

University of Southampton Research Repository ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

An exploration of NK landscapes with neutrality

Nicholas Geard

s336070@student.uq.edu.au

Supervisor: A/Prof Janet Wiles

j.wiles@itee.uq.edu.au

School of Information Technology
and Electrical Engineering
University of Queensland
St Lucia Q 4072

October 2001

Table of Contents

Acknowledgements	iii
Abstract	iv
1. Introduction	1
2. Neutrality in Biology and Evolutionary Computation	3
2.1. Biological Origins	
2.2. Computational Applications	
2.3. Issues	
3. NK: A rugged fitness landscape	9
3.1. Introduction	
3.2. Description of the NK model	
3.3. Practical aspects of simulating NK landscapes	
3.4. Properties of $K = 0$	
3.5. Properties of $K = N-1$	
3.6. Properties of $0 < K < N-1$: Replicating Kauffman's results	
3.6.1. Optima height and walk length	
3.6.2. The <i>massif central</i>	
3.6.3. Local ruggedness	
3.6.4. Global ruggedness – a new view of landscapes	
3.7. Discussion	
3.7.1. Discussion of results	
3.7.2. Discussion of methodology	
4. NKp and NKq: The introduction of neutrality	21
4.1. Introduction	
4.2. Formal description of neutrality	
4.3. NKp model	
4.3.1. Fitness distribution	
4.3.2. Optima height and walk length	
4.3.3. The <i>massif central</i>	
4.3.4. Local ruggedness	
4.3.5. Global ruggedness	
4.3.6. Neutral degree	
4.4. NKq model	
4.4.1. Fitness distribution	
4.4.2. Optima height and walk length	
4.4.3. The <i>massif central</i>	
4.4.4. Local ruggedness	
4.4.5. Global ruggedness	
4.4.6. Neutral degree	
4.5. Discussion	
4.5.1. Differences between NK and NKp	
4.5.2. Differences between NK and NKq	

5. Evolution on neutral landscapes	39
5.1. Introduction	
5.2. Previous work on evolutionary dynamics	
5.3. Evolution by steps	
5.4. Comparative simulations	
5.4.1. Introduction	
5.4.2. Methodology	
5.4.3. Results	
5.5. Diversity maintenance between neutral layers	
5.5.1. Introduction	
5.5.2. Methodology	
5.5.3. Results	
5.6. Discussion	
6. General discussion and conclusions	53
6.1. Contributions	
6.2. Evaluation	
6.3. Conclusions	
6.4. Further Work	
Reference List	57
Appendix A – Calculating fitness values using a hash function	63
A.1. Introduction	
A.2. The first hash function – NKHashA.java	
A.3. The second hash function – NKHashB.java	
A.4. <i>Caveat emptor</i>	
Appendix B – Anomalous results on highly epistatic landscapes	67
B.1. Introduction	
B.2. Discussion	

Acknowledgements

First and foremost I would like to thank my supervisor, Janet Wiles, for her input and encouragement at each stage of this study. I would also like to thank the participants in the Neutral Network Reading Group for useful ideas and discussion. I would also like to extend my appreciation to the School of Information Technology and Electrical Engineering at the University of Queensland for providing the equipment employed in this study. Finally, I would like to express my gratitude to my parents for their support and encouragement over the past year.

Abstract

Recent research into the evolution of RNA molecules has raised awareness of the neutral theory of evolution (Kimura 1983). Several fitness landscape models based upon the NK model (Kauffman 1993) have been proposed to investigate the behaviour of populations evolving on neutral landscapes. In this study, Kauffman's (1993) results pertaining to the NK fitness landscape model are replicated and a new visualisation technique is developed to illustrate the presence of ruggedness and neutrality across a landscape. These results are extended to the neutral NKp (Barnett 1997) and NKq (Newman & Engelhardt 1997) variants of the NK model. It is shown that these two models of neutrality result in landscapes with significant structural differences. An important goal in evolutionary computation is a greater understanding of the relationship between landscape structure and the optimal choice of evolutionary algorithms. A series of simulations is run comparing the performance of a variety of individual and population-based algorithms on a range of neutral and non-neutral landscapes. It is found that hill-climbing algorithms are generally capable of finding better individual solutions, but population-based algorithms tend to find higher average solutions. Finally, it is shown that crossover allows a population on a neutral landscape to maintain a much higher level of genetic diversity on a neutral layer than a population using mutation alone. Furthermore, populations using crossover suffer a less dramatic loss of diversity when a new fitness level is discovered, and recover diversity more rapidly.

1. Introduction

In recent years, a new view of evolutionary search dynamics based upon neutral selection has arisen to challenge the traditional ‘hill-climbing’ paradigm. This new view has been prompted by research in molecular evolution that suggests a large proportion of mutations at a molecular level are selectively neutral. A revaluation is needed of the idea that populations evolve by converging upon local optima in the fitness landscape.

The new view proposes population clusters that wander across layers of uniform fitness under the influence of random mutation and genetic drift. Eventually an individual will chance upon a rare transition point to a more fit level, giving their offspring a distinct selective advantage that will eventually result in a rise in average population fitness. Much of the inspiration for the investigation of neutral landscapes came from research into the structure of RNA folding landscapes in molecular biology (Huynen, Stadler & Fontana 1996). The discovery of neutral features in the search spaces of several real world problems, such as evolving neural nets and hardware evolution, has further established the validity of this approach (Harvey & Thompson 1996).

Within the field of evolutionary computation, one major task has been the definition of the class of problems on which evolutionary search strategies outperform other techniques. This task is necessary in order to establish the value of applying evolutionary algorithms to search and optimisation problems. One approach has been to define general models of search space landscapes and observe and analyse the behaviour of a range of different search strategies upon them. Several of these models involving neutrality have been proposed based upon Kauffman’s well-known NK landscape model (Kauffman 1993, described in Section 3).

At present, there is still a fundamental lack of both empirical data and theoretical understanding concerning the structure of neutral landscapes. Many of the analytical results have been obtained from simplified or idealised landscapes. The extent to which these results apply to the more complex landscapes likely to be found in real world problems has not yet been demonstrated.

The goal of this study was to broaden knowledge of landscapes with neutrality by investigating structural properties of the NKp landscape model (Barnett 1997, 1998, described in Section 4.3) and the NKq landscape model (based on Newman & Engelhardt 1997, described in Section 4.4) using Kauffman’s original NK model for comparison. Initially, Kauffman’s simulations pertaining to NK landscapes were replicated. Next, these simulations were extended to the neutral NKp and NKq landscapes. A series of comparative simulations were carried out to investigate the performance of several different search algorithms on a variety of neutral and non-neutral landscapes. Finally, a simulation was developed to investigate the role of crossover in maintaining genetic diversity in a population evolving on a neutral landscape.

In order to provide some background for neutrality, a summary of the biological context in which the concept was developed and the field of evolutionary computation in which it has been applied is contained in Section 2. Kauffman's NK landscape model is introduced in Section 3 together with a replication of his major results, to establish a baseline for comparison purposes. The two neutral variants of the NK model are introduced in Section 4 with the results of a series of investigations into their structural properties.

At this point, with a more thorough understanding of neutral landscapes, it is appropriate to shift focus from the structure of the landscapes to the performance of evolutionary search algorithms on them. A further review of literature related to evolutionary search processes and dynamics is presented in Section 5, followed by the results of the comparative simulations and investigation into diversity maintenance. Finally, the significance of these results and implications for future work are discussed in Section 6.

2. Neutrality in Biology and Evolutionary Computation

In order to appreciate the significance of neutrality, it is necessary to have some awareness of the biological context in which the concept was developed and also of the field of evolutionary computation in which it has found application. Those with a knowledge of evolutionary biology and the neutral theory of evolution may choose to skip Section 2.1. Similarly, those with a knowledge of evolutionary computation may wish to skip Section 2.2.

2.1. Biological Origins

This section briefly describes the origin of the traditional ‘hill-climbing’ metaphor for evolution and introduces the motivation and research that led to development of the neutral theory of evolution. Some of the implications of this new view for various topics in evolutionary biology are also discussed.

The traditional view held that natural selection was the driving force behind the continual innovation seen in evolution. It proposed that each individual organism could be attributed a fitness value representing its probability of surviving and reproducing in a given environment. The primary source of differences between organisms was in the genetic material they carried and differences at this genotypic level were presumed to be reflected in the phenotype and hence in the selective fitness of the individual. The Fundamental Theorem of Natural Selection (Fisher 1932) states that, so long as there is a variance in fitness across a population and individuals are selected to reproduce on the basis of their relative fitness, the average population fitness will increase over time. The genetic diversity necessary for continuous evolution is ensured by the occurrence of random mutations.

One useful way of visualising the space of possible genotypes in which evolution occurs is to use fitness landscapes (Wright 1931). A fitness landscape is created by specifying a relationship between the set of all possible genotypes, for example, an adjacency matrix based upon an operator such as single-point mutation, and a mapping between the each genotype and its fitness value via a fitness function. The process of evolution can then be described in terms of the movement of populations over the landscape surface (see Lipsitch 1991; Manderick, de Weger & Spiessens 1991; Peliti 1995 for several methods of characterising this relationship). This landscape analogy has been extended to many areas of application, including describing search spaces in combinatorial optimisation problems, the energy function of disordered systems such as spin glasses, and RNA molecule folding space. It was in the context of thinking about fitness landscapes that the idea of evolution as a hill-climbing search process arose.

More recently however, investigations into molecular evolution have uncovered a higher than expected degree of variability at a molecular level. To account for this, Kimura developed the Theory of Neutral Evolution (Kimura 1970, 1983). While the traditional view held that the majority of mutations have some selective effect, Kimura proposed that most single point mutations in biological molecules do not affect the fitness of the molecule to perform its function. One of the first areas to make significant use of Kimura’s idea of selective neutrality was molecular biology. Based upon Maynard Smith’s (1970) early intuitions of protein space, the model of a

molecular *quasispecies* was developed, in which a population on a fitness landscape is pictured as a ‘cloud’ of individuals clustered around a dominant genotype (Eigen, McCaskill & Schuster 1989). Within this population there may exist a wide variety of different genotypes, all of which are expressed in the same phenotype.

Eigen et. al.’s (1989) quasispecies model was supported by further research into RNA evolution (Huynen 1996; Huynen, Stadler & Fontana 1996; Forst, Reidys & Weber 1995). Several approaches have been used to investigate and characterise RNA protein folding landscapes, including simulation, statistical sampling and exhaustive enumeration of sequence structures. Furthermore, a large body of theoretical work relating to landscapes has been developed based upon random graph theory (Reidys, Stadler & Schuster 1997; Reidys, Forst & Schuster 1999; Reidys & Schuster 1999). This research has shown that there is a highly degenerate mapping between sequences (genotypes) and structures (phenotypes) – many protein sequences fold into the same molecular structure. This phenomena leads to a significant level of neutrality in RNA folding landscapes, resulting in several interesting features:

- Sequences folding into identical structures are distributed randomly across the search space; and
- Sequences folding into *all* common structures can be found in a relatively small region of search space.

These features result in a landscape structure that is considerably different to that proposed by the traditional view, in which every genotype maps to a unique phenotype with a distinct fitness value. If the level of neutrality is high enough, there is a high probability that two neighbouring genotypes will map to the same phenotype. As the fitness of an individual is generally based upon its phenotype, this mutation can occur in the absence of selective pressure to increase fitness. Furthermore, if the proportion of neutral mutations is high enough, a series of genotypes linked by these neutral mutations may result, allowing for a variety of neutral features in a landscape. These features have been variously described in the literature as *neutral layers* or *neutral networks* and may take the form of plateaus, terraces, or interconnected ridges stretching through the search space.

One consequence of these neutral layers is that large portions of the fitness landscape can be explored by accumulating neutral mutations and drifting along neutral networks before a potentially beneficial mutation is discovered and exploited by selection. Ohta (1973) claimed that such selective neutrality does not even require that the fitness of molecules remain precisely unchanged. If mutations result in a small enough change in fitness, their effect may never be felt on an evolutionary time scale. Evolution on such a landscape is characterised by long periods (epochs) of stasis during which the population diffuses over the current neutral network, interspersed with occasional fitness increases. At these times, one or more individuals in the population encounter a portal to a higher fitness network and selection favours these higher fitness individuals and their offspring while the remainder of the population dies off. Another consequence of neutrality is that, on at least some classes of landscape involving neutrality, given enough time *any* sequence can be reached from a given starting point and the problem of getting trapped in local optima effectively vanishes.

The concept of neutrality has been used to assist speculation on a variety of issues in evolutionary biology. Several of these issues – speciation and punctuated equilibria – are described below.

Speciation and the evolution of reproductive isolation is one area in which neutrality has been applied (Gavrilets & Gravner 1997; Gavrilets 1997). The traditional view was that different species occupied distinct peaks in a fitness landscape. They were reproductively isolated because any hybrid created by interbreeding would lie in the fitness valley between the two peaks and hence be strongly selected against. Gavrilets recognised several problems with this picture, including the implausibility of the idea that each distinct species has a different fitness, and the improbability of speciation ever occurring if it required a sub-population to migrate across a valley of low fitness. He suggests that, rather than occupying different peaks on a fitness landscape, species are actually subpopulations on a neutral network separated by an accumulation of neutral mutations. Speciation occurs when a population diffuses into two subpopulations with distinct, but equally fit, genotypes, which are unable to interbreed, perhaps because the neutral mutations result in sexually incompatible phenotypes.

A second area in which concepts of neutrality have been used is to support the theory of punctuated equilibria (Eldredge & Gould 1972). On neutral landscapes, evolution follows a distinctive pattern – populations spend long periods of time wandering across neutral networks before finding rare transitions to higher neutral networks. This dynamic has been termed *epochal evolution* by van Nimwegen and Crutchfield (1997). This pattern bears a striking similarity to the palaeobiological phenomena of punctuated equilibria and it has been proposed that the theory of neutral evolution may provide a plausible explanation (see, for example, Fontana & Schuster 1998).

2.2. Computational Applications

Having described the biological context in which the idea of neutrality was developed, it is important to see how this concept was transferred to the field of evolutionary computation.

Evolutionary algorithms were developed throughout the 60's and 70's as a new approach to computer search and optimisation based upon evolution and natural selection (Holland 1975; Goldberg 1989; Fogel 1995). While initial applications centred primarily on function optimisation and related issues, they have also been used as a simulation tool for modelling a wide variety of features of evolutionary biology and population genetics. Although these techniques had been in use before 1975, it was the publication of Holland's "*Adaptation in Natural and Artificial Systems*" with its definition of the Schema Theorem which signalled the birth of a new field. The Schema Theorem, also known as the building block hypothesis, states that evolutionary algorithms work by assembling short, highly fit blocks into longer, more fit blocks using recombination (Holland 1975). Random mutation was generally used as well to maintain a level of genetic diversity in the population and prevent it from converging completely.

As evolutionary algorithms became more widely known and accepted in the computational community, researchers began to ask questions about the class of problems for which they were best suited. The challenge was to find a set of test problems, which were “hard” for traditional search and optimisation techniques but “easy” for genetic algorithms. Research in this direction separated along several paths. One of the early attempts to define what made a problem hard was based upon deceptiveness, the idea that the local structure of a landscape may lead a hill-climbing algorithm away from the global optimum whereas a population based algorithm using genetic operators may still succeed (Goldberg 1993).

The Royal Road problems are another set of landscapes on which evolutionary techniques were expected to outperform traditional hill climbing techniques (Mitchell, Forrest & Holland 1992; Forrest & Mitchell 1993; Mitchell & Holland 1993). These problems exploited the fundamental concept of the building block hypothesis by assigning a solution a fitness value according to the number of ‘blocks’ of bits that had been assembled correctly. The fitness landscapes resulting from these problems were highly neutral (although they were not labelled as such at the time). Any mutation that did not result in the formation (or the destruction) of a block did not result in a fitness change. This lack of feedback was thought to be an impediment to a hill climber because a large number of mutations were required to be present simultaneously before a fitness increase was realised.

In many instances however, evolutionary algorithms performed worse on the Royal Road problems than appropriately modified hill-climbing techniques (Mitchell & Holland 1993). One of the problems for evolutionary algorithms was that crossover, which was assumed to play a primary role in combining blocks, also had the disruptive effect of dismantling blocks before they had become fixed in the population. Another problem was that many of these test suites were unimodal (i.e. lacking local optima). Hence they were simply too easy for hill climbing variants to solve – generally far more efficiently than a population based genetic algorithm. More recently a new model of hierarchically consistent test problems that exploit the capabilities of the crossover operator has been proposed (HIFF and related models, Watson, Hornby & Pollack 1998; Watson & Pollack 1999). Suitable evolutionary algorithms consistently outperform hill-climbing algorithms on these highly multimodal landscapes, primarily because they require that a diverse population to be maintained throughout the evolutionary process in order to prevent convergence on a local optima.

