

Ruggedness and Neutrality - The NKp family of Fitness Landscapes

Lionel Barnett

Centre for Computational Neuroscience and Robotics
Centre for the Study of Evolution
Department of Cognitive and Computing Sciences

University of Sussex, Brighton BN1 9QH, UK

lionelb@cogs.susx.ac.uk

Abstract

It has come to be almost an article of faith amongst population biologists and GA researchers alike that the principal feature of a fitness landscape as regards evolutionary dynamics is “ruggedness”, particularly as measured by the auto-correlation function. In this paper we demonstrate that auto-correlation alone may be inadequate as a mediator of evolutionary dynamics, specifically in the presence of large scale neutrality. We introduce the NKp family of landscapes (a variant on NK landscapes) which possess the remarkable property that varying the degree of neutrality has minimal effect on the correlation structure. It is demonstrated that NKp landscapes feature *neutral networks* which have a “constant innovation” property comparable with the neutral networks observed in models of RNA secondary structure folding landscapes. We show that evolutionary dynamics on NKp landscapes vary dramatically with the degree of neutrality - at high neutrality the dynamics are characterised by population drift along neutral networks punctuated by transitions between networks. The relevance of these models to natural and artificial evolution is discussed.

Introduction

In attempting to address the dynamics of populations of genotypes evolving on fitness landscapes it appears that a specific scenario has become somewhat ingrained in the collective consciousness of researchers - that of a fitness landscape as a rugged, hilly terrain on which populations perform “hill-climbing”. Selective pressure drags a population towards local peaks of relatively high fitness while mutation and recombination search the surrounding landscape by generating new genotypes. But this poses a problem which affects both the biologist and the GA specialist: if selective pressure is strong enough (relative to the disruptive effects of mutation and recombination) to drag a population up a hill, it is also likely to be strong enough to hold it there! How, then, is an evolving

population to avoid becoming trapped on a local hilltop? For the GA worker seeking to optimise a multi-peaked function this is a practical issue and the literature abounds with schemes to avoid the dilemma (Goldberg 1989). For the biologist it is a serious theoretical conundrum, as populations in nature do not seem (at least on macro-evolutionary time-scales) to suffer this fate. It might be claimed that entrapment can be explained away by co-evolution and environmental change but another possibility must be considered - our picture of a fitness landscape as a rugged hilly terrain is misleading and in need of an overhaul.

In both natural and artificial systems a picture is emerging of populations engaged not in hill-climbing but rather drifting along connected networks of genotypes of equal fitness, with sporadic jumps between networks. These “neutral networks” are of particular significance if they have the “constant innovation” property (see below) - for this raises the possibility that (given enough time) almost *any* possible fitness value can ultimately be attained by the population. The scenario of a population trapped on a local hilltop vanishes. It is this new paradigm of evolutionary dynamics which we examine here. It has yet to make a significant impact on the scientific community.

It is, of course, reasonable to ask (both for natural and artificial evolutionary systems) whether such neutral networks actually occur. Comparatively recent developments in evolutionary theory and molecular biology all point to the importance of *selective neutrality* as a significant factor. This work includes Kimura's neutral theory of molecular evolution (Kimura 1983), Eigen's analysis of the molecular quasispecies (Eigen, McCaskill and Schuster 1989; Nowak and Schuster 1989) and recent developments in the understanding of RNA evolution both *in vitro*, in simulation and analytically (Reidys, Stadler and Schuster 1997; Schuster *et al.* 1994; Baskaran, Stadler and Schuster 1996; Grüner *et al.* 1996).

Neutrality has also been detected in various protein models. In molecular biology it is clear that there is often a high degree of redundancy in the coding from genotype to phenotype - there may indeed be redundancy on several levels; e.g. many nucleotide sequences may code for the same amino acid, while many amino acid sequences may code for functionally equivalent proteins. Such coding redundancy will certainly imply the existence of selectively neutral mutation at the molecular level (Crow and Kimura 1970, Kimura 1983). Whether this takes the form of neutral networks with constant innovation is a (highly non-trivial) empirical question. Research into the structure of RNA folding landscapes suggests strongly that such networks may well be a feature of fitness landscapes in molecular biology.

There is also evidence that neutral networks can appear in the fitness landscapes of “difficult” artificial evolution problems; e.g. in the evolution of neural network robot control systems, on-chip hardware evolution (Thompson 1996; Harvey and Thompson 1996) and CA-based landscapes. Ironically it is customary among GA practitioners deliberately to avoid redundancy in the genetic coding of artificial evolution problems.

The NKp landscapes introduced in this paper have the property that altering the degree of neutrality has minimal effect on the ruggedness of the landscape (as measured by the auto-correlation function). They thus provide a useful test-bed for a comparative study of the effects of ruggedness and neutrality on evolutionary dynamics. We begin with some formal definitions.

Neutrality and Ruggedness

All fitness landscapes in this paper are based on fixed-length binary bit-string genotypes. We thus identify a *fitness landscape* of sequence length N with a *fitness function* $f: Q^N \rightarrow \mathbf{R}^+$ where Q^N denotes the binary N -hypercube and \mathbf{R}^+ is the set of real numbers ≥ 0 . The *fitness* of a genotype $g \in Q^N$ is then given by $f(g)$. There is a natural metric, *Hamming distance*, on Q^N defined by: $h(g, g') \equiv$ number of loci (bit-positions) at which g and g' differ. Hamming distance is often referred to in terms of *mutation*. If g, g' are hamming distance d apart we call g' a (d -bit) mutation of g (and vice-versa).

