciences of Complexity, pages 201–220, Addison-Wesley, Redwood, California.
- Derrida, B. (1981). Random energy model—an exactly solvable model of disordered-systems. *Physical Review B*, 24(5):2613–2626.
- Ebner, M. et al. (2001). On Neutral Networks and Evolvability. In *Proceedings of the 2001 Conference on Evolutionary Computation: CEC2001*, pages 1–8, IEEE Press, Piscataway, New Jersey.

- Eldredge, N. and Gould, S. (1972). Punctuated equilibria: An alternative to phyletic gradualism. In Schopf, T., editor, *Models in Paleobiology*, pages 82–115, Freeman, San Francisco, California.
- Elena, S., Cooper, V., and Lenski, R. (1996). Punctuated evolution caused by selection of rare beneficial mutations. *Science*, 272:1802–1804.
- Fogel, L., Owens, A., and Walsh, M. (1966). *Artificial Intelligence through Simulated Evolution*. John Wiley and Sons, New York, New York.
- Forst, C., Reidys, C., and Weber, J. (1995). Neutral networks as model-landscapes for RNA secondary-structure folding landscapes. In Moran, F. et al., editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life, ECAL95*, pages 128–147, Springer, Berlin, Germany.
- Gould, S. and Eldredge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology*, 3:115–151.
- Grüner, W. et al. (1996). Analysis of RNA sequence structure maps by exhaustive enumeration: I Neutral networks, II Structures of neutral networks and shape space covering. *Monatshfte Chem.*, 127:355–374, 375–389.
- Harvey, I. and Thompson, A. (1996). Through the labyrinth evolution finds a way: A silicon ridge. In Higuchi, T., editor, *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware (ICES'96)*, pages 406–422, Springer, Berlin, Germany.
- Holland, J. (1992). *Adaptation in Natural and Artificial Systems*. Second edition. MIT Press, Cambridge, Massachusetts.
- Hordijk, W. (1996). A measure of landscapes. *Evolutionary Computation*, 4(4):335–360.
- Huynen, M. (1996). Exploring phenotype space through neutral evolution. *Journal of Molecular Evolution*, 43:165–169.
- Huynen, M., Stadler, P., and Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Sciences*, 93:394–401.
- Jones, T. and Forrest, S. (1995). Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In Eshelmann, L., editor, *Proceedings of the Sixth International Conference on Genetic Algorithms (ICGA95)*, pages 184–192, Morgan Kaufmann, San Mateo, California.
- Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford, UK.
- Kimura, M. (1983). *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK.
- Kirschner, M. and Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences*, 95:8420–8427.
- Koza, J. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, Massachusetts.
- Lawler, E. et al. (1985). *The Travelling Salesman Problem*. John Wiley and Sons, New York, New York.
- Layzell, P. (1999). Reducing hardware evolution's dependency on FPGAs. In *Proceedings of the Seventh International Conference on Microelectronics for Neural, Fuzzy and Bio-Inspired Systems*, pages 171–178, IEEE Press, Piscataway, New Jersey.
- Layzell, P. (2001). *Hardware Evolution: On the Nature of Electronic Circuits Derived Through Artificial Evolution*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex, UK.