The class of landscape models with which this study is primarily concerned is the NK landscape model (Kauffman 1993, based upon earlier work in Kauffman and Levin 1987). This model was developed in order to illustrate the fitness landscapes that result from chromosomes with varying degrees of epistatic. The fitness contribution of each gene in a genotype is determined by itself and a number of other genes to which it is linked. If there are no epistatic interactions, each gene may be independently optimised and the resulting landscape is smooth. At the other extreme, if the fitness contribution of each gene depends on the value of every other gene, then changing a single gene affects the fitness contributions of every other gene and the landscape is essentially random. In this case it may not be possible to optimise the fitness contribution of every single gene (similar to the property of ‘frustration’ in spin glass systems).

Depending on the number of interactions, landscapes generated using the NK model vary from a smooth landscape with a single peak through to a rugged, hilly terrain with numerous local optima. Kauffman obtained values for the number, density and height of optima, their accessibility from a random starting location and the degree of correlation of landscapes that resulted from this model (Kauffman 1993). Of particular interest was the use of autocorrelation functions developed by Weinberger (1989) to measure the similarity of nearby genotypes in a fitness landscape. The correlation of a landscape measures, given a single fitness value, how much can be inferred about the surrounding fitness values. This level of *ruggedness* was thought to be a primary indicator of the level of difficulty of searching the landscape. In the simplest case (a completely correlated landscape) all slopes lead up to a single optima, in the most complex case (an uncorrelated landscape), the fitness at every point is completely unrelated and no information is provided on the location of the global optima.

The basic NK model was further developed into neutral models by Newman and Engelhardt (1997) and independently by Barnett (1997, 1998). These new models raised questions about the number and size of neutral networks, the nature of the transition points between them and their degree of percolation (the extent to which they span the search space). There has also been a great deal of interest in investigating the behaviour of populations on neutral landscapes, in particular relating to the length of time needed for a population to discover a higher fitness network and the possibility of a population escaping from a local optima via neutral mutations. Doubts have been cast on the validity of using autocorrelation functions to measure the difficulty of searching landscapes containing a high degree of neutrality. On one hand, the presence of neutrality could be thought to increase correlation, as there are a large proportion of genotypes whose neighbours have equal fitness. On the other hand, the fact that any network is arbitrarily close to a given genotype means that large, effectively uncorrelated, fitness jumps are possible. One of the main attractions of neutral networks from a search and optimisation point of view is that they may offer a potential ‘escape route’ from many local optima. In the most extreme case, it is theoretically possible (but practically infeasible) to transform *any* multimodal landscape into a single peaked landscape with the addition of neutrality (Harvey & Thompson 1996).

Research into evolving neural networks for robot control and on-chip hardware evolution (Harvey & Thompson 1996; Thompson 1996) have both suggested that neutrality might also be inherent in complex real world problem domains. It has also been suggested that introducing encoding redundancy, and hence a degree of neutrality, may improve evolutionary search performance in a number of other problem areas (Shipman, Shackleton & Harvey 2000). Not only is the probability of becoming trapped in local optima reduced but also the solutions discovered are frequently more robust.

2.3. Issues

One of the major issues in evolutionary computation at the moment is how to best select an evolutionary algorithm for a particular problem and how to optimise that algorithm’s parameters. As is known from machine learning, no one algorithm will be optimal for all tasks (the “no free lunch” theorem, Wolpert & Macready 1995) and

so the focus of investigation has shifted from finding the “best” algorithm to finding the best algorithm for a particular search space (see, for example, Sharpe 2000).

Barnett (2000) has argued that in many cases, an iterative adaptive walk by a ‘net crawler’ will prove most efficient in landscapes containing neutrality and possessing the additional property of ε -correlation, which applies to landscapes in which networks with higher fitness values are always accessible from the current network (i.e. there are no ‘dead end’ networks from which no higher fitness levels can be reached). Studies of RNA folding landscapes suggest that this may occur in nature. The Royal Road problems also fall into this category.

More understanding is required of problem domains which do not satisfy this property including how this can be determined, whether population based search improves performance, and what part (if any) recombination has to play in evolutionary search. The neutral variants of the NK landscape, Barnett’s (1997) NKp and Newman and Engelhardt’s (1997) NKq models, are sufficiently complex that they would not, in most cases, be expected to satisfy this ε -correlation property.

A further issue is that the structures of the landscapes generated by these models are not yet fully understood. While the two neutral models have been proposed to account for the same biological phenomena, intuition suggests that they may result in landscapes that may have significantly different properties.

The issues of primary interest investigated in this study were:

- The qualitative similarities and differences between neutral and non-neutral landscapes;
- The differences resulting from using the NKp and NKq implementations of neutrality;
- The comparative performance of population based algorithms and simple hill-climbing algorithms; and
- The effects of the various genetic operators such as mutation and crossover on algorithm performance.
- The possibility of crossover allowing greater population diversity to be maintained across transitions between neutral layers.

3. NK: A rugged fitness landscape

As stated in the preceding section, Kauffman's NK landscape model is a useful general model both for investigating the structural properties of landscapes and evaluating the performance of evolutionary algorithms. As the NK model forms the basis on which the neutral variants to be considered in Section 4 are built, it is important to understand the properties of the original model.

3.1. Introduction

Kauffman developed the NK model in *The Origins of Order* (1993) in order to illustrate the connections between selection and self-organisation in evolutionary systems. The NK model is based upon the idea of genetic linkage. Each gene in a genotype makes a contribution towards the fitness contribution of the genotype as a whole, however, the fitness contribution of each gene may be dependent not only on its own allele, but also on the alleles of any number of other genes, either in positions adjacent to it, further away on the genotype, or on another genotype altogether. As the number of these epistatic connections between genes increases, the resulting fitness landscape is changed from being relatively smooth and predictable to increasingly rugged and random.

In this part of the study, the NK landscape model was investigated using simulations. Various structural properties were examined both at the limiting cases and across the range of landscapes that fall in between. These properties include height of optima, lengths of adaptive walks and several techniques for visualising the ruggedness of landscapes. Kauffman's original results were replicated and a new visualisation technique was developed to illustrate structural aspects of the landscapes that appear at a global level.

3.2. Description of the NK model*

For sequences of length N , with Q possible alleles at each position, a *fitness landscape* is defined to be a mapping $f : Q^N \rightarrow R$ termed the *fitness function*. In this paper, only genotypes with 2 possible alleles are considered, so Q^N can be considered to be an N dimensional hypercube. The *fitness* of a particular genotype $g \in G$ is then given by $f(g)$. Two genotypes are considered *neighbours* if they differ by only one position.

An NK landscape has two primary parameters:

N - the length of the genotype, where $N > 0$; and

K - the number of epistatic linkages to each gene, where $0 \leq K < N$.

Each gene i , where $1 \leq i \leq N$, contributes to the total fitness of the genotype depending on the value of its allele and on those of each of the K other genes to which it is linked. It is assigned a *fitness table* mapping each of the 2^{K+1} possible combinations of alleles to a random, real valued number in the range $[0, 1]$ (See Figure 3.1). The fitness contribution of gene i of genotype g is given by $f_i(g)$.

* The mathematical formalism used in this section is based on that used in Barnett (1997).

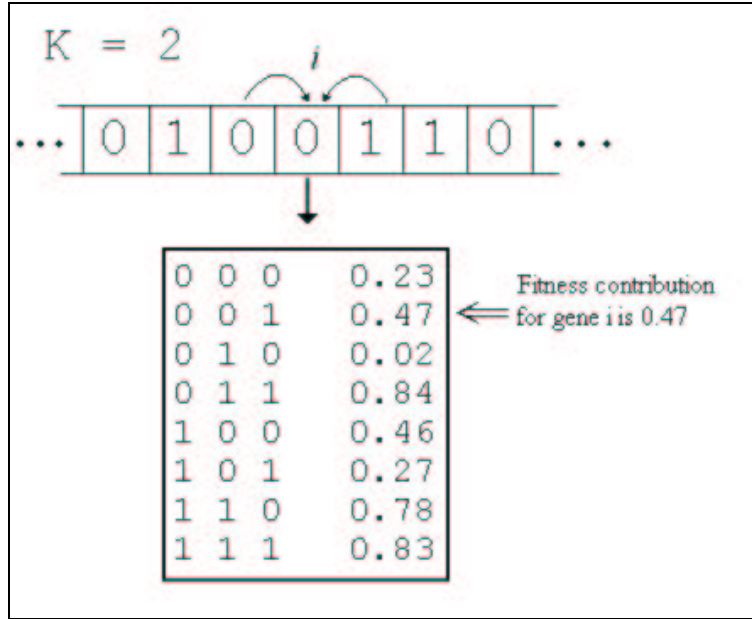


Figure 3.1: A portion of a genotype and fitness table. In this example, $K = 2$, therefore each gene i is linked to two other genes (in this case, the gene to either side). Each gene has its own fitness table with one entry for each possible combination of alleles. (Adapted from Kauffman 1993, Figure 2.2, page 42).

The total fitness of the genotype is the sum of the fitness contributions of each of its genes divided by the total number of genes and thus also falls within the range $[0, 1]$:

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

There are two different ways of assigning the K epistatic links. Using *neighbourhood interaction*, the K links are assigned to the $K/2$ genes immediately to the left and the $K/2$ genes immediately to the right of the gene under consideration. If K is odd, the additional link is (arbitrarily) determined to be to the right of the gene. Using *random interaction*, the K links are assigned at random to K distinct genes from the remaining $N-1$ genes in the genotype. It is also important to note that these links are unidirectional, that is, if the fitness contribution of the gene at position i depends upon the value of the gene at position j , it does not necessarily follow that the fitness contribution of the gene at position j will depend upon the value of the gene at position i .

On any landscape, an *adaptive walk* is considered to be a sequence of genotypes $g_0, g_1, g_2, \dots, g_n$ connected by single bit mutations such that the fitness increases with each successive step, i.e. $f(g_n) > f(g_{n-1})$.

3.3. Practical aspects of simulating NK landscapes

In the simulations carried out for this paper, the NK landscapes were implemented using two different methods. For relatively small landscapes, the fitness tables for each gene were generated using a single random seed at initialisation and used to lookup fitness contributions for the remainder of an evolutionary simulation.

For larger landscapes, particularly as the number of epistatic interactions increased, the storage space required for these tables became prohibitive and a hash function was used to generate consistent random fitness contributions (see Altenberg 1997 for a formal definition). In this method, a single random base was generated for each gene. When a fitness contribution was required, the binary bit string representing the combination of alleles was converted to an integer and the product of this number and the base was used to seed a pseudorandom number generator. The second double from this pseudorandom number generator was taken as the fitness contribution^{*}.

As a wide variety of different simulation techniques were used, the methodology for each of these is described at the appropriate points in Sections 3.6 and 3.7. Initially, the limiting cases of $K = 0$ and $K = N - 1$ will be examined.

3.4. Properties of $K = 0$

In the simplest case where there are no epistatic interactions, there is a single global optimum that can be reached by an adaptive walk from any other point in the fitness landscape and the landscape is relatively smooth – neighbouring genotypes have similar fitness values (see Figure 3.2).

As there are no epistatic interactions, the fitness contribution of each gene depends only upon its own allele and therefore has only two possible values. One of these will be higher than the other, and the global optimum is realised when each gene is set to its higher fitness allele. All other genotypes will have lower fitness and can be changed to the optimum genotype by a sequence of single bit mutations that each result in a fitness increase. It follows therefore that there are no other optima.




Figure 3.2: A representation of a smooth fitness landscape ($N = 8$, $K = 0$). There is a single optima, which can be climbed to from any point in the search space and neighbouring genotypes generally have similar fitness values. Genotypes have been ordered to show the monotonic structure of the landscape.

^{*} Note that this is slightly different from the method described in Altenberg (1997). The second double was taken as the first proved to be insufficiently ‘random’. See Appendix 1 for full details of the algorithm used. The random number generator used was from the `java.util.Random` class in *Java™ 2 Platform Std. Ed. v1.3.1*.

Furthermore as each gene's contribution to the fitness of a genotype is scaled by $1/N$, the maximum fitness difference between two neighbouring genotypes will be $1/N$ (in the limiting case where a gene's fitness contribution flips from 0 to 1 or vice versa).

For a randomly chosen genotype, on average, half of its genes will be set to the optimal allele and the other half will be set to the non-optimal allele. Therefore, on average, it will take $N/2$ steps for an adaptive walk to reach the global optimum.

3.5. Properties of $K = N-1$

Having considered the limiting case where there are no epistatic interactions, the other extreme, where $K = N-1$ is described. In this case, each gene is epistatically linked to every other gene and the fitness contribution of each gene depends on the setting of itself and every other gene. This results in a landscape that is entirely random – knowing the fitness of a genotype provides no information about the fitness of any of its neighbours (see Figure 3.3).

At this limit, landscapes have the following properties:

- There are a large number of local optima.
- The average length of an adaptive walks to local optima is short and scales as $\ln N$.
- Any given genotype can climb to only a small fraction of the local optima.
- Conversely, only a small fraction of genotypes can climb to any given optimum (including the global optima).

A significant feature of these landscapes is that finding the global optimum falls into the class of NP-complete problems. In fact, it has been shown that when epistatic interactions are random, finding the global optima of an NK landscape is NP-complete when $K \geq 3$, and $K = 2$ is an open question (Weinberger 1996).

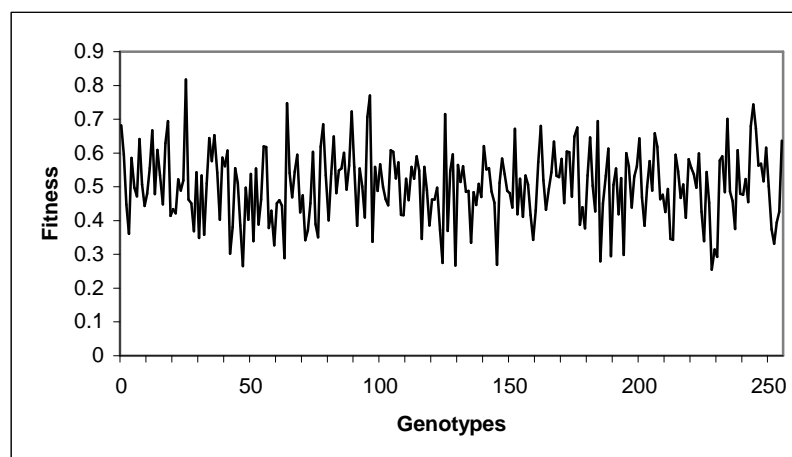


Figure 3.3: A representation of a rugged fitness landscape ($N = 8$, $K = 7$) – there are a large number of optima, each of which is reachable from only a small portion of the search space and neighbouring genotypes tend to have very different fitness values. Genotypes have been ordered numerically.

3.6. Properties of $0 < K < N-1$: Replicating of Kauffman's results

One of the strengths of the NK landscape model is that the parameters can be set to generate landscapes with arbitrary levels of ruggedness between the two extremes described in Sections 3.4 and 3.5. In order to investigate the properties of these landscapes, a series of simulations were run replicating those described in (Kauffman 1993).

3.6.1. Optima height and walk length.

One of the phenomena observed by Kauffman was 'complexity catastrophes', a decrease in the performance of a system (the fitness) as complexity (the number of epistatic interactions) increases. This relationship between performance and complexity can be illustrated by measuring the average height of local optima found by adaptive walks on a range of landscapes. It follows intuitively that, as a landscape becomes more crowded with peaks, the average length of a walk from a random starting point to any one of them will decrease.

For each pair of N and K values, 100 different landscapes were randomly generated. On each of these, 1,000 adaptive walks from random starting points were carried out. In this and the following sections, an adaptive walk consisted of 500 steps. In each step, a one-bit mutant of the current genotype was generated and evaluated. If its fitness was equal to or greater than that of the current genotype, it became the new current genotype. After 500 steps, up to 10 exhaustive searches were carried out on all one-bit mutations of the final genotype, again accepting those with a greater or equal fitness. The aim of this was to ensure that the final genotype was likely to be a local optima. The final fitness found and the average walk length were averaged and plotted (see Figures 3.4 and 3.5). While Kauffman presented his results in tabular form, it was found that plotting graphs illustrated the trends more clearly and allowed for easy comparison with later results on neutral landscapes.




Figure 3.4: The mean fitness of local optima for a range of landscapes. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed (replicated following the methodology in Kauffman 1993, page 55).




Figure 3.5: The mean walk length to local optima for a range of landscapes. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed (replicated following the methodology in Kauffman 1993, page 55).

As K increases, the number of conflicting constraints increases, for example if $K = 1$, the fitness contribution of the gene at position i may depend on the value of the gene at position j , and similarly the fitness contribution of the gene at position j may depend on the value of the gene at position i . It is entirely possible that the combination of alleles that realises the optimal fitness contribution from the gene at position i , results in a suboptimal fitness contribution from the gene at position j . This is similar to the phenomena termed ‘frustration’ in spin-glass systems in physics. As the value of K increases further, the number of these suboptimal compromises increases and hence the average fitness falls (see Figure 3.4). As anticipated, the length of an adaptive walk from a random point to the nearest local optima generally decreased as K increased (see Figure 3.5).