Neutrality

We call a (1-bit) mutation g' of g *neutral* iff $f(g') = f(g)$. This relationship induces a partitioning of Q^N whereby g and g' are in the same equivalence class iff there is a sequence of neutral mutations connecting g and g' ; i.e. there are genotypes $g \equiv g^{(0)}, g^{(1)}, g^{(2)}, \dots, g^{(m)} \equiv g'$ such that $g^{(\alpha)}$ is a 1-bit mutation of $g^{(\alpha-1)}$ for $\alpha = 1, 2, \dots, m$ and $f(g) \equiv f(g^{(0)}) = f(g^{(1)}) = f(g^{(2)}) = \dots = f(g^{(m)}) \equiv f(g')$. The *neutral*

networks of the fitness landscape are defined to be the equivalence classes of this partitioning. We can define a coarser partitioning of Q^N by specifying g and g' to be in the same equivalence class iff $f(g) = f(g')$. We refer to the equivalence classes of this partitioning as *neutral sets*; the neutral networks are the connected components of the neutral sets. Although it is the neutral networks which are of direct relevance to evolutionary dynamics the neutral sets are generally easier to handle analytically; furthermore in many cases of interest the neutral sets consist of few connected components. A word of caution: the “network” terminology may well be misleading. If the frequency of neutral mutation is low there are likely to be very many neutral networks comprising a few, or even single genotypes. Even if there is high neutrality the neutral networks may not resemble networks as much as “clusters”. The *neutral degree* of a genotype g , denoted by $v(g)$ is defined to be the number of neutral mutations of g .

We shall be interested in some notion of *percolation* for neutral networks. It is by no means obvious in what sense percolation may hold relevance for evolutionary dynamics. While it is feasible to transfer the graph-theoretical definition directly to neutral networks, it seems to this author that the related (but distinct) property of “constant innovation rate” introduced by (Huynen 1996) in the context of RNA folding landscapes is likely to be more pertinent; the reasons will hopefully become clear from the discussion of evolutionary dynamics in a later section. Random walks are performed on neutral networks (“neutral walks”) and previously unseen phenotypes (“innovations”) accumulated. The rate of discovery of innovations is then compared to the discovery rate for random walks on the landscape not constrained to a neutral network. Since we are not dealing with phenotypes (in the sense of an intermediate mapping between genotype and fitness) we identify phenotype directly with fitness and consider an innovation to be the discovery of a genotype of previously un-encountered *fitness*.

We thus say that a neutral network has the *constant innovation property* if: (I) the rate of discovery of innovations remains approximately constant for a reasonably large number of steps - what Huynen terms “perpetual innovation” - and (II) the rate of discovery is comparable with that of an unconstrained random walk. Below we investigate this property rather than conventional percolation. It should be noted that constant innovation is indeed distinct from percolation - it is not difficult to construct fitness landscapes with neutral networks that percolate in the graph-theoretical sense, but fail one or both of the above criteria (Jakobi 1996).

Ruggedness

The most frequently encountered measure of ruggedness of a fitness landscape is the *auto-correlation function*. It is often defined in terms of fitness values at successive steps along random walks (Weinberger 1990; Kauffman 1993) but, as remarked in (Stadler 1996) "...it seems to be rather contrived to invoke a stochastic process in order to characterise a given function [i.e. the fitness function] defined on a finite set". We thus use the definition below, apparently first proposed in (Eigen, McCaskill and Schuster 1989).

Let $f: Q^N \rightarrow \mathbf{R}^+$ be a fitness landscape. We first define the mean fitness of the landscape:

$$(1) \quad \bar{f} \equiv 2^{-N} \sum_{g \in Q^N} f(g)$$

the fitness variance:

$$(2) \quad \sigma_f^2 \equiv 2^{-N} \sum_{g \in Q^N} (f(g) - \bar{f})^2$$

and for $d = 1, 2, \dots, N$ the set:

$$(3) \quad Q^N(d) \equiv \{(g, g') \in Q^N \times Q^N \mid \text{lh}(g, g') = d\}$$

Thus $Q^N(d)$ is the set of pairs of genotypes in Q^N Hamming distance d apart. We now define the auto-correlation function to be:

$$(4) \quad \rho(d) \equiv \frac{1}{\sigma_f^2} \frac{1}{|Q^N(d)|} \sum_{(g, g') \in Q^N(d)} (f(g) - \bar{f})(f(g') - \bar{f})$$

for $d = 1, 2, \dots, N$. For consistency we also set $\rho(0) \equiv 1$.

Note: We stress that the quantities \bar{f} , σ_f^2 and $\rho(d)$ are not statistics but simply real numbers associated with a fitness landscape. There appears to be some confusion in the literature on this issue; auto-correlation is sometimes defined by averaging fitness, etc. over *ensembles* of landscapes, e.g. the family of all NKp landscapes with fixed N , K and p (Fontana *et. al.* 1993, Weinberger 1990). In this paper shall we use angle brackets exclusively to indicate that a mean (expectation) is to be taken of a quantity considered as a random variable defined on the sample space of all possible NKp landscapes with fixed N , K and p .

The NKp Family of Fitness Landscapes

We begin by reviewing the construction of an NK landscape (Kauffman 1993). Let $N > 0$ be the genotype length and let $0 \leq K < N$. N and K are fixed during the construction. To each locus on the genotype (i.e. a position $1 \leq i \leq N$ on the bit-string) we assign

independently and at random K distinct loci (excluding the locus under consideration). These loci, plus the locus i itself, are said to be *epistatically linked* to locus i . The idea is that a locus i makes a contribution to the total fitness of a genotype which depends on the value of the allele (0 or 1) at each of the $K+1$ loci epistatically linked to locus i . To each such combination of alleles (there are 2^{K+1} in all) a fitness contribution is assigned as a real number drawn independently and uniformly at random from the interval $[0,1]$. We can think of this as the association of a *fitness table* F_i with each locus i ; for a genotype $g \in Q^N$, given the sequence of alleles $\epsilon_i(g) = a_1 a_2 \dots a_{K+1}$, say, at the loci epistatically linked to locus i the fitness contribution of locus i is given by $F_i(\epsilon_i(g))$, which we also denote by $f_i(g)$.

Finally, to calculate the fitness of an entire genotype the fitness contributions of all loci are summed and the result divided by N to normalise the fitness to the range $[0,1]$. In the above notation:

$$(5) \quad f(g) \equiv \frac{1}{N} \sum_{i=1}^N f_i(g)$$

In summary, an NK landscape is fully specified by N , K , the particular assignment of epistatic links and the contents of the N fitness tables.