- Lipsitch, M. (1991). Adaptation on rugged landscapes generated by iterated local neighboring genes. In Belew, R. and Booker, L., editors, *Proceedings of the Fourth International Conference on Genetic Algorithms (ICGA91)*, pages 128–135, Morgan Kaufmann, San Mateo, California.
- Manderick, B., Weger, M., and Spiessens, P. (1991). The genetic algorithm and the structure of the fitness landscape. In Belew, R. and Booker, L., editors, *Proceedings of the Fourth International Conference on Genetic Algorithms (ICGA91)*, pages 143–150, Morgan Kaufmann, San Mateo, California.
- Marr, D. (1976). Artificial intelligence—a personal view. Technical Report AIM 355, MIT AI Lab, MIT, Cambridge, Massachusetts.
- Marrow, P. (1999). Evolvability: Evolution, computation, biology. In Wu, A., editor, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program (GECCO-99 Workshop on Evolvability)*, pages 30–33, Morgan Kaufmann, San Mateo, California.
- Naudts, B. and Kallel, L. (2000). A comparison of predictive measures of problem difficulty in evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 4(1):1–15.
- Newman, M. and Engelhardt, R. (1998). Effects of selective neutrality on the evolution of molecular species. *Proceedings of the Royal Society of London, B*, 265:1333–1338.
- Provine, W. (1986). *Sewall Wright and Evolutionary Biology*. University of Chicago Press, Chicago, Illinois.
- Rechenberg, I. (1973). *Evolutionsstrategie*. Friedrich Frommann Verlag, Stuttgart, Germany.
- Reidys, C. and Stadler, P. (2001). Neutrality in Fitness Landscapes. *Applied Mathematics and Computation*, 117(2–3):321–350.
- Rosé, H., Ebeling, W., and Asselmeyer, T. (1996). The Density of States—a Measure of the Difficulty of Optimization Problems. In Voigt, H.-M. et al., editors, *Proceedings of the Fourth Conference on Parallel Problem Solving from Nature: PPSN IV. Lecture Notes in Computer Science*, 1141:208–217, Springer, Berlin, Germany.
- Shackleton, M., Shipman, R., and Ebner, M. (2000). An Investigation of Redundant Genotype-Phenotype Mappings and their Role in Evolutionary Search. In *Proceedings of the 2000 Congress on Evolutionary Computation: CEC2000*, pages 493–500, IEEE Press, Piscataway, New Jersey.
- Sibani, P. and Pederson, A. (1999). Evolution Dynamics in Terraced NK Landscapes. *Europhysics Letters*, 48:34–52.
- Smith, T., Husbands, P., and O’Shea, M. (2001a). Not measuring evolvability: Initial exploration of an evolutionary robotics search space. In *Proceedings of the 2001 Congress on Evolutionary Computation: CEC2001*, pages 9–16, IEEE Press, Piscataway, New Jersey.
- Smith, T., Husbands, P., and O’Shea, M. (2001b). Evolvability, neutrality, and search difficulty in evolutionary robotics. *Biosystems*. Submitted to special issue on evolvability. Also available as Cognitive Science Research Paper 535, University of Sussex, Sussex, UK.
- Smith, T., Husbands, P., and O’Shea, M. (2001c). Neutral Networks and Evolvability with Complex Genotype-Phenotype Mapping. In Kelemen, J. and Sosík, P., editors, *Advances in Artificial Life, Sixth European Conference: ECAL 2001*, pages 272–281, Springer, Berlin, Germany.
- Smith, T., Husbands, P., and O’Shea, M. (2001d). Neutral Networks in an Evolutionary Robotics Search Space. In *Proceedings of the 2001 Congress on Evolutionary Computation: CEC2001*, pages 136–145, IEEE Press, Piscataway, New Jersey.
- Sniegowski, P., Gerrish, P., and Lenski, R. (1997). Evolution of high mutation rates in experimental populations of *E. coli*. *Nature*, 387:703–705.

- Stadler, P. (1996). Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20:1–45.
- Stadler, P. and Schnabl, W. (1992). The landscape of the traveling salesman problem. *Physics Letter A*, 161:337–344.
- Stadler, P. and Wagner, G. (1998). Algebraic theory of recombination spaces. *Evolutionary Computation*, 5(3):241–275.
- Taddei, F. et al. (1997). Role of mutator alleles in adaptive evolution. *Nature*, 387:700–702.
- Thompson, A. (2001). Neutrality in Evolutionary Hardware Experiments. In press.
- Turney, P. (1999). Increasing evolvability considered as a large-scale trend in evolution. In Wu, A., editor, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program (GECCO-99 Workshop on Evolvability)*, pages 43–46, Morgan Kaufmann, San Mateo, California.
- van Nimwegen, E., Crutchfield, J., and Huynen, M. (1999). Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences*, 96:9716–9720.
- Vassilev, V. (1997). Information analysis of fitness landscapes. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life: ECAL97*, pages 116–124, MIT Press, Cambridge, Massachusetts.
- Vassilev, V. and Miller, J. (2000). The advantages of landscape neutrality in digital circuit evolution. In Miller, J. et al., editors, *Proceedings of the Third International Conference on Evolvable Systems: From Biology to Hardware (ICES'2000)*. *Lecture Notes in Computer Science*, 1801:252–263, Springer, Berlin, Germany.
- Vassilev, V., Fogarty, T., and Miller, J. (2000). Information characteristics and the structure of landscapes. *Evolutionary Computation*, 8(1):31–60.
- Wagner, G. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976.
- Weinberger, E. (1990). Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336.
- Weinberger, E. (1991). Local Properties of Kauffman's N-K Model: A Tunably Rugged Energy Landscape. *Physical Review A*, 44(10):6399–6413.
- Wilke, C. (2001). Adaptive evolution on neutral networks. *Bulletin of Mathematical Biology*, 63:715–730.
- Wilke, C. et al. (2001). Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, 412:331–333.
- Wright, S. (1932). The Role of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. In Jones, D., editor, *Proceedings of the Sixth International Congress on Genetics*, Volume 1, pages 356–366. Reprinted in Ridley, M. (1997). *Evolution*. pages 32–40, Oxford University Press, Oxford, UK.