3.6.2. The *massif central*

For low values of K , the local optima are clustered in the same region of search space, termed a *massif central* by Kauffman (1993). As K increases, the local optima diffuse evenly over the fitness landscape. Kauffman suggested that this level of global structure in the landscape, the highest optima being nearest to one another and to the global optimum, might have particular implications for the usefulness of recombination as a genetic operator. Once several good solutions appear in a population, the existence of a massif central structure would increase the probability of better solutions being discovered by using crossover to interpolate between the current peaks.




Figure 3.6: The correlation between the fitness of local optima and their Hamming distance from the fittest optima found on 3 NK landscapes ($N = 96$, (a) $K = 2$; (b) $K = 4$, (c) $K = 8$) (replicated using the methodology described in Kauffman 1993, page 60).

Three landscapes with $N = 96$ and three different values of K (2, 4 and 8) were generated. On each landscape 1,000 adaptive walks (each of 500 steps) were carried out and the final genotypes and fitnesses recorded. Each final fitness reached was plotted against the genotype's Hamming distance* from the global optima. The global optimum in this case was defined to be the fittest optima discovered after the 10,000 walks (see Figure 3.6).

On the smoothest landscape ($K = 2$), the highest local optima are found closest to the global optima (see Figure 3.6(a)). This is less apparent when the landscape is slightly more rugged (see Figure 3.6(b)) and the effect disappears completely when $K = 8$ (see Figure 3.6(c)). At this point, the highest local optima are as likely to be found at a great distance from the global optima as the lower local optima.

3.6.3. Local ruggedness

The maximum possible fitness difference between genotypes is given by $\frac{K+1}{N}$.

Therefore, as K increases with respect to N , the maximum possible fitness difference increases, resulting in the landscape becoming more rugged. Kauffman illustrated this by superimposing a graph of the fitness of a single local optima over a graph of all of its one-bit mutation neighbours. This provides a visual representation of the correlation between neighbouring points on a fitness landscape.

As K increases the fitness difference between a local optima and its neighbouring points also increases (see Figure 3.7). This results in optima being more difficult to locate, as there is less information in the local neighbourhood to guide search towards them.

* The Hamming distance between two genotypes is given by the number of bit positions at which they differ.

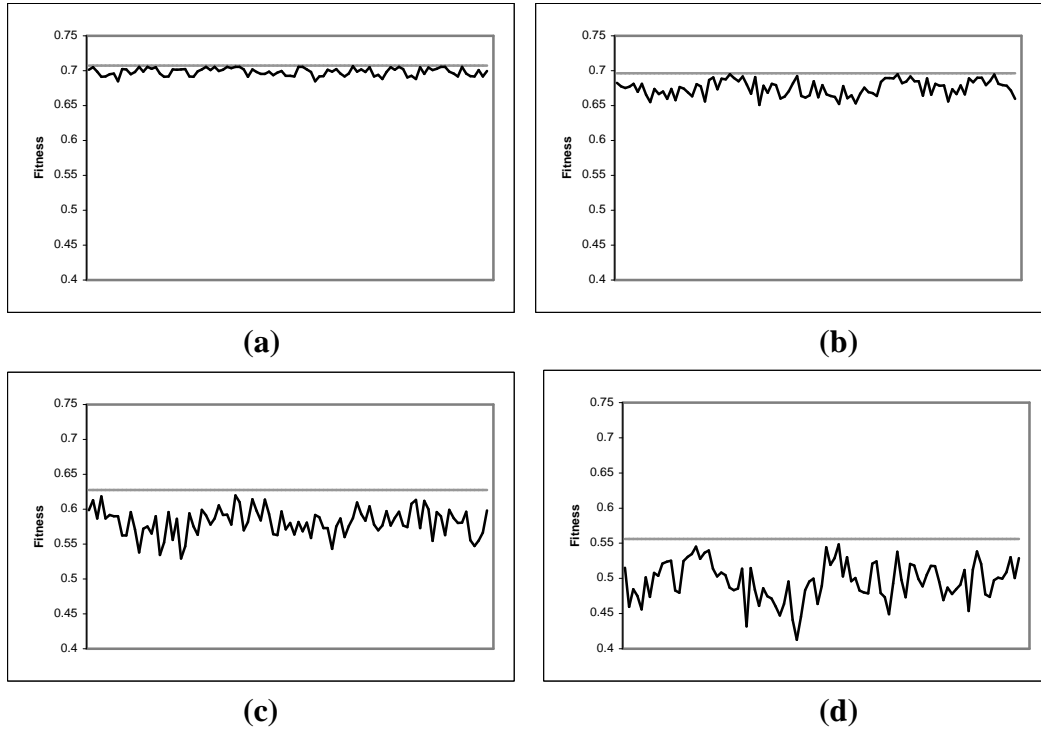


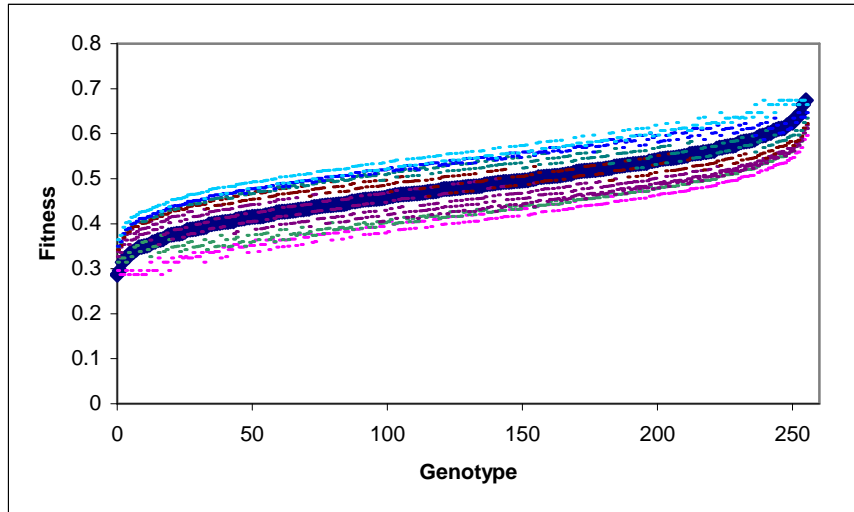
Figure 3.7: The ruggedness of NK landscapes where $N = 96$ for increasing values of K , (a) $K = 2$; (b) $K = 4$; (c) $K = 8$; (d) $K = 95$. The graphs show the fitness of a random local optima (horizontal line) and the fitness of all 96 of its one-bit mutation neighbours (replicated following the methodology of Kauffman 1993, page 57)*.

3.6.4. Global ruggedness – a new view of landscapes


Although the visualisation technique presented in Section 3.6.3 gives a good idea of what happens to an NK landscape as K increases at a local level, it is more difficult to visualise what happens to the structure of a landscape at a global level as K increases.

A novel approach to visualising landscapes was developed by taking a sample of random points on the landscape (or an exhaustive enumeration for sufficiently small landscapes) and plotting a fitness distribution overlaid with the fitness of each of genotype's one-bit mutation neighbours (see Figure 3.8). Comparing graphs of landscapes with a variety of different K values illustrates quite clearly the gradual disintegration of structure in the landscape as K increases. In the case of a smooth landscape there is a high degree of correlation between the fitness of a point and the fitness of each of its one-bit mutation neighbours (see Figure 3.8(a)). At any given point in the landscape, the adjacent points are only slightly higher or slightly lower – a trend that persists across the landscape, from valleys to peaks (with the obvious exception that valley points have no lower neighbours and peaks have no higher neighbours). In the most rugged landscape there is no correlation between neighbouring points at all (see Figure 3.8(c)). At any point in the landscape, the adjacent points are as likely to be significantly higher or lower than the current point as they are to be similar.


* Note that in these graphs, each point represents the fitness of a single neighbour of the optima under consideration. There is otherwise no relationship between adjacent points. They are connected simply to provide a visual approximation of the ruggedness of the landscape.



(a)



(b)



(c)

Figure 3.8: Plot of fitness values of each genotype in a set of NK landscape (bold line) overlaid with the fitness of each genotype's one-bit mutation neighbours. $N = 8$, (a) $K = 0$; (b); $K = 2$; (c) $K = 7$.

3.7 Discussion

The results of the simulations matched Kauffman's original results to an acceptable level. Due to the significant random component of the model, identical quantitative results were not expected, however, the observed qualitative trends were similar. The novel visualisation technique described in Section 3.7 was found to provide a useful graphical representation of the ruggedness of landscapes at a global level. The remainder of this section will focus on several unexpected results and methodological issues

3.7.1. Discussion of results

An interesting feature of the results (noted by Kauffman 1993) appears in the investigation of local optima height (see Section 3.6.1 and Figure 3.4). The greatest average local optima height occurs when there is a certain level of ruggedness in the landscape (generally at $K = 2$). This was initially surprising as it was expected that the highest optima would occur in the smooth landscape ($K = 0$) and that conflicting constraints would lower the average optima height towards the landscape average as the number of interactions increased. Furthermore, the average optima height does not fall below that of the smooth landscape ($K = 0$) until the landscape becomes significantly more rugged (around $K = 16$).

Looking more closely at the way in which fitnesses are calculated however, the reason for this becomes clearer. The fitness contribution of each gene is chosen uniformly and at random from the range $[0, 1]$. For a genotype of length 100, the smooth ($K=0$) case, requires 200 random numbers to be generated (one for each of the 0 and 1 alleles at each locus). While each of these genes can be optimised independently, there is a good chance that at a significant proportion of the loci, the choice will be between two low fitness alleles.

As the landscape becomes more rugged, the pool of possible fitness contributions grows larger. When $K = 2$, 400 possible fitness contributions are generated and at each loci there is a choice between four different values. The result of this is that, on average, more alleles will have the possibility of realising a high fitness contribution. As K continues to increase, the size of the pool of possible fitness contributions increases exponentially (number of random numbers generated = $2^{K+1} N$).

The down side of increasing the number of interactions is that, as described in Section 3.5, the optimal choice may not always be able to be made – optimising one gene may result in the others with which it interacts being set to suboptimal values. As K increases, so does the number of compromises that need to be made.

At a certain point however, the benefit of having a greater range of fitness contributions to choose from can be balanced against the need to compromise on these choices, and the resulting landscape will allow a population achieve a higher average fitness. This feature is important for comparison between the original NK landscape model and the neutral NKp and NKq variations (introduced in Section 4), which result in landscapes that are optimal at different values of K .

3.8.2. Discussion of methodology

One serious methodological issue that arose in the simulations was the possibility of obtaining different results depending upon the method used to generate fitness contributions. Due to the limitations of computational speed and memory, full enumeration of the fitness tables was not feasible for landscapes with large values of N and K . The hash function method (Altenberg 1997, described in Section 3.3) was used in these cases, but some anomalies were noted in the results it produced.

Initially, the Altenberg's method was implemented by generating a random number ('base random') for each location on the genotype. The bit string resulting from the gene under consideration and its K linkages was converted to an integer and the sum of this and the base random was used to seed a random number generator from which the fitness contribution was obtained.

This method resulted in the numbers being used to seed the random number generator being very close together (for each gene, they fell within the range [base random, base random + 2^N]). It was discovered that the random numbers generated with these seeds were extremely similar (they tended to be identical up to the third or fourth significant digit). This problem was minimised by using the product of the integer representation and the base random (as opposed to the sum) and taking the second random number in the sequence (as opposed to the first) as the fitness contribution (see Appendix 1 for the algorithm used).

There still exists the possibility of some unwanted correlation between fitness contributions. One obvious solution is to use a more robust random number generator that does not result in any correlation between random number sequences produced by similar seeds. A full investigation of random number generation was outside the scope of this study.

4. NKp and NKq: The addition of neutrality

Having described the original NK landscape model in Section 3, it is now possible to introduce the two neutral variations, NKp and NKq.

4.1. Introduction

One of the problems encountered in research into both natural and artificial evolution is the possibility of a population converging on local optima and becoming trapped, unable to evolve any further. In biological evolution this doesn't seem to happen. Even if certain suboptimal traits become fixed in a population, the species is not necessarily prevented from undergoing further evolution. For those interested in evolutionary computation as a search or optimisation technique this is a problem because it means that the evolutionary search algorithms used will often become trapped at suboptimal solutions.

One possible solution to these problems is the theory of selective neutrality. Recent research into molecular evolution has suggested that most single point mutations are selectively neutral, that is, they have no effect on the fitness of the molecule. This feature, coupled with the high-dimensionality of most search spaces means that there will frequently be a neutral 'escape route' from local optima (Conrad, Ebeling & Volkenstein 1993).

Two approaches that have been suggested to introduce neutrality into NK fitness landscapes are:

- The NKp landscape proposed by Barnett (1997, 1998), which introduces neutrality by setting a certain proportion (p) of the entries in a genotype's fitness tables to 0;
- The NKq^{*} landscape proposed by Newman and Engelhardt (1997), in which the entries in a genotype's fitness table are quantised to an integer value (in the range $[0, q)$).

This part of the study examined both of these models of neutral landscapes, comparing them to the original NK model and to each other. The focus was on differences in their structural properties and considering possible implications for evolutionary search. The general approach was to extend the methodology used by Kauffman (replicated in Section 3) to each of the neutral landscape models to provide a common form of data for comparison. In general the methodology used for each simulation was identical to that described in the corresponding parts of Section 3.6. Any significant differences are described in the report at the appropriate point.

^{*} In Newman & Engelhardt's (1997) study, the parameter determining the number of levels of quantisation of the fitness contributions was termed F . That parameter has been renamed q in this paper and the resulting landscape labelled NKq ("quantised" NK landscape) purely for reasons of aesthetic symmetry with the NKp landscape model. Newman & Engelhardt's model is otherwise unchanged.

4.2. Formal description of neutrality

This section defines several terms that will be used throughout the remainder of this report.

A *neutral mutation* is a one-bit mutation from g to g' where $f(g) = f(g')$.

A *neutral network* is a set of genotypes such that each genotype is linked to one or more others by neutral mutations. This is also referred to as a *neutral layer* in some of the literature.

The *neutral degree* of a genotype is the number of different neutral mutations it may undergo.

On any neutral landscape, a *neutral walk* is a sequence of genotypes $g_0, g_1, g_2, \dots, g_n$ connected by single bit mutations such that there is no change in fitness with each successive step, i.e. $f(g_n) = f(g_{n-1})$. Generally a neutral walk is of interest if it spans the search space, as this suggests that the neutral network percolates through the search space to some degree.

4.3. NKp model

The NKp model proposed by Barnett (1997, 1998) simulates neutrality by setting a proportion p of the entries in a gene's fitness table to 0. Therefore there are a number of different allele combinations that will result in that gene making no fitness contribution to the genotype as a whole. When $p = 0$, there are no zero-entries in the fitness tables and the landscape is equivalent to one produced using the NK model. As $p \rightarrow 1$ the number of zero-entries in the fitness tables increases and the landscape produced becomes more neutral. At the limit of $p = 1$, all the fitness entries are set to zero and a flat landscape results.

4.3.1. Fitness Distribution

The fitness of a genotype on an NKp landscape is calculated in the same way as one on an NK landscape, by summing the fitness contributions of each of the genes and dividing by the length of the genotype. Due to the increasing number of genes making zero fitness contributions as $p \rightarrow 1$, the average fitness will tend to fall as the amount of neutrality increases. As the actual number of zero-entries is probabilistic, there is no straightforward way of scaling the resulting fitness back to the range $[0, 1]$, which makes fitness comparisons between landscapes with different levels of neutrality difficult.

To demonstrate the resulting difference, the fitness values of a random sample of 10,000 points on a variety of smooth and rugged, neutral and non-neutral landscapes were plotted (see Figure 4.1). The most noticeable difference is the general drop in average fitness when neutrality is introduced.

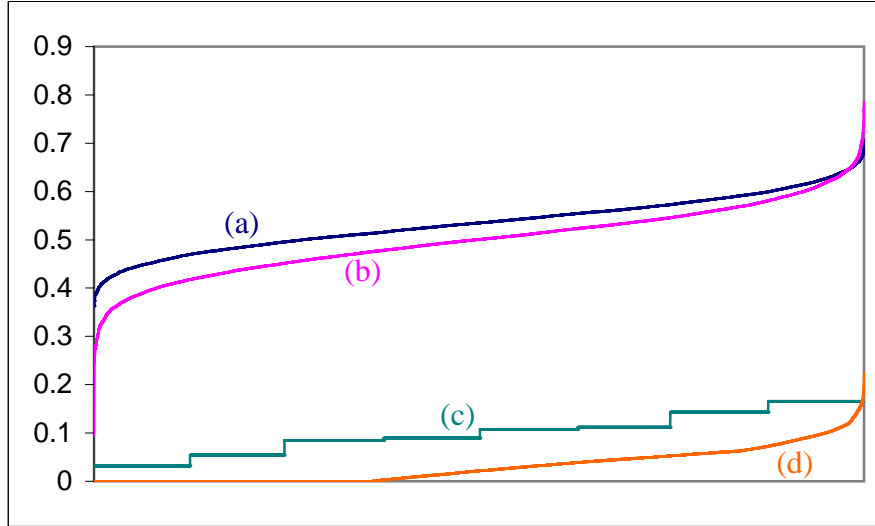


Figure 4.1: Fitness distributions for a variety of smooth and rugged NKp landscapes: (a) $N = 16, K = 0, p = 0.0$; (b) $N = 16, K = 15, p = 0.0$; (c) $N = 16, K = 0, p = 0.94$; (d) $N = 16, K = 15, p = 0.94$. Fitness distributions are based upon a random sample of 10,000 points.