It is clear from the construction that there is (almost surely) *no* neutral mutation on an NK landscape - for if two genotypes differ at some locus the respective fitness contributions for that locus will be drawn from different fitness table entries which will (almost surely) be different. There is, however, a "natural" way to introduce neutrality into the model, via the following biologically-inspired argument: the NK model assumes that *every possible* combination of alleles at the loci epistatically linked to a given locus gives rise to a positive contribution to fitness. In nature, however, it seems plausible that many (if not most) combinations of alleles will make *no* contribution to fitness. We could reflect this in the NK model by specifying that the fitness table entry corresponding to such an allelic combination be equal to zero. Thus motivated we proceed as follows: a new parameter $0 \leq p \leq 1$ is introduced to represent the probability that an arbitrarily allelic combination makes no contribution to fitness. Explicitly, when assigning values to the fitness tables we set each entry to 0 independently with probability p . If an entry is not set to zero it is assigned uniformly randomly from the range $[0,1]$ as before. We refer to the resulting landscape as an *NKp landscape*. The case $p = 0$ corresponds to a normal NK landscape, while $p = 1$ corresponds to a completely flat landscape (all fitness table entries are zero).

Please note that due to space constraints most results in the following sub-sections are quoted without proof.

Neutral structure of NKp landscapes

Many of the results quoted below depend on the following observation (which holds almost surely):

$$(6) \quad \text{if } g, g' \in Q^N \text{ then } f(g) = f(g') \Leftrightarrow \text{for all } i \text{ such} \\ \text{that } f_i(g) \neq 0 \text{ we have } \varepsilon_i(g) = \varepsilon_i(g')$$

It is evident that the possibility of neutral mutation arises on an NKp landscape. A calculation yields for the probability that an arbitrary mutation on an arbitrary NKp landscape be neutral:

$$(7) \quad p_{\text{neutral}} = p^2 \left(1 - \frac{K}{N-1} (1-p^2) \right)^{N-1}$$

For large sequence length N this is well approximated by:

$$(8) \quad p_{\text{neutral}} \approx p^2 e^{-K(1-p^2)}$$

Thus for long genotypes the probability that a mutation is neutral is roughly independent of the genotype length and drops off exponentially with increasing epistasis K . A problem with p_{neutral} however, is that neutrality is not spread uniformly over the landscape - in fact an NKp landscape is by no means uniform in its structure, but may be decomposed naturally into subsets corresponding to genotypes with a particular number of zeroes in their fitness tables. Thus for an NKp landscape $f: Q^N \rightarrow \mathbf{R}^+$ and $g \in Q^N$ we define:

$$(9) \quad \zeta(g) \equiv \text{number of loci } i \text{ for which } f_i(g) = 0$$

and for $n = 0, 1, \dots, N$ we define:

$$(10) \quad Z_n(f) \equiv \{g \in Q^N \mid \zeta(g) = n\}$$

Next we note that for $g \in Z_n(f)$ the fitness of g is the sum of $N-n$ independent random variables uniformly distributed on $[0,1]$. Thus the expected fitness for a $g \in Z_n(f)$ for some NKp landscape is:

$$(11) \quad \langle f(g) \rangle_{g \in Z_n(f)} = \frac{N-n}{2N}$$

A calculation gives for the expected neutral degree of a $g \in Z_n(f)$ with $1 \leq n \leq N$:

$$(12) \quad \langle v(g) \rangle_{g \in Z_n(f)} = np(1-q)^{N-n} [1 - (1-p)q]^{n-1}$$

where we have set $q \equiv \frac{K}{N-1}$. Comparing (11) and (12)

we find that for large N the neutral degree of genotypes in an NKp landscape drops off roughly exponentially with increasing fitness - the "higher up" the landscape we go the less neutrality we can expect to encounter. This can also be seen from the observation that for all $g \in Q^N$ we must have $v(g) \leq \zeta(g) \leq N(1 - f(g))$.

We also estimate the sizes of the sets $Z_n(f)$. A

calculation yields:

$$(13) \quad \langle |Z_n(f)| \rangle = 2^N \binom{N}{n} p^n (1-p)^{N-n}$$

It is easy to show that the subsets $Z_n(f)$ have the following useful property: if $\Gamma \subseteq Q^N$ is a neutral set (or indeed a neutral network) then $\Gamma \subseteq Z_n(f)$ for some n . We calculated the expected size of neutral sets contained in $Z_n(f)$ in the following sense: define the random variable S_n (for fixed N, K and p) to be the size of the (unique) neutral set containing a genotype uniformly randomly selected from $Z_n(f)$ for a randomly selected NKp landscape $f: Q^N \rightarrow \mathbf{R}^+$. [Note that S_n does *not* represent the size of a neutral set randomly selected from some $Z_n(f)$ - it will differ due to the variance of neutral set size within the $Z_n(f)$'s]. We have:

$$(14) \quad \langle S_n \rangle =$$

$$\sum_{i=0}^n \sum_{j=0}^i \sum_{k=0}^j \sum_{m=0}^k (-1)^{i+j+k+m} \binom{n}{i} \binom{i}{j} \binom{j}{k} \binom{k}{m} \theta_{N,K,j,m} 2^{j-k} p^i$$

$$\text{where } \theta_{N,K,j,m} \equiv \left[\frac{\binom{N-1-(j-m)}{K}}{\binom{N-1}{K}} \right]^{N-j}$$

if $j - m \leq N - 1 - K$ and zero otherwise. If K is small compared to N it was found empirically that:

$$(15) \quad \langle S_n \rangle \approx \exp \left(N \log(2p) \frac{\left(e^{\frac{2K}{p} \frac{n}{N}} - 1 \right)}{\left(e^{\frac{2K}{p}} - 1 \right)} \right)$$

This indicates that $\langle S_n \rangle$ scales roughly as $\mathcal{O}(e^{e^n})$ for

small K . Fig 1a and 1b below plot the formula (14) for $N = 30$. In Fig 1a $K = 4$ and p is varied; in Fig 1b $p = 0.9$ and K is varied. The formula (14) may also be used to estimate the mean *number* of neutral sets in $Z_n(f)$.