Another feature of this landscape model is that the structure of the landscape is radically altered, particularly when the landscape is very rugged. The high proportion of fitness entries set to zero in a very neutral landscape (for example, where $p = \frac{N-1}{N}$), results in up to approximately one third of the possible genotypes having a total fitness of zero (in the $K = N - 1$ limiting case – see fitness distribution (d) in Figure 4.1).

This suggests an interaction between the K and p parameters of an NKp landscape that warrants further analytical investigation.

4.3.2. Optima height and walk length

Results for average optima height and walk length were obtained using the same method as in Section 3.6.1 above.

While the general trend towards decreasing optima height and decreasing walk length was similar to that found for NK landscapes (compare Figures 4.2 and 4.3 with Figures 3.4 and 3.5 above), the value of K at which the average optima height is maximised varies. For NK landscapes, an epistatic linkage of 2 generally resulted in the maximum average optima height, whereas for NKp this value varied from 4 to 8 depending on the length of the genotype (see Figure 4.2).




Figure 4.2: The mean fitness of local optima for a range of NKp landscapes where $p = 0.99$. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed.




Figure 4.3: The mean walk length to local optima for a range of NKq landscapes where $q = 2$. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed.

Similarly, the addition of neutrality results in significantly shorter walk lengths on smooth landscapes (see Figure 4.3). It must be noted that in many cases, the “local optima” which has been discovered may well not be a local optima at all, but rather the edge of a neutral plateau from which further improvements could be made at some more distant point.

4.3.3. The *massif central*

As in Section 3.6.2 the presence of a *massif central* was tested by plotting the height of 1,000 local optima against their Hamming distance from the global optima.

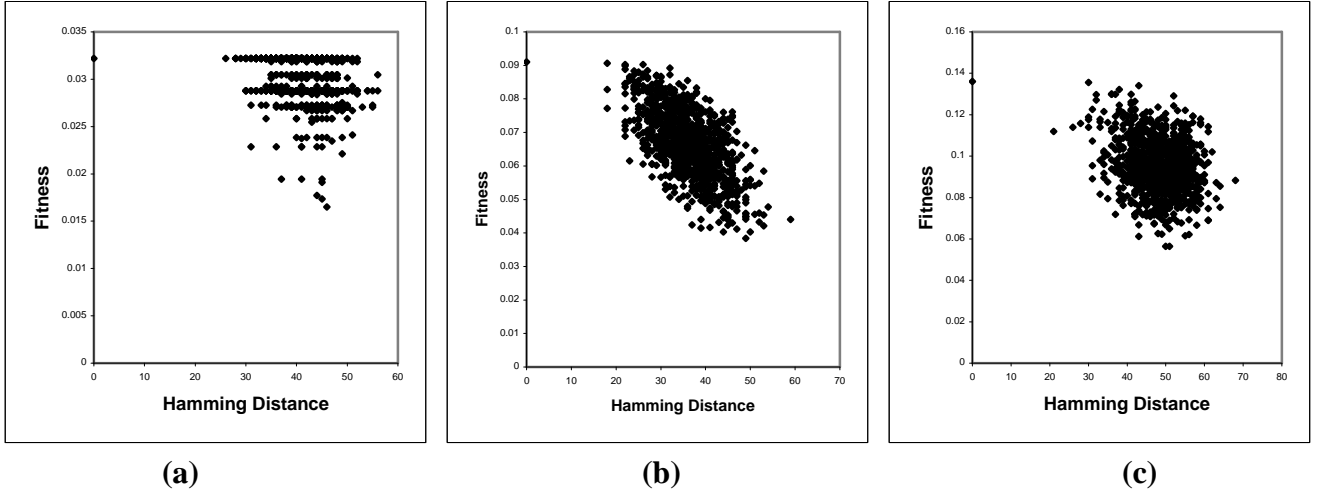


Figure 4.4: The correlation between the fitness of local optima and their Hamming distance from the global optima on 3 NKp landscapes ($N = 96$, $p = 0.99$, (a) $K = 2$; (b) $K = 4$; (c) $K = 8$.

It appears that the *massif central* phenomena of local optima being clustered together and close to the global optima in landscapes with low K also appears in NKp landscapes. One feature that emerged from this simulation (see Figure 4.4 (a)) is the large number of points with equal fitness to the global optima. This suggests that the global optima in this case is not necessarily an isolated peak, but rather one or more plateaus or ridges extending across the landscape.

4.3.4. Local ruggedness

As in Section 3.6.3, the ruggedness of NKp landscapes was investigated by comparing the fitness of a random local optima with the fitnesses of each of its one-bit mutation neighbours (see Figure 4.5).

In contrast to the non-neutral landscapes (see Figure 3.7 above), most single bit mutations from a local optima result in no fitness change. Especially for low values of K , local optima exist as plateaus or ridges rather than as isolated peaks. At the same time, in a sufficiently rugged landscape, many fitness levels may be found in a fairly small neighbourhood of a randomly chosen point (see Figure 4.5 (d)). This illustrates one of the features of neutral, that any structure is accessible within a short distance of an arbitrary point in the search space.

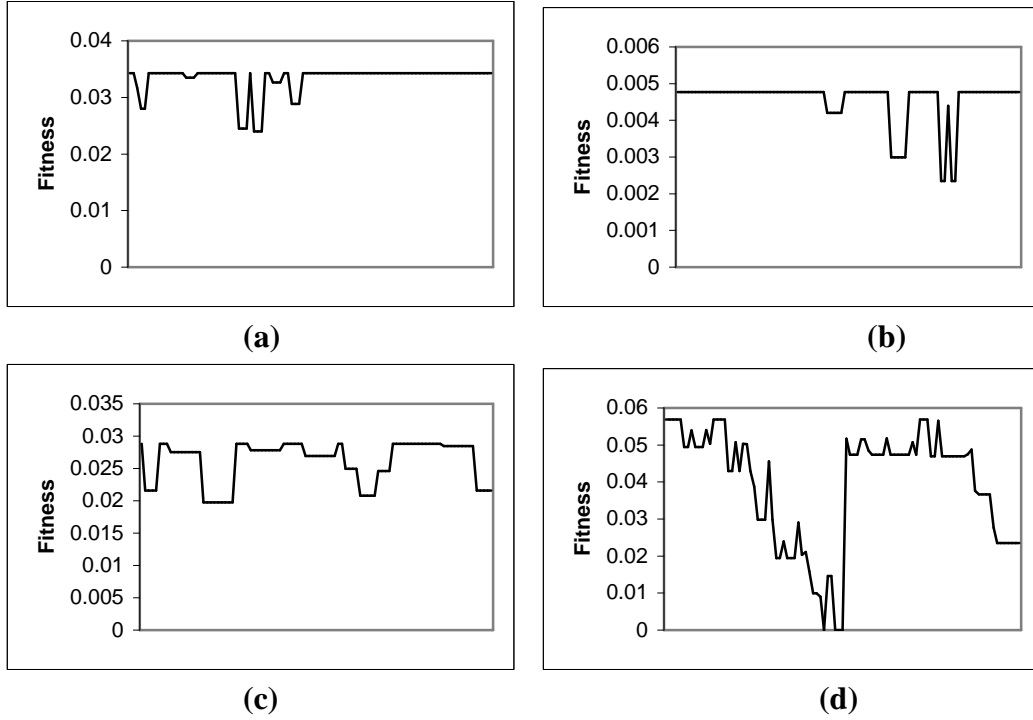


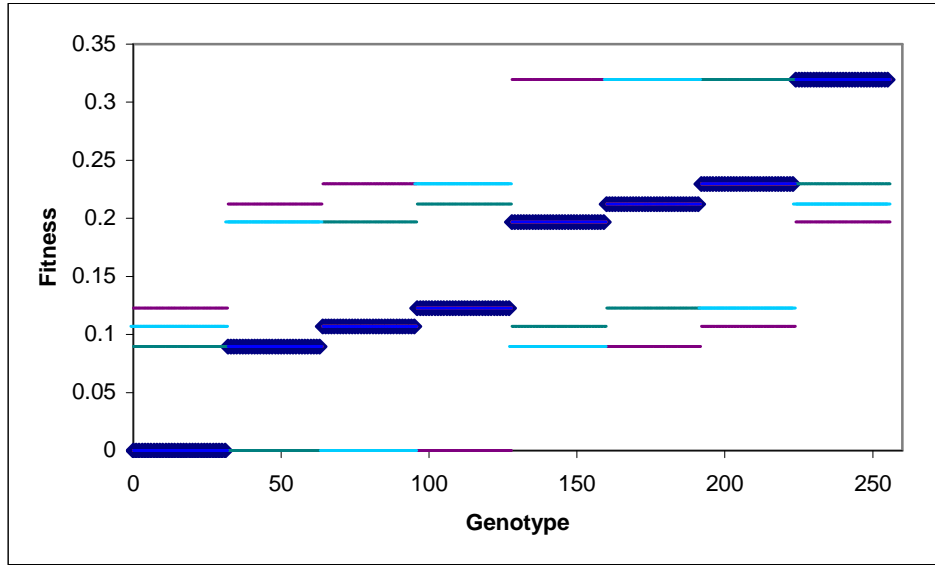
Figure 4.5: The ruggedness of NKp landscapes where $N = 96$ and $p = 0.99$ for increasing values of K , (a) $K = 2$; (b) $K = 4$; (c) $K = 8$; (d) $K = 95$. The graphs show the fitness of a random local optima (horizontal line) and the fitness of all 96 of its one-bit mutation neighbours.

4.3.5. Global ruggedness

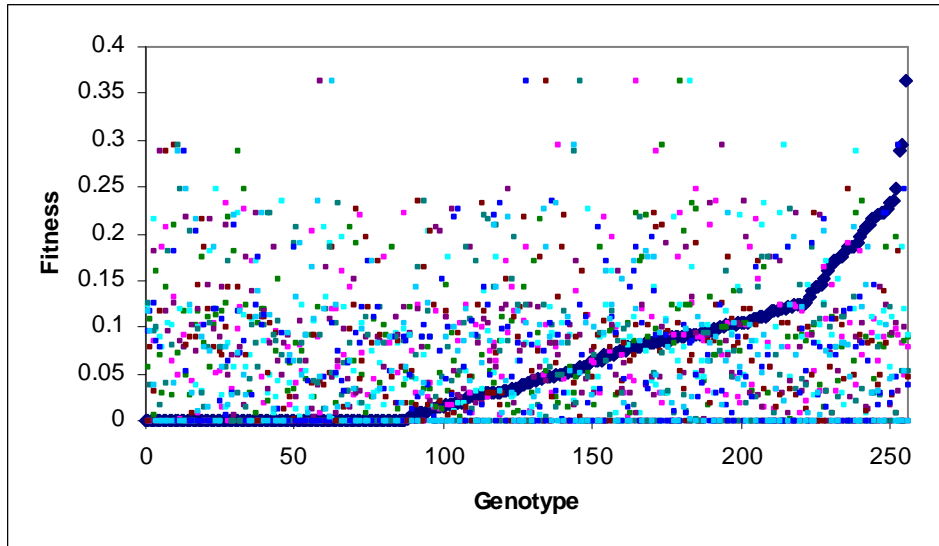
The same novel technique used to visualise neighbourhood relationships in NK landscapes may also be used for neutral NK landscapes.

This allows the neutral layers (and the areas of the landscape from which they can be accessed) to be clearly seen as horizontal lines across the graph (see Figure 4.6). Caution must be used in interpreting these graphs however, as the existence of several genotypes with the same fitness does not necessarily imply that they are connected by a sequence of neutral mutations. The equal fitness genotypes may be isolated points or be grouped in several small clusters rather than forming one large connected network. Using the measurement of neutral degree described in Section 4.3.6 below in conjunction with these graphs allows for a more accurate understanding of the landscapes.

One interesting feature of these graphs is that they allow significant changes in the structure of landscapes to be recognised visually. For instance the graphs in Figure 4.7 suggest that some sort of ‘phase change’ may occur when neutrality is changed from $\frac{N-3}{N}$ to $\frac{N-4}{N}$. It is hoped that results obtained from evolutionary simulations upon these landscapes will support the idea of critical values of parameters, around which the structure of landscapes (and potentially the dynamics of populations on them) changes radically.

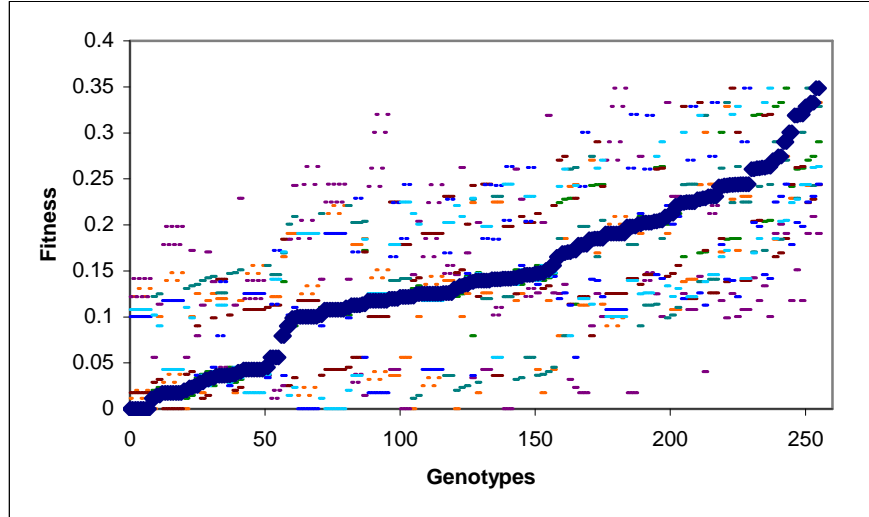


(a)




(b)

Figure 4.6: Plot of fitness values of each genotype in (a) a smooth ($N = 8$, $K = 0$, $p = 0.875$) and (b) a rugged ($N = 8$, $K = 7$, $p = 0.875$) NKp landscape (bold line) overlaid with the fitness of each genotype's one-bit mutation neighbours.



(a)



(b)

Figure 4.7: Critical levels of neutrality. Plot of fitness values of each genotype in an NKp landscape (bold line) overlaid with the fitness of each genotype's one-bit mutation neighbours. $N = 8$, $K = 2$, (a) $p = 0.625 ((N-3)/N)$; (b) $p = 0.5 ((N-4)/N)$.


4.3.6. Neutral degree

One further measurement of the structure of neutral landscapes that may be obtained is the neutral degree. By comparing the number of mutations that are beneficial, neutral or detrimental at each fitness level, it is possible to see where on a landscape the areas of maximum neutrality lie.


Relatively small landscapes ($N = 8$) were used so that each genotype in the search space could be enumerated and each of its possible mutations checked. Two landscapes were investigated, one relatively smooth ($K = 2$) and another maximally rugged ($K = 7$). The genotypes were binned into categories according to their fitness. The first category represents an average of the 0 fitness genotypes. Each successive category represents an average of the genotypes falling into the next 0.05 fitness range (see Figure 4.8).

As would be expected in both cases, the number of beneficial mutations is maximal at the lowest fitness level, falling to zero at the highest fitness level. Conversely, the number of detrimental mutations rises from zero at the lowest fitness level to the maximum possible value of 8 (for an 8 bit string) at the highest fitness level.

In the less rugged landscape, the degree of neutrality decreases in proportion with the fitness level, indicating that the size of neutral layers decreases as the landscape is climbed. In the rugged landscape however, the only neutral layer is found at zero. All of the non-zero fitness genotypes are found on rugged peaks rising out of this flat floor. As indicated in Section 4.3.5 above, some of these genotypes may have equal fitness, but because they are not connected by a sequence of neutral mutations, they do not form a neutral layer.



(a)



(b)

Figure 4.8: Average number of detrimental, neutral and beneficial mutations at each fitness level. (a) a smooth neutral landscape (N=8, K=2, p=0.875); (b) a rugged neutral landscape (N=8, K=7, p=0.875). Each fitness level represents a 0.05 fitness range.

4.4. NKq model

The NKq model introduced by Newman and Engelhardt (1998) takes the different approach of quantising the fitness contributions of each gene. Rather than being random real numbers in the range $[0, 1]$, they are random integers in the range $[0, q]$ where q defines the number of levels of quantisation of each fitness contribution. The maximally neutral case occurs when $q = 2$ and there are only two possible fitness contributions a gene can make (0 or 1). As $q \rightarrow \infty$ neutrality decreases and the landscapes approach those produced using the NK model.

4.4.1. Fitness distribution

The fitness of a genotype is calculated by a modified version of the original fitness equation for NK landscapes:

$$f(g) = \frac{1}{N(q-1)} \sum_{i=1}^N f_i(g)$$

This results in the fitness being scaled back to the range $[0, 1]$, meaning that direct comparisons between landscapes with different levels of neutrality can be made.

To illustrate, the fitness values of a random sample of 10,000 points on a variety of smooth and rugged, neutral and non-neutral landscapes were plotted.

In contrast to the NKp model of neutrality introduced in Section 4.3, introducing neutrality using the NKq model results in a much less dramatic change to the structure of the landscape (see Figure 4.9). Essentially, by quantising the range of possible fitness levels, the NKq model smoothes out many of the local irregularities in a standard NK landscape into flat terraces while leaving the global structure of the landscape relatively unaltered.




Figure 4.9 comparison of fitness distributions for a variety of smooth and rugged NKp landscapes with and without neutrality: (a) $N = 16, K = 0, p = 0.0$; (b) $N = 16, K = 15, p = 0.0$; (c) $N = 16, K = 0, p = 0.94$; (d) $N = 16, K = 15, p = 0.94$. Fitness distributions are based upon a random sample of 10,000 points.

4.4.2. Optima height and walk length

Results for average optima height and walk length were obtained using the same method as in Section 3.6.1 above (see Figures 4.10 and 4.11).

The general trend relating increasing K to decreasing average optima heights and decreasing walk lengths remains similar to that found for NK landscapes (see Figures 4.10 and 4.11). One of the main differences is that the average heights of local optima are, on the whole, greater on an NK q landscape than on an NK landscape. (This feature can also be seen from the graph of fitness distributions in Figure 4.9).




Figure 4.10: The mean fitness of local optima for a range of NK q landscapes where $q = 2$. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed.




Figure 4.11: The mean walk length to local optima for a range of NK q landscapes where $q = 2$. Each data point is averaged over 100 randomly generated landscapes, on each of which 1,000 adaptive walks were performed.

Furthermore, adaptive walks are generally somewhat shorter on an NKq landscape than an NK landscape. This follows directly from the properties of neutrality. Local optima, rather than consisting of a single point are likely to be a cluster of points extending in some sort of ridge or plateau through the landscape. Therefore they will be close to a greater number of points in the landscape and the average length of a walk to them from a random starting point will be shorter. This may also explain the slight increase in average walk length from $K = 0$ to $K = 2$.

4.4.3. The *massif central*

As in Section 3.6.2 the presence of a ‘*massif central*’ in NKq landscapes was tested by plotting the height of 10,000 local optima against their Hamming Distance from the global optima.

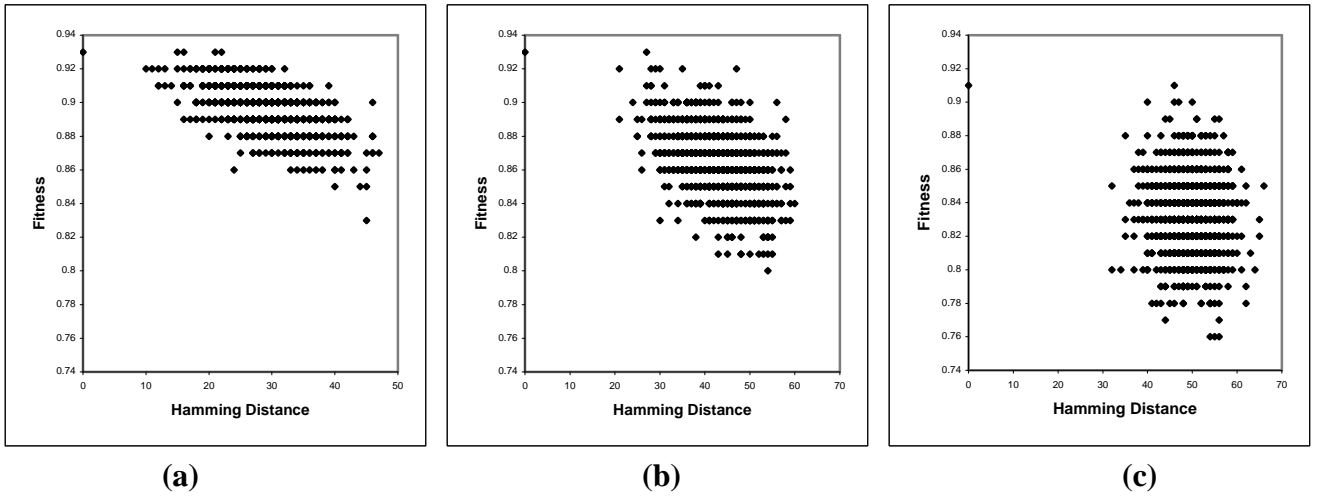


Figure 4.12: The correlation between the fitness of local optima and their Hamming Distance from the global optima on 3 NKq landscapes ($N = 96$, $q = 2$, (a) $K = 2$; (b) $K = 4$; (c) $K = 8$).

As was expected, considering the minimal impact neutrality has on global structure in the NKq model, the *massif central* also appears in NKq landscapes (see Figure 4.12). The quantised nature of the landscape is also quite visible.

4.4.4. Local ruggedness

As in Section 3.6.2, the ruggedness of NKq landscapes was investigated by comparing the fitness of a random local optima with the fitnesses of each of its one-bit mutation neighbours.

As K increases, the fitness differences between local optima and neighbouring genotypes also increase (see Figure 4.13), though possibly more slowly than in the NK landscape model (see Section 3.6.3).


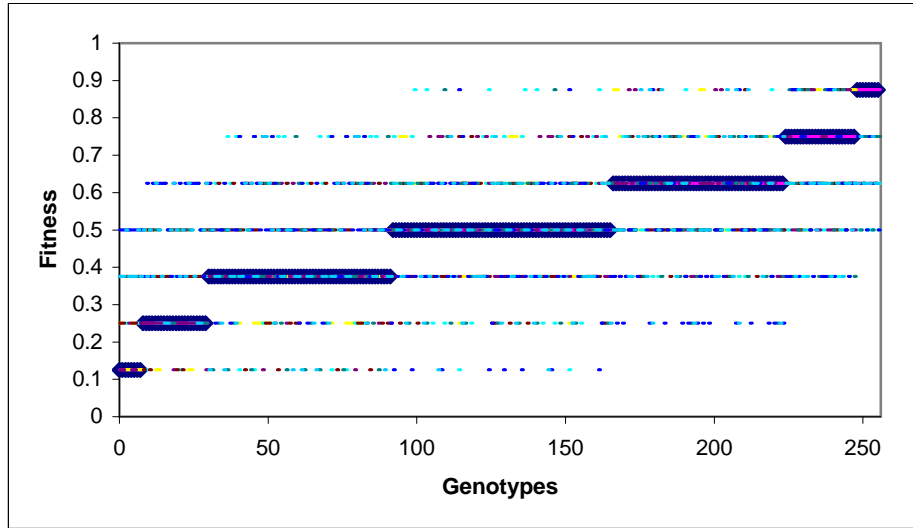


Figure 4.13: The ruggedness of NKq landscapes where $N = 96$ and $q = 2$ for increasing values of K , (a) $K = 2$; (b) $K = 4$; (c) $K = 8$; (d) $K = 95$. The graphs show the fitness of a random local optima (horizontal line) and the fitness of all 96 of its one-bit mutation neighbours.


While some neutrality is visible in the graphs in Figure XXX (as evidenced by the points at which a neighbouring point has the same fitness as the optima, or where two neighbouring genotypes have the same fitness), there is not as much as may be expected. One possible reason for this is that, on NKq landscapes, the maximum amount of neutrality tends to be found around the landscape average (see Section 4.4.7. below) rather than around the peaks.

4.4.5. Global ruggedness

The new technique developed to visualise ruggedness on a global scale was applied to NKq landscapes. Again, the neutral levels can be seen clearly (see Figure 4.14). The difference in the way neutrality is distributed across a landscape is also apparent. Whereas neutrality on the rugged NKp landscape was found primarily at the lower fitness levels (see Figure 4.6 above), it is more evenly distributed around the centre of the rugged NKq landscape, falling off as fitness increases or decreases (see Figure 4.14 (b)).



(a)




(b)


Figure 1: Plot of fitness values of each genotype in (a) a smooth ($N = 8$, $K = 0$, $q = 2$) and (b) a rugged ($N = 8$, $K = 7$, $q = 2$) NKq landscape (bold line) overlaid with the fitness of each genotype's one-bit mutation neighbour.

4.4.6. Neutral degree

As in section 4.3.7, the neutral degree of the landscape was investigated by exhaustively enumerating a relatively small landscape and recording the number of detrimental, neutral and beneficial mutations at each level (see Figure 4.15). This was somewhat simpler using the NKq landscape model, as the fitness of each genotype is already quantised into a known number of levels depending on the genotype length and amount of neutrality (the number of possible levels is given by $qN - N + 1$. see Newman & Engelhardt 1997 for working). Note however, that not all of these possible fitness levels will necessarily appear in a particular landscape.



(a)



(b)

Figure 4.15: Average number of detrimental, neutral and beneficial mutations at each fitness level. (a) a smooth neutral landscape ($N = 8$, $K = 2$ and $q = 2$); (b) a rugged neutral landscape ($N = 8$, $K = 7$ and $q = 2$)

In the smooth landscape, the neutral degree remains similar across the landscape (see Figure 4.15 (a)). In the rugged landscape however, the average number of neutral mutations is maximal at the middle fitness level (see Figure 4.15 (b)). This indicates that the largest neutral layer occurs at the average fitness level of the landscape, and the size of the neutral layers decreases as you move up or down the landscape.

4.5 Discussion

The most interesting discovery from this part of the research is that not only does the structure of landscapes generated using the NKp and NKq models differ significantly from that of those generated using the standard NK model, they are also quite different to each other. It may be concluded that while adding neutrality to a landscape does introduce certain general properties, the way in which that neutrality is implemented is also critical to the structure of landscape generated. This section highlights some of the major differences between the various landscape models

4.5.1. Differences between NK and NKp

The first major difference between the NK and NKp landscape models is that the maximum average optima height occurs at a much higher level of epistatic interaction on NKp landscapes than on NK landscapes (compare Figures 3.4 and 4.2). At the same time, average walk lengths are significantly shorter (compare Figures 3.5 and 4.3). For low levels of epistatic interaction (in particular the $K = 0$ smooth case) walk lengths are extremely short, suggesting that the optima in these landscapes are very large and easily accessible from many different regions in the search space. This intuition is supported by investigations into the neighbourhood of local optima (see Figure 4.5), suggesting that local optima exist in the form of large plateaus or ridges extending throughout the search space.

Another critical feature of the NKp landscape model is the fact that there is a relationship between the ruggedness and neutrality, such that the landscape structure varies significantly as the values of K and p are altered (as can be seen in the fitness distributions in Figure 4.1). At low levels of epistatic interaction, the similarity of the ‘massif central’ results suggests that global landscape structure is similar. However, as K increases, the number of zero-fitness genotypes increases, changing the shape of the landscape dramatically by creating a large flat zero-fitness floor. It is expected that the magnitude of this effect on the fitness distribution could be decreased by altering the fitness value to which the proportion p of the fitness table entries are set (say from 0 to 0.5).

4.5.2. Differences between NK and NKq

Using the NKq model to introduce neutrality has a far less radical effect on the landscape than using the NKp model. The evidence suggests that by quantising the fitness contributions, the global landscape structure remains generally unchanged. At a local level, hillsides move from being smooth slopes to level terraces and optima change from isolated peaks to flat plateaus and ridges.

In comparison with the NK landscape, the average NKq landscape fitness is similar, but average peak heights (and conversely, valley depths) are higher (and lower) reflecting the extreme nature of the fitness contributions (for $q = 2$, the fitness contributions are restricted to 0 and 1 as opposed to intermediate real values on either NK or NKp).

One interesting feature in NKq landscapes is the initial rise in average walk lengths from $K = 0$ to $K = 2$. A possible explanation for this is that the global optimum in a

smooth NKq landscape is a large connected component extending through the search space. As this component first begins to break up (the ruggedness increases) the optima initially become more difficult to reach. As the landscape continues to become more rugged, the increase in number of optima results in the walk length to reach them dropping, as in the NK landscape. Furthermore, if a large proportion of the search space is the global optima plateau, it would be expected that a significant number of the population of hill climbers would be initialised already at the global optima.

This conclusion is also supported by the change in neutral degree as ruggedness increases. For low levels of K , the landscape is equally neutral across the full range of fitness levels. As K increases, the neutrality is confined to the fitness levels around the landscape average and the peaks and valleys revert to being isolated points.

5. Evolution on neutral landscapes

Having covered in detail the structures of both the NK landscape and the neutral NKp and NKq variants, it is now possible to shift attention to the evolution of populations on these landscapes.

5.1. Introduction

Based on the differences in structure between the NK, NKp and NKq landscapes, it was expected that there would be a number of significant differences in the dynamics and performance of populations evolving on the landscapes produced by the original NK model and the neutral NKp and NKq variants. There is a considerable amount of research concerned with the dynamics of evolution (see Section 5.2 for an overview). However, much of this has focused on the effects of specific algorithms, specific operators or specific landscapes. There is not yet been any comprehensive attempt to compare different algorithms across neutral and non-neutral landscapes.

A related issue with which this study is concerned is the value of crossover as a genetic operator. There has been much debate questioning whether or not recombination improves the performance of evolutionary algorithms. Most of the research in this area has considered the role of crossover in constructing and disrupting building blocks. Another important benefit of crossover which has not received due attention is its ability to increase the level of genetic diversity in a population.

This phase of the study involved a series of simulations comparing several different evolutionary algorithms on a variety of neutral and non-neutral landscapes. A further set of simulations was run to investigate the ability of crossover to maintain diversity in a population across transitions between fitness layers. This section provides a brief overview of previous work on evolutionary dynamics and the principle processes and mechanisms affecting populations. This is followed by an illustration of the primary differences between evolution on neutral and non-neutral landscapes. The results of the comparative simulations and investigation into the effect of crossover on diversity maintenance are then presented and discussed.

5.2. Previous work on evolutionary dynamics

Ever since the metaphor of adaptive landscapes was proposed (Wright 1932), there has been an interest in finding methods of describing the way in which populations move across these landscapes under the force of selective pressure and genetic operators such as mutation and recombination. Several different approaches have been suggested ranging from representing a population using high-dimensional vectors (Nix & Vose 1991, Vose & Liepens 1991) through to statistical mechanics (Prügel-Bennett & Shapiro 1994). The balance that must be achieved is between a form of analysis that is too low-level, becoming quickly overwhelmed by large population sizes and genotype lengths, and one that is too high-level, requiring simplifications that result in too much information being lost.

One of the more complete attempts at developing a formal framework for analysing evolutionary dynamics has been the statistical dynamics approach, which combines a

dynamical systems approach with a statistical physics and stochastic processes approach (van Nimwegen, Crutchfield & Mitchell 1997, 1998; van Nimwegen, Crutchfield and Huynen 1999; van Nimwegen & Crutchfield 1999). This approach involves calculating statistical data on a population that can then be analysed non-linearly. An important consequence of the statistical dynamics approach has been an increased awareness of the interdependence of genetic operators and the difficulty this creates when attempting to optimise parameter settings. The remainder of this section describes each of the operators and pressures that result in evolution.

In the biological world, *mutation* has long been seen as a secondary operator whose purpose is to prevent the total loss of genetic information due to convergence. This was originally the view of the evolutionary computation community, who saw the recombination of building blocks as the primary evolutionary process. However, the new picture emerging is of mutation as the driving force behind evolution. In landscapes with neutrality, there is a balance between mutation leading populations to diffuse across neutral networks until they discover a link to a higher fitness network and selection ensuring that the population clusters around the fittest genotype thus far discovered. Research into the optimum mutation rate to minimise this search time (Ochoa, Harvey & Buxton 1998) suggests that it may be related to the idea of an error threshold – the mutation rate at which fit individuals are lost faster than selection can replace them (Eigen et. al. 1989; Nowak & Schuster 1989). Van Nimwegen and Crutchfield (1998) arrive at similar results and Barnett (2000) derives analytical results for the optimal mutation rates for a population-of-one *netcrawler* on a particular class of landscapes.

When evolutionary algorithms were first developed, *recombination* was thought to be the main operator in genetic search due to its role in constructing advantageous combinations of genes (Holland 1975). There is a significant body of research that has found that crossover improves the performance of evolutionary algorithms (Syswerda 1989, Schaffer & Eshelman 1991; Radcliffe 1991; Hordijk & Manderick 1995, Eiben & Schipper 1996, Suzuki 1997; Boerlijst 1996; Watson 2000). However, it has recently been suggested that recombination may be unimportant in evolutionary search except as an error-repair mechanism or macro-mutation operator (a large step across the fitness landscape, most probably to an uncorrelated region of search space). It has been shown that recombination has negligible effect on population dynamics on particular classes of landscapes with neutrality (van Nimwegen, Crutchfield and Mitchell 1997).