It would appear to be difficult to derive analytically an estimate for the size, number and distribution of the neutral *networks* in NKp landscapes; in lieu the results on neutral sets are helpful - the author suspects that the neutral sets comprise, on the whole, few connected components. The reader is referred to (Barnett 1997) for an empirical analysis of neutral networks on "small" NKp landscapes. See also (Grüner *et al.* 1996) for a detailed analysis of neutral networks on RNA secondary structure folding landscapes.

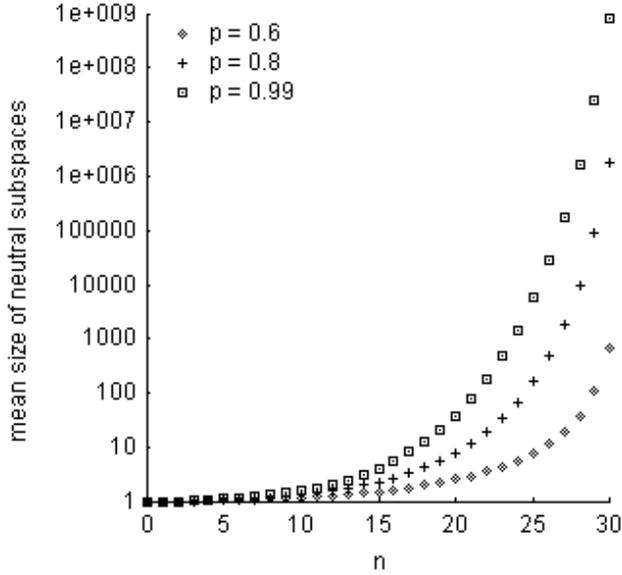


Fig 1a

Estimated mean sizes of neutral sets in $Z_n(f)$ as computed from (14) for $N = 30$, $K = 4$ and several values of p .

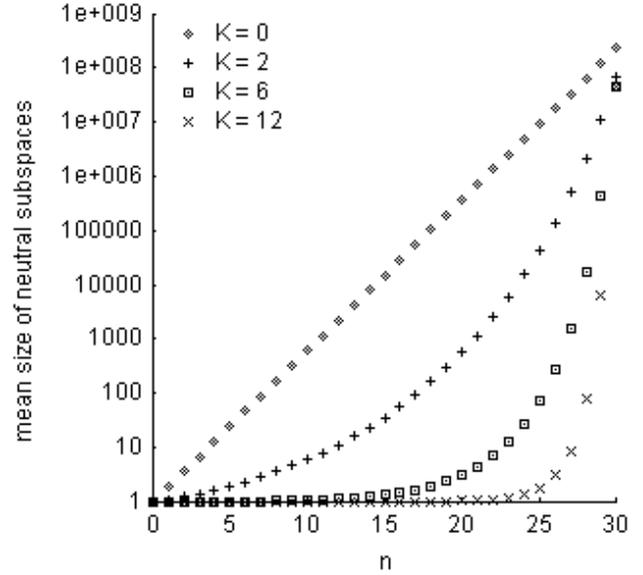


Fig 1b

Estimated mean sizes of neutral sets in $Z_n(f)$ as computed from (14) for $N = 30$, $p = 0.9$ and several values of K .

Correlation structure of NKp landscapes

The most surprising result regarding the correlation structure of NKp landscapes is the minimal effect of the neutrality parameter p on the auto-correlation function. The quantity $\langle \rho(d) \rangle$ (the auto-correlation ensemble mean) was estimated by sampling for a variety of N , K and p values. The results consistently indicated seemingly negligible dependence on p . Indeed, so small is the variation with p , that it was initially thought by the author that $\langle \rho(d) \rangle$ is invariant with respect to p . However more stringent statistical testing, in particular the Student's t-test (Press *et al.* 1992) which measures the significance of a difference of means, indicated a small but significant departure from invariance. The significance is smaller for large N and it may be the case that $\langle \rho(d) \rangle$ is invariant with respect to p in some sense "in the limit" of large N . The ensemble definition of auto-correlation, (as distinct from the ensemble mean - see note above), was also tested for p -invariance. The results suggest that it is a true invariant. The derivation of an analytical expression for the ensemble auto-correlation function in (Fontana *et al.* 1993) suggests that this is indeed the case.

There is one particular class of NKp landscapes for which it is possible to calculate $\rho(d)$ explicitly: this is the case where, out of all the N fitness tables there is only a single entry of non-zero fitness. $\rho(d)$ for these "degenerate" NKp landscapes is given by:

$$(16) \quad \rho_{\text{deg}}(d) = \frac{(P-1)^2 \alpha(d) - 2(P-1)\beta(d) + \gamma(d)}{P(P-1) \binom{N}{d}}$$

where:

$$(17) \quad \alpha(d) \equiv \begin{cases} \binom{L}{d} & \text{if } d \leq L \\ 0 & \text{otherwise} \end{cases}$$

$$(18) \quad \beta(d) \equiv \binom{N}{d} - \alpha(d)$$

$$(19) \quad \gamma(d) \equiv (P-1) \binom{N}{d} - \sum_{k=\text{Max}(1, d-L)}^{\text{Min}(d, K+1)} \binom{K+1}{k} \binom{L}{d-k}$$

and we have set $P = 2^{K+1}$ and $L = N-K-1$. See Appendix A.2 of (Barnett 1997) for details. Surprisingly, this turns out to be a remarkably good estimate for $\langle \rho(d) \rangle$. Fig 2 plots $\rho_{\text{deg}}(d)$ for $N=60$ and a few K values.