Effective genetic search may be characterised as a balance between exploration and exploitation. The genetic operators – mutation and recombination – aid exploration by moving a population over an adaptive landscape and maintaining a diversity of genetic material. The degree to which an evolutionary algorithm exploits successful solutions is generally determined by the method of *selection* used. While fitness proportional selection is the original and most commonly used selection method, it has been shown that the resulting selection pressure is not consistent over an evolutionary run and variations such as ranking, tournament and sigma scaling may improve search performance (Goldberg 1989; Whitley 1991; Mitchell 1996)

Random *genetic drift* is a stochastic process that affects the proportions of genes in a finite population. If a population's size is sufficiently small, and each pair of parents

only produce a limited number of offspring, the chances of the frequency of a particular gene staying the same from generation to generation is extremely small. Over long periods of time, this can lead to the genetic convergence of a population and even the creation of distinct species, as the forces of genetic drift operate independently on isolated sub-populations. An extension of this concept is the idea of *neutral drift*, in which the proportions of genes in a population on a neutral layer can undergo random fluctuations without any contribution by selective pressure. The behaviour of populations on a flat landscape was investigated (Derrida & Peliti 1991) in an attempt to characterise the way in which a population diffuses across a flat layer under the forces of mutation and random drift. Intuition suggests that these results, which demonstrate how complicated evolutionary dynamics can be even in the absence of selection pressure, would also extend to the behaviour of a population on a large neutral layer.

One of the continuing challenges to the evolutionary computing community is the setting of parameters such as population size, mutation rate, encoding and landscape parameters in such away that the resulting evolutionary search is optimised. The *no free lunch* theorem (Wolpert & Macready 1995) has been extended to demonstrate that, even for the limited class of real world problems, there is still no one algorithm which will be optimal for all problems (Sharpe 2000). Following this realisation, the challenge has shifted from finding the “holy grail” algorithm that will solve all problems to understanding the relationship between a particular problem landscape and the performance of evolutionary algorithms on it. For instance, Barnett (2000) demonstrated that a netcrawler is sufficient for search spaces without local optima. However, the addition of neutrality to a landscape does not necessarily imply an absence of local optima, and it is possible that there are neutral landscapes where a larger population is required to achieve optimal results.

5.3. Evolution by steps

Some of the primary differences between evolution on neutral and non-neutral landscapes can be apprehended immediately by comparing graphs of the average fitness over an evolutionary run (see Figures 5.1 and 5.2).




Figure 5.1: The average population fitness over 1500 generations on a landscape with no neutrality (Parameters: $N = 60$, $K = 12$ and $p = 0.0$ with a population of 200 and mutation rate of 0.0167).





Figure 5.2: The average population fitness over 1500 generations on a landscape with neutrality (Parameters: $N = 60$, $K = 12$ and $p = 0.98$ with a population of 200 and mutation rate of 0.0167).


A population on a non-neutral NK landscape moves uphill for a short period, but soon becomes trapped in the region of a fitness peak (see Figure 5.1). This outcome would be acceptable if the peak happened to be the global optimum. Clearly however, if the landscape is rugged, the chance of this occurring is very small. Furthermore, there is only a very small probability that the population will be able to escape this local optima through random mutation and continue climbing.

A population evolving on a neutral NK landscape with the same level of ruggedness undergoes a considerably different history (see Figure 5.2). There are long periods of stasis interspersed with brief transition periods where the fitness increases. The periods of stasis occur when the population is searching the current neutral network for a transition point to a higher fitness network. The sudden jump in fitness represents one or more individuals discovering one of these transition points, at which point selection pressure converges the population around them.

This dynamic can be observed by measuring the diversity of a population and comparing it to the average population fitness (see Figure 5.3). The measure of diversity that has been used here is the number of unique genotypes appearing in a population at each generation. Other measures of diversity, such as the average distance of each genotype from the population mean, can be used. The diversity rises up to the point at which a new fitness level is discovered, at which point it falls dramatically. The mechanisms of this dynamic are investigated further in Section 5.5.



(a)



(b)

Figure 5.3: Population diversity and population fitness on (a) an NKp landscape ($N = 60$, $K = 12$, $p = 0.99$, population size = 200, mutation rate = 0.0167) and (b) a Royal Staircase landscape ($N = 4$, $K = 8$, population size = 1000, mutation rate = 0.001). While the same pattern of diversity growth and convergence occurs on both landscapes, it can be observed more clearly on the Royal Staircase landscape (see Section 5.5.2 for a formal definition of this landscape model).

A secondary feature that can be seen from Figures 5.1 and 5.2 is the increased robustness of populations on landscapes with neutrality. While the average fitness of the population on the non-neutral landscape fluctuates considerably, that of the population on the neutral landscape is much more stable. Rather than an optimum consisting of a ‘peak’ it is more likely to be a ‘ridge’ or ‘plateau’, and it has been demonstrated (van Nimwegen and Crutchfield 1997) that a population will tend to evolve towards the region with the highest neutral degree of the current network, where it is at less danger of ‘falling off’ onto a lower network.

5.4. Comparative Simulations

5.4.1. Introduction

The primary focus of this part of the study was trying to gain some understanding of which method of evolutionary search is optimal under which circumstances.

In order to answer this question, a series of simulations was run on NKp and NKq landscapes (using NK landscapes as a benchmark). Varying degrees of neutrality were used in an attempt to demonstrate the similarities and differences between NK, NKp and NKq as regards the average population fitness and the best fitness achieved. A variety of different search techniques were used, including hill climbers and populations using either mutation alone or mutation with crossover.

A secondary purpose of these simulations was to investigate further any differences in the structural properties of the landscapes that may be indicated by observing the performance of search algorithms on them.

5.4.2. Methodology

The landscape models used were the NK, NKp and NKq models described above (Sections 3 and 4). Each model was tested using both neighbourhood and random interactions. The genotype length (N) was set at 100. The number of interactions (K) was varied from 0 (smooth landscape) to 99 (rugged landscape), with more emphasis placed on the lower values of K (0, 2, 4, 8, 16, 32, 64, 99).

Four different algorithms were used, a hill climber, population based with mutation and population based with mutation and single point crossover and population based with mutation and uniform crossover. For each algorithm on each landscape, the best fitness found and the average final fitness were measured. These were averaged over 20 different randomly generated landscapes.

The hill climber simulations used a population of 200 (independent) hill climbers, each of who started at a random location and took 1500 steps. The mutation rate used was 0.01, so an average of 1 gene per genotype was flipped at each step. The new genotype was accepted if its fitness was equal to or greater than that of the old genotype. The best fitness was the best fitness in the population after 1500 steps. The length of time until this fitness was first discovered was also recorded. The average fitness was the average of the 200 hill climbers after 1500 steps.

The population-based algorithms used a population of 200 and were run for 1500 generations. The mutation rate used was 0.01. Sigma-scaling selection was used to ensure a continuous pressure on the population. In the simulations using crossover, two offspring were created from each mating. Using uniform crossover, there was a 50% probability that the allele at each locus came from either parent. Using single-point crossover, a single crossover point was chosen at random and the offspring inherited half of their genotype from one parent and half from the second parent. The best fitness was the best fitness in the population at any point in the run. The generation at which this fitness was reached was recorded. The average fitness was the average fitness of the population after 1500 generations.

5.4.3. Results

The best and average fitnesses found by each algorithm on each landscape are presented in Figures 5.4, 5.5 and 5.6. Only the results for those landscapes with $K \leq 16$ are included, as several anomalies were noted in results of simulations run on landscapes with higher levels of epistasis. A further discussion of these landscapes is included in Appendix B.

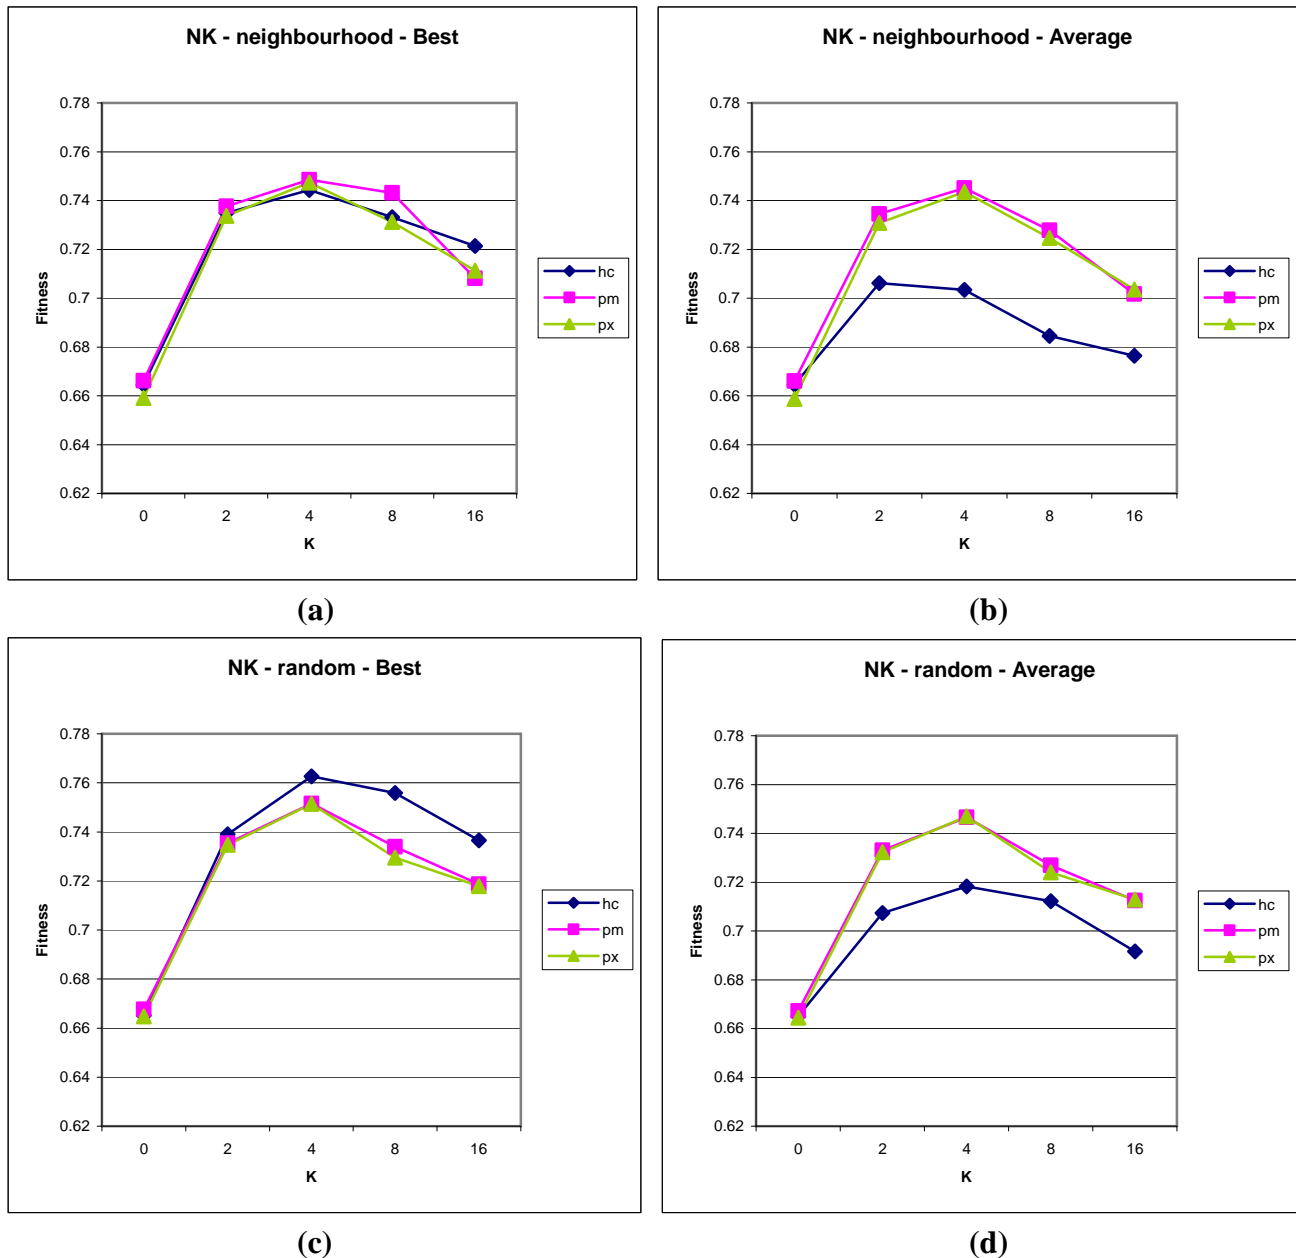





Figure 5.4: Comparative simulation results for NK landscapes with neighbourhood interactions (a) best fitness found, (b) average fitness of final population, and random interactions (c) best fitness found, (d) average fitness of final population. See text for full details.




(a)



(b)



(c)




(d)


Figure 5.5: Comparative simulation results for NKp landscapes with neighbourhood interactions (a) best fitness found, (b) average fitness of final population, and random interactions (c) best fitness found, (d) average fitness of final population. See text for full details.

The results of the simulations using population-based algorithms with single point crossover were generally similar to those of the populations using uniform crossover and have been omitted from the graphs for clarity (see Hordijk & Manderick 1995 for an investigation into crossover on NK landscapes).


Several general trends were observed in the comparative simulations. The hill climber algorithm tended to find equivalent solutions to those found by the population based algorithms at low levels of epistasis. As epistasis increased however, the hill climber consistently found better solutions (see graphs (a) and (c) of Figures 5.4, 5.5 and 5.6). The population based algorithms however, tended to have better average solutions (see graphs (b) and (d) of Figures 5.4, 5.5 and 5.6).




(a)



(b)



(c)



(d)

Figure 5.6: Comparative simulation results for NKq landscapes with neighbourhood interactions (a) best fitness found, (b) average fitness of final population, and random interactions (c) best fitness found, (d) average fitness of final population. See text for full details.

The hill climber algorithm tended to find better solutions on landscapes with random epistatic linkages than on those with neighbourhood linkages (see graphs (a) and (c) of Figures 5.4, 5.5 and 5.6).

The idea that there is a significant structure between the NKp and NKq implementations of neutrality was reinforced by the distribution of the best and average fitnesses found by each algorithm (see Figures 5.5 and 5.6). Whereas the NKq landscape was similar to the NK landscape in that the highest best and average fitnesses were found when $K = 4$, the NKp landscape differed considerably, with the best fitnesses found at $K = 8$.

5.5. Diversity maintenance between neutral layers

5.5.1. Introduction

As illustrated in Section 5.2 above, evolution on landscapes with neutrality follows a distinctive pattern. There are long periods in which average population fitness remains stable and population diversity increases, during which time the population is dispersing over the current fittest neutral layer. These epochs are interspersed with relatively short periods in which average population fitness increases and population diversity drops. This dynamic is caused by the discovery of a higher fitness layer by one or more individuals in the population. As their fitness is now higher than that of the rest of the population, they and their descendents are selected to reproduce more frequently, and their genotypes come to dominate the population. This is potentially a problem in a biological context – a significant loss of diversity in a host population could provide parasites with an advantage in a host-parasite co evolutionary system. In an evolutionary computation context, this loss of diversity could lead to premature convergence on a suboptimal fitness peak if critical genetic material is lost.

One way in which the diversity loss across the transition to a higher fitness level could be minimised is if recombination is used. If, during the transition, mating occurred between individuals on the higher fitness level and those on the lower fitness level, some of the genetic diversity contained in the old population at the lower fitness level should be transferred to the new population at the higher fitness level.

A series of simulations were designed to measure the drop in diversity across the first fitness increase in a Royal Staircase landscape (van Nimwegen & Crutchfield 1998). The Royal Staircase landscape model was used in preference to a neutral NK variant because the transitions between layers are clearer and more predictable. As demonstrated in Section 5.2 above, the drop in diversity across transitions between fitness layers is similar using both models.

5.5.2. Methodology

The Royal Staircase landscape described in (van Nimwegen & Crutchfield 1998) was used because, like the NKp and NKq models, it contains a high level of neutrality, but it also has the additional advantage of being highly predictable. The fitness levels and neutral degree are known in advance. This allowed the issue of diversity to be investigated in isolation from other structural issues associated with NK landscapes. It is expected that the results obtained here will also apply to population behaviour across fitness transitions on neutral NK landscapes.

The Royal Staircase fitness function is defined as follows (van Nimwegen & Crutchfield 1998)*:

1. Genotypes are specified by binary strings $s = s_1 s_2 \dots s_L, s_i \in \{0,1\}$, of length $L = NK$.
2. Reading the genotype from left to right, the number $I(s)$ of consecutive ones is counted.

* van Nimwegen and Crutchfield's (1998) original phrasing has been retained, some terminology has been altered for compatibility with the remainder of this report.

3. The fitness $f(s)$ of string s with $I(s)$ consecutive ones, followed by a zero, is $f(s) = 1 + \lfloor I(s)/K \rfloor$. The fitness is thus an integer between 1 and $N + 1$.

A landscape with $N = 4$ and $K = 8$ (i.e. 4 blocks of 8 bits = genotype of length 32) was used. The population size was 100 and was initially converged (all individuals identical) on a random genotype. The mutation rate was 0.001 and single point crossover at a rate of either 0.0 or 0.7 was used. Fitness proportional selection was used.

At each generation, the average population fitness, total population diversity (number of unique genotypes) and population diversity at each fitness level was recorded. The first and second derivatives of the average fitness were then used to locate the time frame over which the fitness increase occurred. The fall in total diversity across this time period, measured as a proportion of the diversity before the discovery of the new level, was then calculated.