It is also worth remarking that the *variance* of $\rho(d)$ (considered, for each d , as a random variable over the sample space of NKp landscapes with fixed N , K and p) is fairly small, particularly for small d . This implies in particular that $\rho_{\text{deg}}(d)$ as defined by (16) is a useful estimate of the auto-correlation for a specific NKp

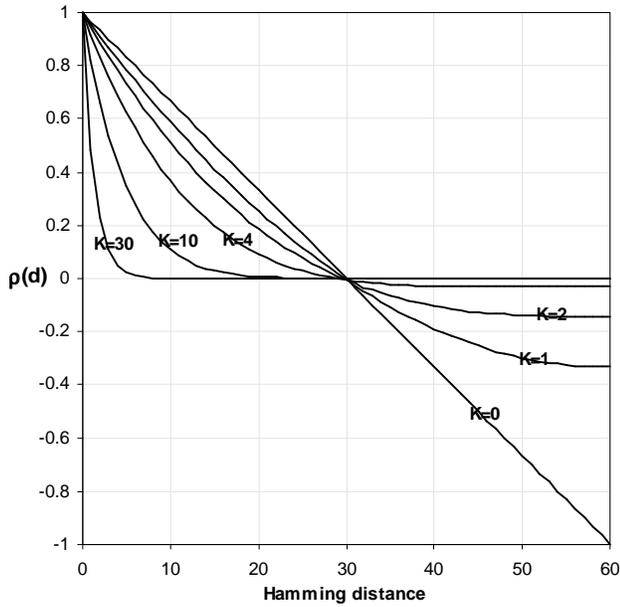


Fig 2
The function $\rho_{\text{deg}}(d)$ for $N=60$

landscape (and indeed for a specific NK landscape). These results appear at first sight to be paradoxical; we might expect that the high proportion of pairs of genotypes with equal fitness would tend to produce higher correlation with increasing p . However, in a certain sense, NKp landscapes actually become *more* rugged with increasing p . Genotypes of high fitness are comparatively rare; hence a near neighbour of a high fitness genotype is likely to be a genotype with more zeros in its fitness tables (particularly when epistasis is high), and hence of far smaller fitness. On average these effects seem to cancel each other out; it was found that the *covariance* of the fitness of pairs of genotypes Hamming distance d apart scales approximately the same as the fitness *variance*. More precisely, it was found empirically that for some function $\phi(d)$ which depends on N and K but not on p , the mean fitness “auto-covariance” of an NKp landscape is given to a high approximation by:

$$(20) \quad \langle \text{cov}(d) \rangle \approx \phi(d)(1-p)(1+3p)$$

The fitness variance is just $\text{cov}(0)$.

Constant innovation on NKp landscapes

Neutral walks were performed on NKp landscapes as previously described. One such test is plotted in Fig 3, which may be compared with the corresponding plot

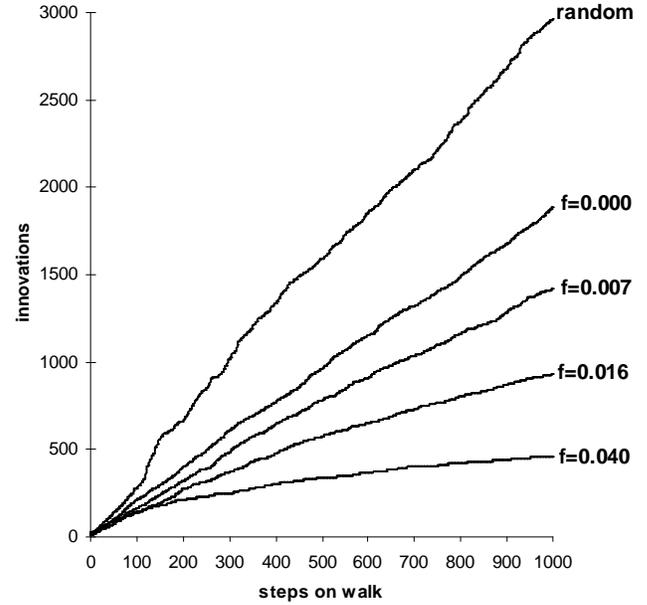


Fig 3
Cumulative innovations on an NKp landscape ($N=60$, $K=14$, $p=0.99$) for 1000-step random and neutral walks. f -values indicate the fitness of the corresponding neutral network.

for an RNA folding landscape in (Huynen 1996). The approximate linearity and slope of the graphs indicate that the constant innovation property does indeed hold on NKp fitness landscapes, at least for neutral networks of modest fitness. As we climb higher up the landscape, however, the innovation rate falls and innovations “peter out” sooner; this is consistent with our earlier findings that the degree of neutrality and expected size of neutral networks fall as fitness increases.

This concludes our analysis of the structure of NKp landscapes. In the next Section we examine their evolutionary dynamics.

Evolutionary Dynamics on NKp Landscapes

Firstly note that in this paper we only consider fixed-size populations evolving under mutation and selection; recombination is not considered. The evolutionary algorithm employed is conventional “fitness-proportional with roulette-wheel selection”, as follows: let population size be M . To construct the population at the next generation from the current population we perform M selections (with replacement) from the current population, such that the probability of a genotype being selected is proportional to its fitness. Every genotype in the new population is then mutated with a per-locus probability m , where m is the (fixed) mutation rate. Typically we take

$M = 200$ and $m = 0.001$.

A previous study (Barnett 1997) investigated in some detail the dynamics of adaptive evolution on several abstract fitness landscapes featuring neutral networks, including NKp landscapes. The picture that emerges is strikingly similar to that described in (Huynen, Stadler and Fontana 1996) for RNA secondary structure folding landscapes and we conjecture that such dynamics are generic for landscapes with neutral networks which have the constant innovation property. A brief summary is as follows: most of the time the population (at reasonably low mutation rates) is largely confined to a specific neutral network, (corresponding to Huynen *et al.*'s "dominant phenotype") on which it drifts at a characteristic rate which is related to the population size, mutation rate and degree of neutrality of the network (see below). During such "metastable" episodes (van Nimwegen, Crutchfield and Mitchell 1997) diffusion is qualitatively similar to diffusion on a flat (i.e. completely neutral) landscape; the latter situation is analysed mathematically in (Derrida and Peliti 1991), where it is found that stochastic effects of selection and mutation typically cause the population to fragment into clusters or sub-populations of genotypes, each cluster sharing a recent common ancestor. Such clustering is also a feature of populations diffusing on neutral networks in non-flat landscapes. Mutation generates new genotypes that explore neighbouring networks. If a genotype of higher fitness (i.e. on a higher-fitness neutral network) is discovered then, if selection pressure is strong enough relative to mutation, the population may, with a certain probability, transfer *en masse* to the higher neutral network. There is also the possibility that a population may, through stochastic effects, "fall off" its current network to a lower-fitness network. The probabilities of attaining or maintaining a given network are related to what has been termed the "phenotypic error threshold" (Forst, Reidys and Weber 1995) by analogy with the classical "genotypic" error threshold for single-peak fitness landscapes (Eigen, McCaskill and Schuster 1989).