5.5.3. Results


The results of the simulations indicated that the drop in diversity across the transition between fitness layers was much lower when crossover was used (see Table 5.1). Whereas the population using mutation alone lost almost half of its diversity on average, the population using mutation and crossover lost only around 20% of its diversity on average.

Table 5.1: Drop in diversity across the transition between the first and second fitness layers (given by the difference between the number of unique genotypes before and after the transition, divided by the number of unique genotypes before the transition). Average and standard deviation calculated from 20 runs. The Royal Staircase landscape was used ($N = 4$, $K = 8$, population size = 100, mutation rate = 0.001, crossover rate = 0.0 or 0.7).


	Mutation only	Crossover and mutation
Average	0.48	0.22
Standard deviation	0.17	0.12

The simulations also illustrated clearly some of the differences between the dynamics of populations evolving using mutation only and those evolving using mutation and crossover (see Figure 5.7).

The rise in diversity during the first 20 generations was similar in both cases. This is due to the fact that crossover will not have a significant effect until there is a certain level of diversity, due to mutation, in the population. After this point however, the differences are striking. Random mutation is likely to result in new genotypes that differ in only one or two places from the remainder of the population; therefore there is a distinct possibility of back mutations (mutations which return a previously mutated genotype back towards the population average). Crossover however, allows new genotypes whose hamming distance from the remainder of the population is much greater, and hence whose probability of undergoing back mutation is much smaller. Whereas the diversity in a population using only mutation stabilises around 25, the diversity of a population using crossover continues to rise until the majority of individuals have unique genotypes.



(a)



(b)

Figure 5.7: Diversity changes across fitness level transitions. An initially converged population of 100 evolving on a Royal Staircase landscape ($N = 32$, $K = 8$) with a mutation rate of 0.001 and a crossover rate of (a) 0.0 and (b) 0.7. Initial growth in diversity is similar, however with mutation only, a ceiling on diversity is soon reached. With crossover, a much higher level of diversity is possible. In both cases, diversity is lost during the transition from a lower fitness level to a higher fitness level. However, this diversity is recovered much more rapidly by the population using crossover.

The next notable difference occurs at the transition point between two fitness levels. The proportion of the population on each fitness level is similar in both cases, as this is a result of the selection pressure rather than the genetic operators used. While both populations suffer a drop in diversity, a population using crossover recovers much more quickly. Assuming that only one individual discovers the new fitness level, if crossover is not used all the diversity in the rest of the population on the lower fitness level is lost as these individuals are quickly selected against. The population on the higher fitness level will be descended from the individual who discovered the first block and will now have to “rediscover” the lost diversity. The growth in diversity on the higher fitness level is likely to be somewhat slower than the initial growth in

diversity as a certain proportion of mutations will either occur to individuals still on the lower fitness level, or will act to lower the fitness of an individual on the higher fitness level. The overall level of diversity is also likely to be lower, as the discovered block must remain converged in the population.

The population using crossover however, is able to “transfer” a lot of the diversity from the lower level to the higher level. So long as crossover does not interfere with the discovered block, a cross between an individual on the higher level and an individual on the lower level will result in an offspring on the higher fitness level who carries genetic material not inherited from the individual who discovered the first block. The growth in diversity at this point is therefore much more rapid than the initial growth in diversity. Once again, the overall level of diversity will be lower due to the presence of the first block throughout the population.

5.5. Discussion

Some of the general trends that are apparent across all the landscapes include:

- The best and average fitnesses found by each algorithm peak at $K = 4$ for the NK and NKq landscapes and $K = 8$ for the NKp landscape.
- The hill climber algorithm tends to find equivalent solutions to those found by the population-based algorithms at low levels of epistasis. As epistasis increases however, the hill climber consistently finds better solutions.
- The population algorithms tend to have better average solutions.
- The hill climber algorithm tends to find better solutions on landscapes with random epistatic linkages than on those with neighbourhood linkages.

The difference between the two different types of neutrality is further emphasised in these results. While the NKq results show similar trends to the NK results as K varies, the NKp results are significantly different. In particular, the NKp implementation of neutrality results in an optimal landscape that has a much higher level of epistatic interaction.

The significantly better performance of the hill climbing algorithms (at least as regards the best solution found) may be due to the fact that each hill climber is independently able to explore the region of the landscape around where it is generated and therefore increase the chances of an optimal solution being found (at the expense of many becoming trapped in lower local optima and reducing the average fitness found). The population-based solutions however, converge rapidly to a particular region of the search space (probably that around the most fit solution in the first few generations of an evolutionary run) that may not necessarily contain the optimal solution. This is more likely to occur at higher levels of epistasis when misleading optimal solutions may be discovered far from the global optimum.

While the initial conclusion may be that crossover has no benefits as a genetic operator, the second set of simulations demonstrates its usefulness. A population evolving using crossover and mutation is able to maintain a consistently higher level of diversity across an evolutionary run than a population using mutation alone. Furthermore, a population using crossover recovers more rapidly from the loss of diversity caused by convergence when a new fitness layer is discovered.

These results raise several issues in relation to the role of crossover in biological populations. For example, it may be plausibly argued that a population in which every fitness increase was accompanied by dramatic genetic convergence would be far more susceptible to invasion by parasites that evolved to exploit the dominant population genotype. By allowing diversity from a lower fitness population to be transferred to a higher fitness population, crossover enables populations to guard more carefully against exploitation by a parasite species. From an evolutionary computation point of view, the value of crossover lies in the fact that it allows for a much greater level of diversity in a population – thus increasing the probability of further fitness increases being discovered. One potential area for further work could be an investigation of whether crossover does increase the rate at which fitness increases are discovered. Some of the results relating to rate of diversity growth and the maximum level of diversity sustainable in a population could also benefit from a more analytic investigation, possibly in relation to existing models in population genetics.

6. General discussion and conclusions

6.1. Contributions

The contributions made by the study fall into two main categories, those related to landscape structure, and those related to the algorithm performance.

The following contributions have been made to an understanding of the structure of neutral landscapes:

- Kauffman's (1993) results relating to the NK landscape model were replicated successfully, establishing a baseline for comparison with the neutral NKp and NKq variants.
- A new visualisation technique was developed to illustrate landscape ruggedness and neutrality at a global level.
- The methodology used by Kauffman to investigate the structure of NK landscapes was extended to the neutral NKp and NKq variants, allowing for a direct comparison between structural properties of neutral and non-neutral landscapes.
- It was discovered that the NKp and NKq implementations of neutrality have significantly different structural properties.
- NKq landscapes are similar to NK landscapes with local ruggedness smoothed into neutral terraces:
 - o When K is small (a smooth landscape), the neutral degree remains stable across the different fitness levels of the landscape.
 - o When K is large (a rugged landscape), the neutral degree is greatest around the average landscape fitness decreases as you move up or down the fitness levels.
- NKp landscapes are structurally quite different to NK landscapes due to an interaction between the values of K and p .
 - o When K is small (a smooth landscape), neutral degree is greatest at the lowest fitness level and decreases as fitness increases.
 - o When K is large, the neutral degree is largest on the zero fitness 'floor' of the landscape, with little, if any, neutrality found as fitness increases.
- Whereas NK and NKq landscapes are generally optimal (i.e., realise the highest best and average fitness peaks) at around $K = 4$, NKp landscapes are optimal when they contain a greater number of epistatic interactions ($K = 8$ or $K = 16$).

The following contributions have been made to an understanding of algorithm performance:

- Across a wide range of neutral and non-neutral landscapes, hill climber algorithms consistently outperform population-based algorithms with regard to the highest fitness found on an NK landscape.
- Population-based algorithms, although not finding the highest fitness peaks, tend to find better solutions on average.
- Crossover does not have a significant impact on the performance of population-based algorithms on either neutral or non-neutral NK landscapes.
- Hill climber algorithms tend to perform better on NK landscapes with random (as opposed to neighbourhood) epistatic interactions.

Finally, the initial steps towards a new understanding of the role of crossover were taken. Crossover allows for a greater level of genetic diversity to be maintained in a population on a neutral layer and also minimises the amount of diversity that is lost when a population discovers a transition to a higher fitness layer.

6.2. Evaluation

The simulations demonstrate that computational evolution on a neutral landscape follows a predictable pattern. Initially, selection acts to rapidly converge a randomised population around the fittest individual, ensuring that individuals on suboptimal peaks are not selected. The genetic operators, mutation and crossover, then cause a population to diverge across the current neutral layer until one or more individuals discover a higher fitness layer. The portion of a population on the lower fitness layer is rapidly selected against, and a new population comprises primarily of descendents of the individual who first discovered the new layer.

Due to the random nature of an initial population and the selection process, there is a possibility that the initial convergence may occur in a suboptimal region of the landscape, hindering the search for the global optimum. The independence of a series of hill climbers can allow for a more thorough exploration of different areas of the landscape. This advantage comes at the cost of many of a population of independent hill climbers becoming trapped on suboptimal fitness peaks. On an neutral landscape, a hill climber may cease to evolve either when a fitness level is reached from which there are no transitions to higher levels, or when a plateau is reached that is so large that the transition to a higher fitness layer is not discovered within the time frame of the simulation.

The behaviour of a population undergoing a transition from a lower fitness level to a higher fitness level differs depending on the genetic operators used. If mutation is the only operator used, the population suffers a dramatic drop in genetic diversity at this transition point. This drop is due to the fact that the population on the higher fitness level will generally all be descended from the first individual to discover the new level*. All of the diversity in the population on the lower fitness level is lost. If

* Except for the relatively infrequent cases where more than one individual discovers the new fitness level.

crossover is also used, recombination between the individual who discovers the new fitness level and the remainder of the population on the lower fitness level may occur. Recombination between individuals on different fitness levels allows some of the genetic diversity of the population on the lower fitness level to be transferred to individuals on the higher fitness level. Therefore the drop in diversity is not so dramatic, and the recovery of diversity occurs more rapidly.

6.3. Conclusions

Based on the literature reviewed, neutrality is a plausible and appealing notion from a biological perspective. From a practical point of view, it may be concluded that there are benefits to recognising neutrality where it occurs in problem domains and tailoring evolutionary search algorithms to exploit it. However, experience with evolving populations on NKp and NKq landscapes suggests that achieving optimal performance on neutral landscapes is highly sensitive to landscape and population parameter settings. Therefore, adding neutrality to existing problem domains is not necessarily a straightforward process.

One of the strongest conclusions that may be drawn from the results of this study is that the way in which neutrality is implemented has a critical impact on both the structure of the resulting landscapes and the performance of evolutionary search algorithms on them. While NKq is qualitatively similar to NK in many respects, NKp is considerably different from both NK and NKq. It can not be claimed that either one or the other of these models is necessarily more useful or a more accurate representation of biological reality. Neutrality may operate at many levels in nature, from molecules up to organisms, and different models may be appropriate in different situations.

Finally, it was found that crossover confers diversity maintenance benefits upon a population that are not necessarily reflected in simple fitness-based measures of performance. This discovery reinforces the idea that the choice of algorithms and operators needs to be made in the context of the particular computational problem to be solved, or the particular biological phenomenon being demonstrated. In the same way that there may be more than one useful model of neutrality, there are likely to be a variety of optimal choices of algorithms and operators depending on the nature of the system being modelled.

6.4. Further work

There are several directions further research could take at this point. Some of the more interesting possibilities include:

- A further investigation of the non-performance based advantages of crossover on neutral landscapes, such as diversity maintenance in host-parasite models.
- An investigation of the anomalous results that emerged when populations were evolved on landscapes with a large number of epistatic interactions.
- Characterising the interaction between the K and p parameters of NKp landscapes analytically.
- An exploration of techniques to recognise neutrality in real problem domains and mechanisms to exploit it.

Reference List

- Altenberg, L. (1997) NK Fitness Landscapes. In T. Back, D. Fogel & Z. Michalewicz (Eds.), *The Handbook of Evolutionary Computation*. Oxford: Oxford University Press.
<http://pueo.mhpc.edu/~altenber/FTP/LeeNKFL.ps> (17/10/01)
- Barnett, L. (1997) *Tangled webs: Evolutionary dynamics on fitness landscapes with neutrality*. Master's thesis, University of Sussex.
<ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/lionelb/FullDiss.ps.gz> (17/10/01)
- Barnett, L. (1998) Ruggedness and neutrality - the NKp family of fitness landscapes. In C. Adami, R. K. Belew, H. Kitano and C. Taylor (Eds.), *ALIFE VI, Proceedings of the Sixth International Conference on Artificial Life* (pp. 18-27). Cambridge, MA: The MIT Press.
http://www.cogs.susx.ac.uk/users/lionelb/download/alife6_paper.ps.gz (17/10/01)
- Barnett, L. (2001) Netcrawling – Optimal Evolutionary Search with Neutral Networks. *Proceedings of the 2001 Congress of Evolutionary Computation (CEC2001), May 27-31, Seoul, Korea* (pp. 30-37). Piscataway, NJ: IEEE Press.
<http://www.cogs.susx.ac.uk/users/lionelb/download/netcrawl.ps.gz> (17/10/01)
- Boerlijst, M. C., Bonhoeffer, S. & Nowak, M. A. (1996) Viral quasi-species and recombination. *Proc. R. Soc. Lond. B*, 263, 1577-1584.
- Conrad, M., Ebeling, W., Volkenstein, M. (1993) Evolutionary thinking and the structure of fitness landscapes. *Biosystems*, 27, 125-128.
- Crow, J. F., & Kimura, M. (1970) *An introduction to population genetics theory*. New York: Harper and Row.
- Darwin, C. (1859) *The Origin of Species by means of Natural Selection*. Oxford: Oxford University Press.
- Derrida, B. and Peliti L. (1991) Evolution in a Flat Fitness Landscape. *Bull. Math. Biol.*, 53, 355-382.
- Eiben, A. E., & Schippers, C. (1996) Multi-parent's niche: n-ary crossovers on NK-landscapes. In H. M. Voigt, W. Ebeling, I. Rechenberg and H. P. Schwefel (Eds.), *Proceedings of the 4th Conference on Parallel Problem Solving from Nature*, number 1141 in Lecture Notes in Computer Science, (pp. 319-328). Berlin: Springer-Verlag.
<http://www.wi.leidenuniv.nl/~gusz/ppsn96.ps.gz> (17/10/01)
- Eigen, M., McCaskill, J. & Schuster, P. (1989) The Molecular Quasispecies. *Adv. Chem. Phys.*, 75, 149-263.
- Eldredge, N. & Gould, S. J. (1972) Punctuated equilibria: an alternative to phyletic gradualism. In Schopf, T. J. M., (ed.) *Models in Paleobiology*. pp. 82-115. San Francisco: Freeman.
- Fogel, D. B. (1995) *Evolutionary Computation: Towards a New Philosophy of Machine Intelligence*. New York: IEEE Press.

Fontana, W. & Schuster, P. (1998) Continuity in evolution: On the nature of transitions. *Science*, 280, 1451-1455.

Forrest, S. & Mitchell, M. (1993) Relative building-block fitness and the building block hypothesis. In D. Whitley (Ed.), *Foundations of Genetic Algorithms 2* (pp. 109-126). San Mateo, CA: Morgan Kaufmann.
<http://www.santafe.edu/~mm/Forrest-Mitchell-FOGA.ps.gz> (17/10/01)

Forst, C. V., Reidys, C., & Weber, J. (1995) Evolutionary dynamics and optimization: Neutral networks as model-landscapes for RNA secondary-structure folding-landscapes In F. Moran, A. Moreno, J. J. Merelo & P. Chacon (Eds.), *Third European Conference on Artificial Life*, number 929 in Lecture Notes in Artificial Intelligence. Berlin: Springer-Verlag.
<http://www.santafe.edu/sfi/publications/Working-Papers/95-10-094.ps> (17/10/01)

Gavrilets, S. (1997) Evolution and speciation on holey landscapes. *Trends in Ecology and Evolution*, 12, 307-312.

Gavrilets, S. & Gravner, J. (1997) Percolation on the fitness hypercube and the evolution of reproductive isolation. *J. Theor. Biol.*, 184, 51-64.

Goldberg, D. E. (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley.

Harvey, I. & Thompson, A. (1996) Through the labyrinth evolution finds a way: A silicon ridge. In T. Higuchi, (Ed.), *Proceedings of The First International Conference on Evolvable Systems: From Biology to Hardware (ICES96)*. Berlin: Springer-Verlag.
ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/sil_ridge.ps.gz (17/10/01)

Holland, J. H. (1975) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. Ann Arbor, MI: University of Michigan Press.

Hordijk, W. & Manderick, B. (1995) The usefulness of recombination. In F. Moran, A. Moreno, J. J. Merelo & P. Chacon (Eds.), *Third European Conference on Artificial Life*, number 929 in Lecture Notes in Artificial Intelligence, (pp. 908-919). Berlin: Springer-Verlag. <http://www.santafe.edu/~wim/Research/usefulness.ps.gz> (17/10/01)

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.
<http://www.tbi.univie.ac.at/papers/Abstracts/95-01-02.ps.gz> (17/10/01)

Huynen, M. A. (1996) Exploring phenotype space through neutral evolution. *J. Mol. Evol.*, 43, 165-169.
<http://www.santafe.edu/sfi/publications/Working-Papers/95-10-100.ps> (17/10/01)

Kauffman, S. A. (1993) *The Origins of Order - Organization and Selection in Evolution*. New York: Oxford University Press.