NKp landscapes afford a unique opportunity to investigate the form of adaptive evolution with tuneable neutrality. Differences in dynamical behaviour observed for fixed N and K values, but different values of p, cannot be ascribed to the correlation structure as we know this to be virtually invariant under change of p. In this paper we concentrate on one particular aspect of the dynamics, that of population diffusion.

To this end we measured the *diffusion coefficient* of the population *centroid* at successive generations of an evolutionary run. The *centroid* of a population P of genotypes on an N-dimensional hypercube is a real-valued

N-dimensional vector $\mathbf{c} \in \mathbf{R}^N$ defined by $c_i \equiv \frac{1}{N} \sum_{g \in P} g_i$

where $g_i = 0$ or 1 is the allele of g at locus i. \mathbf{c} may be thought of as the centre of mass of the population, considered as a set of points (weighted by their multiplicity in the population) on the hypercube embedded in the vector space \mathbf{R}^N . The diffusion coefficient is defined as the square of the Euclidean distance (in \mathbf{R}^N) travelled by the population centroid per generation. It measures the rate at which the population drifts through the landscape. In practice this quantity tends to fluctuate rapidly from generation to generation; in the graphs below we plot a rolling average over the previous 100 generations to smooth it out. It is possible to estimate the diffusion coefficient by assuming that the diffusion rate will be similar to that on a *flat* landscape of dimension equal to the (mean) neutral degree of the current neutral network (Huynen, Stadler and Fontana 1996). For population size M, mutation rate m and neutral degree ν the estimated value is given by¹:

$$(21) \quad D = \frac{\nu m}{1 + 2Mm}$$

In (Barnett 1997) this was found to be a good estimate for small values of m and a further refinement was suggested.

While the diffusion coefficient tells us about how "fast" the population is wandering about the landscape it does not tell us very much about how "far" it is wandering; e.g. on a single peaked landscape the centroid may move quite rapidly but remain in the locality of the peak. We will be especially interested in the actual distances travelled by the centroid when comparing population dynamics on low and high neutrality landscapes and thus have need for a measure of actual distances travelled by the centroid. To this end we also computed a "time-lagged" diffusion coefficient, which we define to be the square of the distance between the centroid "now" and its position t_{lag} generations previously. In all experiments $t_{\text{lag}} = 100$ was used and the time-lagged coefficient smoothed over 100 generations prior to plotting.

Figs 4 and 5 illustrate typical evolutionary runs over 3000 generations on NKp landscapes for $N=60$, $K=12$ and $p=0.99$ and 0 respectively. In both cases the population size was $M = 200$ and the mutation rate $m=0.001$. Apart from the population mean fitness we also plot the mean neutral degree (for the $p=0.99$ case), diffusion coefficient and time-lagged diffusion. In the $p=0.99$ case the graphs

¹ The formula given in (Huynen, Stadler and Fontana 1996) is for RNA sequences which have four allelic values (it also contains an error - the 5 should be replaced by a 6) and must be adjusted for the binary case. Furthermore, their definition of the centroid works out at twice the magnitude of ours so we must divide their diffusion coefficient by 2.

bears out the picture of evolutionary dynamics outlined above, with periods of metastability punctuated by transitions to higher fitness neutral networks clearly visible. In (Barnett 1997) it is demonstrated that during these periods the population is indeed largely confined to a specific neutral network and that the population drifts and clusters as described above. As we would expect the degree of neutrality falls as fitness increases. The apparently random fluctuations in the diffusion rate reflect the stochasticity of drift and clustering in the population, although overall there is a correlation between diffusion rate and neutral degree as suggested by (21).

During the transitions between neutral networks the time-lagged diffusion increases sharply. This may be ascribed to a “bottleneck effect” as the steep increase in selection pressure occasioned by the discovery of a fitter genotype strongly converges the population around the new genotype. This phenomenon, also known as “hitchhiking”, has been studied in other fitness landscapes, particularly the so-called “Royal Road” landscapes (van Nimwegen, Crutchfield and Mitchell 1997). Mutation then reasserts itself and the population resumes neutral diffusion on the new network. Since (most) neutral networks have the constant innovation property it is unlikely that the drifting population will exhaust the supply of previously unseen (and thus potentially higher fitness) neighbouring genotypes. Evolutionary search may potentially continue unabated; the question is how long it is likely to take before neutral drift discovers a gateway to a higher network (van Nimwegen, Crutchfield and Mitchell 1997).

Fig 5 tells a different story. The landscape is now rugged and multi-peaked with many local optima (Kauffman 1993). The population climbs rapidly up the landscape until it reaches a local optimum at which still higher optima are too rare in the locality to be easily discovered by mutation. At this point the population is effectively trapped - the search for fitter genotypes becomes worse than random search, as the population is confined to the locality of a local optimum. This is indicated by the lagged diffusion, which is significantly lower than for the neutral case.

Conclusions

We have seen that the dynamics of adaptive evolution on fitness landscapes in the presence of neutral networks with the constant innovation property have a distinctly different flavour from the case of ruggedness without neutrality. The scenario of entrapment by local optima is evaded; adaptation is characterised by neutral drift punctuated by transitions between networks rather than local hill-climbing. Furthermore, the formation of sub-populations allows a population to search diverse areas of

a fitness landscape in parallel.