Kauffman, S. A., & Levin, S. (1987) Towards a General Theory of Adaptive Walks on Rugged Landscapes. *Journal of Theoretical Biology*, 128, 11-45.

Kimura, M. (1983) *The Neutral Theory of Molecular Evolution*. Cambridge, UK: Cambridge University Press.

Lipsitch, M. (1991) Adaptation on Rugged Landscapes Generated by Iterated Local Interactions of Neighbouring Genes. In R. K. Belew & L. B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms* (pp. 128-135). La Jolla, CA: Morgan Kaufmann.

Manderick, B., de Weger, M., & Spiessens, P. (1991) The genetic algorithm and the structure of the fitness landscape. In R. K. Belew & L. B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms* (pp. 222-229). La Jolla, CA: Morgan Kaufmann.

Maynard Smith, J. (1970) Natural selection and the concept of a protein space. *Nature*, 225, 563.

Mitchell, M., Forrest, S., and Holland, J. H. (1992) The royal road for genetic algorithms: Fitness landscapes and GA performance. In F. J. Varela and P. Bourguin (Eds.), *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life* (pp. 3-10). Cambridge, MA: MIT Press/Bradford Books.
<http://www.santafe.edu/~mm/RoyalRoad.ps.gz> (17/10/01)

Mitchell, M., & Holland, J. H. (1993) When will a genetic algorithm outperform a hill climbing algorithm? In Forrest, S. (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms*. (pp. 647-654). La Jolla, CA: Morgan Kaufmann.
<http://www.santafe.edu/~mm/ga-hillclimb.ps> (17/10/01)

Mitchell, M. (1996) *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press.

Naudts, B., Suys, D., & Verschoren, A. (1997) Epistasis as a Basic Concept in Formal Landscape Analysis. In Th. Baeck (Ed.), *Proceedings of the 7th International Conference on Genetic Algorithms* (pp. 65-72), San Francisco, CA: Morgan Kaufmann.
<http://islab.ruca.ua.ac.be/files/publications/icga97.ps.gz> (17/10/01)

Newman, M. & Engelhardt, R. (1998) Effect of neutral selection on the evolution of molecular species. *Proc. R. Soc. London B.*, 256, 1333-1338. <http://xxx.lanl.gov/ps/adap-org/9712005> (17/10/01)

Nix, A. E. & Vose, M. D. (1991) Modelling genetic algorithms with Markov chains. *Annals of Mathematics and Artificial Intelligence*, 5.

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.

- Ochoa, G. and Harvey, I. & Buxton, H. (1998) Error thresholds and their relation to optimal mutation rates. In D. Floreano, J-D. Nicoud, F. Mondada (eds), *Proceedings of the Fifth European Conference on Artificial Life (ECAL99)*. Springer-Verlag.
<ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/ecal99.ps.gz> (17/10/01)
- Ohta, T. (1973) Slightly deleterious mutant substitutions in evolution. *Nature*, 246, 96-98.
- Peliti, L. (1995) *Fitness Landscapes and Evolution*. Lecture given at the NATO ASI on Physics of Biomaterials: Fluctuations, Self-Assembly and Evolution, Geilo (Norway), March 27-April 9, 1995. <http://xxx.soton.ac.uk/ps/cond-mat/9505003> (17/10/01)
- Prügel-Bennett, A. & Shapiro, J. L. (1994) Analysis of genetic algorithms using statistical mechanics. *Physical Review Letters*, 72(9), 1305-1309.
- Radcliffe, N. J. (1991) Forma Analysis and Random Respectful Recombination. In R. K. Belew & L. B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms*. La Jolla, CA: Morgan Kaufmann.
- Reidys, C., Stadler, P. F., & Schuster, P. (1997) Generic Properties of Combinatory Maps: Neutral Networks of RNA Secondary Structure. *Bull. Math. Biol.*, 59, 339-397.
<ftp://ftp.itc.univie.ac.at/pub/PREPRINTS/rei-95-02.ps.gz> (17/10/01)
- Reidys, C., Forst, C., & Schuster, P. (1999) *Replication and mutation on neutral networks*. Santa Fe Institute Preprint 98-04-036.
<http://www.santafe.edu/sfi/publications/Working-Papers/00-11-061.ps> (17/10/01)
- Reidys, C., & Stadler, P. F. (1999) *Neutrality in fitness landscapes*. Santa Fe Institute preprint 98-10-089.
<http://www.santafe.edu/sfi/publications/Working-Papers/98-10-089.ps.gz> (17/10/01)
- Schaffer, J.D., & Eshelman, L.J. (1991) On crossover as an evolutionary viable strategy. In R. K. Belew and L. B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms*, (pp. 61-68). San Francisco, CA: Morgan Kaufmann.
- Sharpe, O. J. (2000) *Towards a Rational Methodology for Using Evolutionary Search Algorithms*. PhD Thesis, University of Sussex.
<http://www.cogs.susx.ac.uk/users/olivers/Thesis/thesis.html> (17/10/01)
- Shipman, R., Shackleton, M., & Harvey, I. (2000) The use of neutral genotype-phenotype mappings for improved evolutionary search. *BT Technol. J.*, 18(4), 103-111.
<ftp://ftp.cogs.susx.ac.uk/pub/users/inmanh/bt.pdf> (17/10/01)
- Suzuki, H., (1997) The optimum recombination rate that realizes the fastest evolution of a novel functional combination of many genes. *Theoretical Population Biology*, 51, 185-200.

Thompson, A. (1996) Silicon evolution. In J. R. Koza, D. E. Goldberg, D. B. Fogel and R. L. Riolo, (Eds.), *Genetic Programming 1996: Proceedings of the First Annual Conference, July 28-31, 1996, Stanford University*, (pp. 444-452). Cambridge, MA: MIT Press.

van Nimwegen, E., Crutchfield, J. P. & Mitchell, M. (1997) Finite populations induce metastability in evolutionary search. *Physics Letters A*, 229(2), 144-150.
<http://www.santafe.edu/projects/evca/Papers/fpimies.html> (17/10/01)

van Nimwegen, E., Crutchfield, J. P., & Mitchell, M. (1998) Statistical dynamics of the Royal Road genetic algorithm. In A. Eiben and G. Rudolph, (Eds.), *Theoretical Computer Science, Special issue on Evolutionary Computation*.
<http://www.santafe.edu/projects/evca/Papers/StatDynRRGA.html> (17/10/01)

van Nimwegen, E. & Crutchfield, J. P. (1998) Optimising epochal evolutionary search: Population-size independent theory. Santa Fe Institute Working Paper 98-06-046.
<http://www.santafe.edu/projects/evca/Papers/oecspsit.html> (17/10/01)

van Nimwegen, E., Crutchfield, J., & Huynen, M. (1999) Neutral evolution of mutational robustness. *Proc. Nat. Acad. Sci. USA*, 96, 9716-9720.
<http://www.santafe.edu/projects/evca/Papers/nemr.html> (17/10/01)

van Nimwegen, E., & Crutchfield, J. (1999) *Metastable Evolutionary Dynamics: Crossing Fitness Barriers or Escaping via Neutral Paths?* Working Paper 99-07-041. Santa Fe Institute. <http://www.santafe.edu/projects/evca/Papers/med.html> (17/10/01)

Vose, M., & Liepens, C., (1991) Punctuated Equilibria in Genetic Search. *Complex Systems*, 5, 31-44

Watson, R. A., Hornby G. S., & Pollack, J., B. (1998) Modeling Building-Block Interdependency. In Proc. of Parallel Problem Solving from Nature V. Springer. Pages 97-106. http://www.demo.cs.brandeis.edu/papers/hiff_ppsn.ps.gz (17/10/01)

Watson, R. A., & Pollack, J. B. (1999). Hierarchically-Consistent Test Problems for Genetic Algorithms. In Angeline et. al., (Eds.), *Proceedings of 1999 CEC*, (pp. 1406-1413). IEEE Press. http://www.cs.brandeis.edu/~richardw/papers/geccolate_hcnp.ps.gz (17/10/01)

Weinberger, E. D. (1990) Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63, 325-336.

Whitley, D. (1989) The GENITOR algorithm and selection pressure. In J. D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp 116-121). San Mateo, CA: Morgan Kaufmann.

Wolpert, D. H., & Macready, W. G. (1995). *No free lunch theorems for search*. Technical Report 95-02-010, Santa Fe Institute.
<http://www.santafe.edu/sfi/publications/Working-Papers/95-02-010.ps> (17/10/01)

Appendix A – Calculating fitness values using a hash function

A.1. Introduction

The simulation of NK landscapes requires the creation and storage of large tables of fitness values. Each of the N genes can make one of 2^{K+1} possible fitness contributions. As N and more particularly K become large, the generation and storage of these tables becomes computationally unfeasible.

An alternative method for generating fitness contributions using a hash function is described by Altenberg (1997). This involves converting the portion of the genotype under consideration from a binary value to an integer and using this integer to seed a random number generator. Altenberg's (1997) algorithm was implemented in Java (see NKHashA.java code below).

A.2. The first hash function – NKHashA.java

Code

```
import java.lang.Math;
import java.util.Random;

/****
 * NKHashA - generates fitness contributions for a particular gene.
 * Written by Nic Geard, June 2001. Based on an algorithm described by
 * Altenberg (1997)
 ****/

class NKHashA {
    private int base; // the base seed for the random number generator

    public NKHashA(Random rand) {
        base = rand.nextInt();
    }

    /****
     * getContribution - returns the first double generated by a random
     * number generator seeded by the sum of the base seed and the index.
     *
     * The index is an integer representation of the binary string
     * consisting of the value of the gene under consideration and the K
     * genes to which it is epistatically linked.
     ****/

    public double getContribution(int index) {
        Random rnd = new Random(base + index);
        return rnd.nextDouble();
    }

    public static void main(String[] args) {
        Random r = new Random();
        NKHashA t = new NKHashA(r);

        for(i=0; i<args[0]; i++)
            System.out.println(t.getContribution(i));
    }
}
```

Output

```
java NKHashA 8
0.5544185737011045
0.5546873190177194
0.5545977223443532
0.5541498283844896
0.5540602466122846
0.5543289919288995
0.5542393952555333
0.5537915012956697
```

As can be seen from the output of the `NKHashA.java` function, there is a high degree of correlation between the fitness contributions generated. To minimise this, the algorithm was modified to use the *product* of the base and the index to seed the random number generator, and the *second* random number in the resulting sequence was taken as the fitness contribution (see `NKHashB.java` code below).

A.3. The second hash function – `NKHashB.java`

Code

```
import java.lang.Math;
import java.util.Random;

/****
 * NKHashB - generates fitness contributions for a particular gene.
 * Written by Nic Geard, June 2001. Based on an algorithm described by
 * Altenberg (1997)
 ****/

class NKHashB {
    private int base; // the base seed for the random number generator

    public NKHashB(Random rand) {
        base = rand.nextInt();
    }

    /****
     * getContribution - returns the second double generated by a random
     * number generator seeded by the produce of the base seed and the index.
     *
     * The index is an integer representation of the binary string
     * consisting of the value of the gene under consideration and the K
     * genes to which it is epistatically linked.
     ****/

    public double getContribution(int index) {
        Random rnd = new Random(base * index);
        rnd.nextDouble();
        return rnd.nextDouble();
    }

    public static void main(String[] args) {
        Random r = new Random();
        NKHash t = new NKHash(r);

        for(i=0; i<args[0]; i++)
            System.out.println(t.getContribution(i));
    }
}
```

Output

```
java NKHashB 8
```

```
0.2405364156714858  
0.4613891869680744  
0.6412802388987969  
0.3011201988804495  
0.9136292136199583  
0.1763451538999552  
0.8331273201109255  
0.8963327294050983
```

As can be seen from output of the `NKHashB.java` function, there is much less correlation between the fitness contributions generated.

A.4. *Caveat emptor*

As noted in the body of the report (see Section 3.3) some discrepancies between the results found on NK landscapes generated using fully enumerated lookup tables and those found on landscapes generated using the hash function method were noted. Further work is required to investigate the properties of the random number generator used and establish the validity of using hash functions to generate fitness contributions on NK landscapes with large values of N and K .

Appendix B – Anomalous results on highly epistatic landscapes

B.1. Introduction

The comparative simulations presented in Section 5.4 of this report were run on a variety of NK, NKp and NKq landscapes with genotype length N fixed at 100 and the number of epistatic interactions K varied from 0 to 99. It was expected that the best and average fitness found on these landscapes would peak at certain low values of K and fall off as $K \rightarrow N - 1$. The results however, revealed that for certain values of K , the algorithms performed significantly more poorly than anticipated (see Figure B.1).

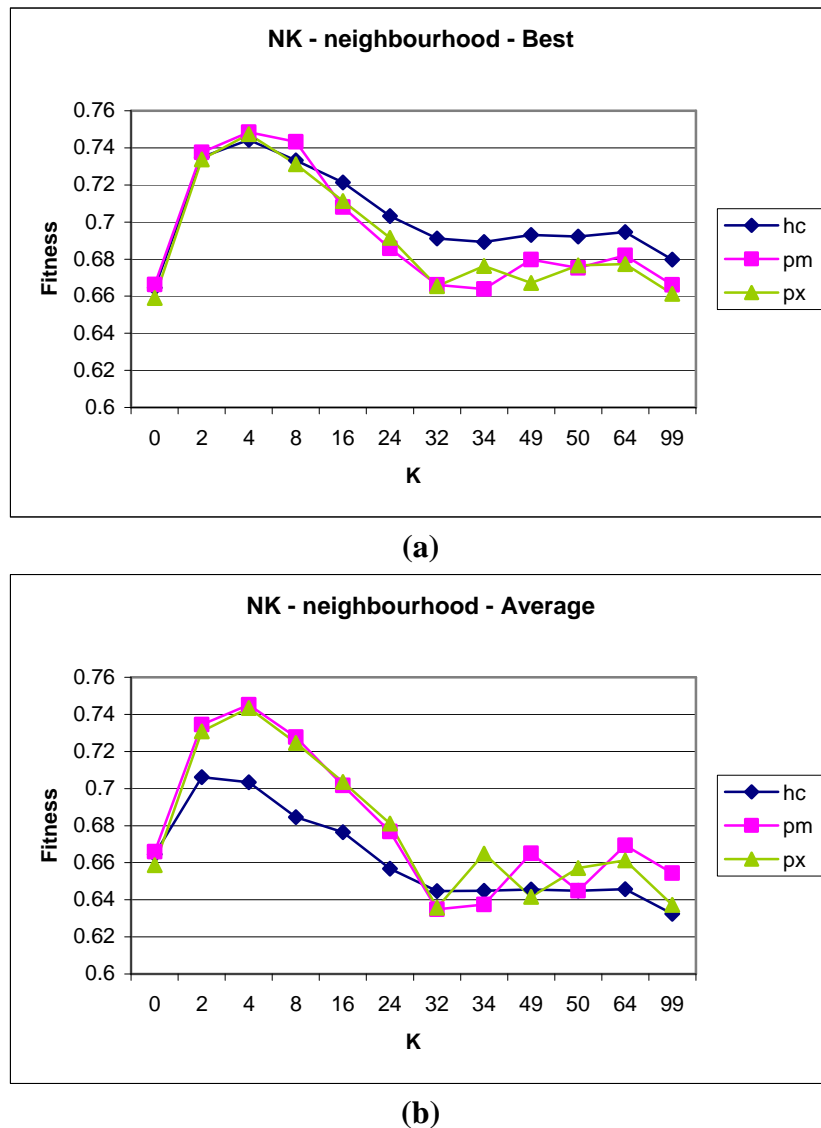


Figure B.1: Anomalous results at certain levels of epistasis. These results were obtained using the same method as described in Section 5.4.2 (values averaged over 20 runs, $N = 100$, population size = 200, number of steps/generations = 1000, mutation rate = 0.001, uniform crossover). hc = hill climber, pm = population with mutation only, px = population with mutation and crossover.

B.2. Discussion

Several possible explanations for these results were considered. The first was that the unexpected results were due to properties of the random number generator and the hash function method used to generate fitness contributions on landscapes with high values of K .

The second possibility considered was that there might be some implicitly modular structure to NK landscapes. If $K + 1$ is a large factor of N (e.g. $K = 32$ or $K = 49$) then it is possible that the genotype may contain several nearly independent blocks (i.e. with few epistatic interactions between them). The discovery of one of these fitness blocks, particularly if it causes a large fitness increase, may result in *hitch-hiking* genes becoming fixed in the population and preventing discovery of further blocks. The fact that the phenomenon was much more pronounced when population-based algorithms rather than hill climbers were used supported this possibility. Further investigations were carried out by running simulations with critical values of K (see Figure B.1) but initial results were inconclusive.

A more thorough investigation of this phenomenon is necessary, but was outside the scope of this study.