Regarding natural evolution, as argued in the Introduction the issue of selective neutrality is becoming difficult to ignore. Even though the concerns of population geneticists and molecular biologists may often seem far removed from our abstract fitness landscapes it is pointed out in (Huynen, Stadler and Fontana 1996), for example, that one issue of prime interest to evolutionary biology, that of the fixation rate of nucleotide substitutions, is closely related to the population diffusion rate. One general approach that suggests itself is to “reverse engineer” theoretical results; thus a theoretical estimate of the diffusion rate might be deployed to determine the degree of neutrality in a natural evolving system.

There are also pungent implications for artificial evolution. The GA community has long been fixated on correlation structure as the primary factor in the efficacy of evolutionary search. It may be of benefit to GA practitioners to exploit the open-endedness and parallelism implicit in adaptation on neutral networks. One could, for example, envisage schemes whereby the mutation rate is optimised on-line for maximal rates of drift whilst staying below the (local) phenotypic error threshold. Perhaps, also, a change of attitude to the issue of coding of an optimisation problem may be fruitful. Whereas the instinct of many workers is to minimise coding redundancy as an extra burden on a search procedure, they may be thus dooming their search to the fate of entrapment by local optima. Of course it is not to be supposed that there is a “free lunch” involved - redundancy alone certainly does not imply neutral networks with constant innovation. However, it seems that some hard optimisation problems feature neutral networks in a natural way (Thompson 1996). A fascinating area for research would be to investigate in what sense neutral networks might be “intrinsic” to a search problem.

The NKp family of landscapes, aside from the intriguing near-invariance of the auto-correlation function, may hopefully prove to be a useful test-bed for the study of neutral evolution, given the combination of tuneable ruggedness and neutrality. Areas that suggest themselves for further research include the extension of results on neutral networks to include recombination, the issue of “nearly neutral” mutation and the effects of “noisy” fitness. It would also be of great interest to ascertain to what extent the pattern of evolutionary dynamics that emerges from RNA folding landscapes and NKp landscapes is in any sense “generic”. A promising approach may be to employ techniques from statistical mechanics, as applied with some success by (van Nimwegen, Crutchfield and Mitchell 1997) to the Royal Road fitness landscapes, which feature neutrality, albeit without the constant innovation property.

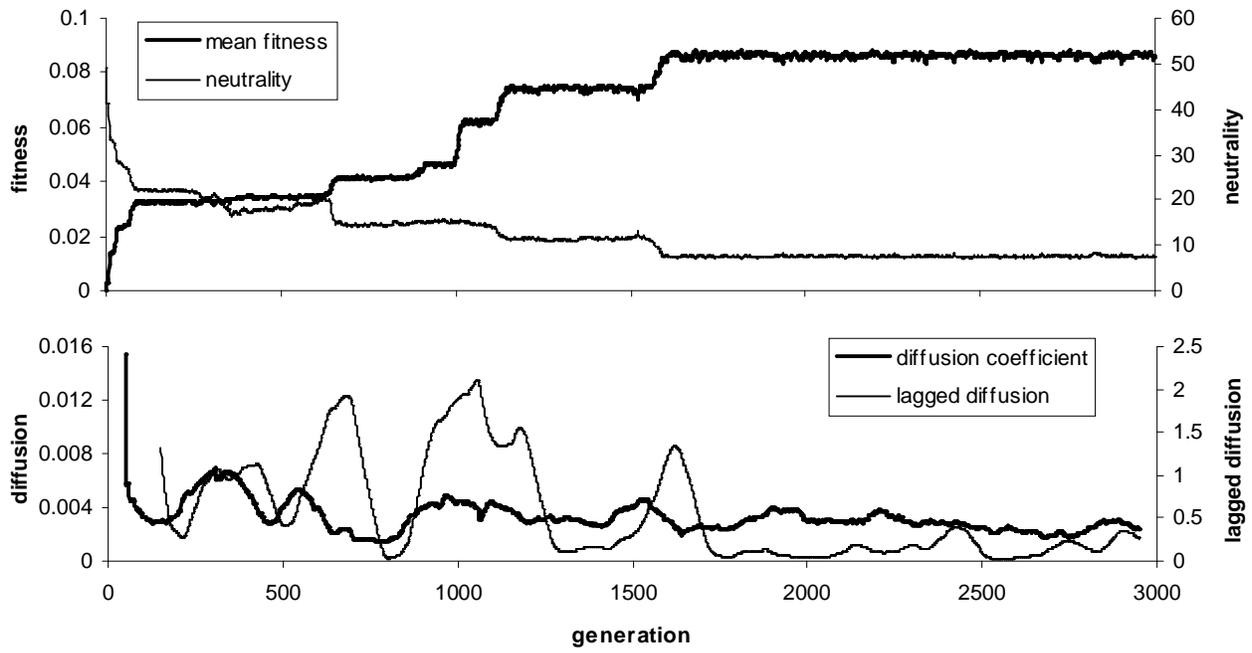


Fig 4

An evolutionary run on an NKp landscape ($N=60$, $K=12$, $p=0.99$) population size 200, mutation rate 0.001

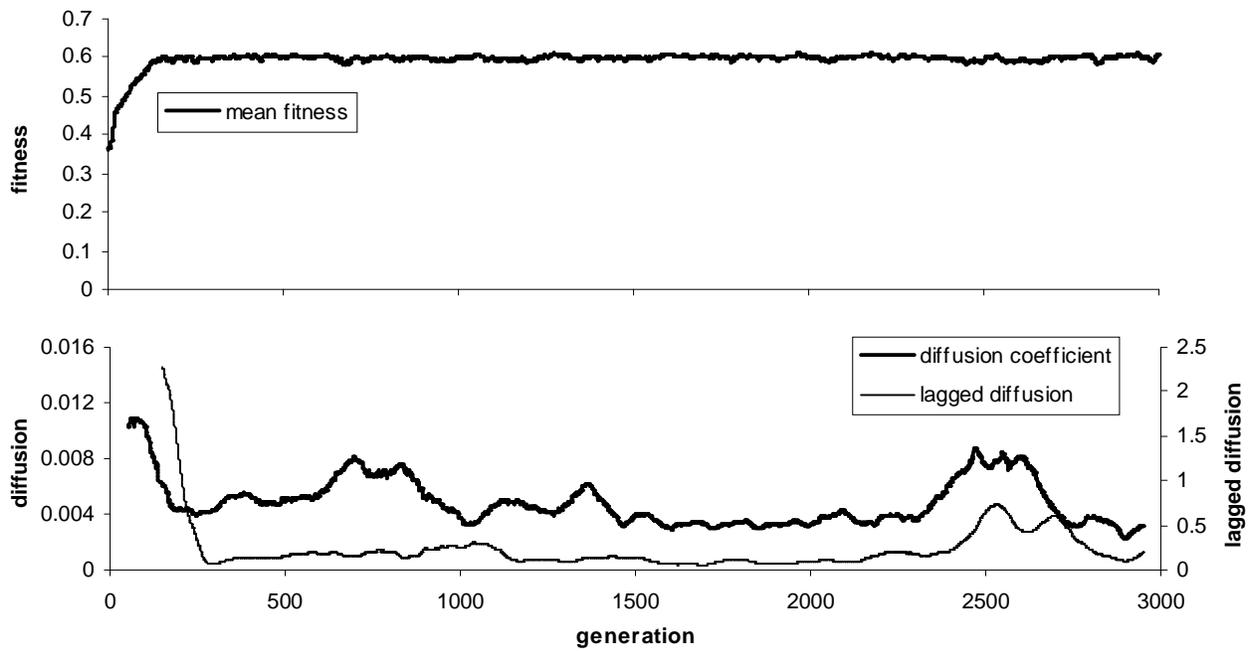


Fig 5

An evolutionary run on an NKp landscape ($N=60$, $K=12$, $p=0$) population size 200, mutation rate 0.001

Acknowledgements

The author would like to thank Inman Harvey and many others at the Sussex University CCNR for crucial discussions, and Dave Harper for invaluable assistance with statistical issues.

References

- Barnett, L. C. 1997. *Tangled Webs - Evolutionary Dynamics on Fitness Landscapes with Neutrality*. MSc. diss., School of Cognitive and Computing Sciences, Sussex Univ. UK.
- Baskaran, S., Stadler, P.F. and Schuster, P. 1996. *Approximate Scaling Properties of RNA Free Energy Landscapes*. J. Theor. Biol. 181: 299-310.
- Crow, J.F. and Kimura, M. 1970. *An Introduction to Population Genetics Theory*. Harper & Row, New York.
- Derrida, B. & Peliti, L. 1991. *Evolution in a Flat Landscape*. Bull. Math. Biol. 53:355-382.
- Eigen, M., McCaskill, J. & Schuster, P. 1989. *The Molecular Quasispecies*. Adv. Chem. Phys. 75:149-263.
- Forst, C.V., Reidys, C. & Weber, J. 1995. *Neutral Networks as Model-Landscapes for RNA Secondary-Structure Folding-Landscapes*. Lecture Notes in Artificial Intelligence, vol. 929: Advances in Artificial Life (Morán, F., Moreno, A., Merelo, J.J. & Chacón eds.), Springer-Verlag, Berlin.
- Goldberg, D.E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.
- Fontana, W., Stadler, P.F., Bornberg-Bauer, E.G., Griesmacher, T., Hofacker, I.L., Tacker, M., Tarazona, P., Weinberger, E.D. & Schuster, P. 1993. *RNA Folding and Combinatory Landscapes*. Phys. Review E, 47(3): 2083-2097.
- Grüner, W., Giegerich, R., Strothman, D., Reidys, C., Weber, J., Hofacker, I. L., Stadler, P. F. and Schuster, P. 1996. *Analysis of RNA Sequence Structure Maps by Exhaustive Enumeration: I. Neutral Networks. II. Structures of Neutral Networks and Shape Space Covering*. Monatshefte Chem. 127: 355-374 & 375-389.
- Harvey, I. and Thompson, A. 1996. *Through the Labyrinth Evolution Finds a Way: A Silicon Ridge*. Proc. 1st Internatl. Conf. Evol. Sys.: From Biology to Hardware (ICES 96). Springer-Verlag.
- Huynen, M.A 1996. *Exploring Phenotype Space through Neutral Evolution*. J. Mol. Evol. 43: 165-169.
- Huynen, M.A., Stadler, P.F. & Fontana, W. 1996. *Smoothness Within Ruggedness: The Role of Neutrality in Adaptation*. Proc. Natl. Acad. Sci. (USA) 93:397-401.
- Jakobi, N. 1996. *Encoding Scheme Issues for Open-Ended Artificial Evolution*. Proc. Parallel Processing in Nature, Voigt, H-M., Ebeling, W., Rechenberg, I. & Schwefel, H-P (eds.) Springer-Verlag: pp 52-61.
- Kauffman, S. A., 1993. *The Origins of Order - Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK.
- van Nimwegen, E., Crutchfield, J.P. & Mitchell, M. 1997. *Statistical Dynamics of the Royal Road Genetic Algorithm*. Santa Fe Institute Pre-print 97-04-035, Santa Fe, NM, USA.
- Nowak, M. & Schuster, P. 1989. *Error Thresholds of Replication in Finite Populations - Mutation Frequencies and the Onset of Muller's Ratchet*. J. Theor. Biol. 137:375-395.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T. & Flannery, B.P. 1992. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, Cambridge, UK.
- Reidys, C.M., Stadler, P.F. & Schuster, P. 1997. *Generic Properties of Combinatory Maps - Neutral Networks of RNA Secondary Structures*. Bull. Math. Biol. 59: 339-397.
- Schuster, P., Fontana, W., Stadler, P.F. & Hofacker, I.L. 1994. *From Sequences to Shapes and Back - A Case Study in RNA Secondary Structures*. Proc. Roy. Soc. (London)B 255:279-284.
- Stadler, P.F. 1996. *Landscapes and their Correlation Functions*. J. Math. Chem. 20: 1-45.
- Thompson, A. 1996. *Silicon Evolution*. Proc. Genetic Programming, Koza, J., R. et al. eds. 96:444-452 MIT Press.
- Weinberger, E.D. 1990. *Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference* Biol. Cybern. 63:325-336.