Adaptive evolution in static and dynamic environments

Thesis

UNRESTRICTED

Adaptive Evolution in Static and Dynamic Environments

Anthony John Hirst BEng(Hons)

Department of Psychology, Open University

1997

Submitted in partial fulfillment of requirements for the degree of *Doctor of Philosophy*

Primary subject area: Evolutionary Computation

Author number: M7164646
Date of submission: 19 September 1997
Date of award: 21 December 1998

Abstract

This thesis provides a framework for describing a canonical evolutionary system. Populations of individuals are envisaged as traversing a search space structured by genetic and developmental operators under the influence of selection. Selection acts on individuals' phenotypic expressions, guiding the population over an evaluation landscape, which describes an idealised evaluation surface over the phenotypic space. The corresponding valuation landscape describes evaluations over the genotypic space and may be transformed by within generation adaptive (learning) or maladaptive (fault induction) local search.

Populations subjected to particular genetic and selection operators are claimed to evolve towards a region of the valuation landscape with a characteristic local ruggedness, as given by the runtime operator correlation coefficient. This corresponds to the view of evolution discovering an evolutionarily stable population, or quasi-species, held in a state of dynamic equilibrium by the operator set and evaluation function. This is demonstrated by genetic algorithm experiments using the NK landscapes and a novel, evolvable evaluation function, *The Tower of Babel.* In fluctuating environments of varying temporal ruggedness, different operator sets are correspondingly more or less adapted.

Quantitative genetics analyses of populations in sinusoidally fluctuating conditions are shown to describe certain well known electronic filters. This observation suggests the notion of Evolutionary Signal Processing. Genetic algorithm experiments in which a population tracks a sinusoidally fluctuating optimum support this view. Using a self-adaptive mutation rate, it is possible to tune the evolutionary filter to the environmental frequency. For a time varying frequency, the mutation rate reacts accordingly. With local search, the valuation landscape is transformed through temporal smoothing. By coevolving modifier genes for individual learning and the rate at which the benefits may be directly transmitted to the next generation, the relative adaptedness of individual learning and cultural inheritance according to the rate of environmental change is demonstrated.

Acknowledgments

To all those who have helped me thus far, I would like to offer my utmost thanks. Specifically -

Jon Rowe, who took on the challenge of organising a collection of loosely related ideas into a rather more coherent whole, and greatly aided my appreciation of the ESP model.

Martin Le Voi for listening as I struggled to explain what I really, really meant, and who helped bring this thesis to a presentable conclusion.

Inman Harvey for helping me organise my thoughts on ruggedness.

Russell Anderson for his comments on learning and signal processing techniques.

Judy Greene, for kickstarting the write up just when it was needed.

George Kiss, for introducing me to the OU and encouraging me to feed on new ideas, & Hank Kahney, for offering encouragement in his own unique way - would that you both had seen this day.

Trevor Collins and Adrian Thompson for their invaluable comments, stimulating on-topic conversation and valued encouragement.

Paul MacEachern and Richard Colenso for coffees, chats & somewhere to stay.

The Open University, for funding my research.

And of course -

Emma, for a life outside thesis, Jess, for coming into our lives and putting it all into perspective, and Mum and Dad, for being there always.

Table of Contents

List of Figures

List of Tables

Chapter 1 - Introduction

"Prefacing stories a reader has not yet read, since it demands the analysis of plots that it may be inconvenient to deal with in advance, is a somewhat impossible task..."

JL Borges, in "Afterword" to The Book of Sand.

However...

"In Science - in fact, in most things - it is usually best to begin at the beginning...
"And so...I'll give you the Axioms of Science. After that, I shall exhibit some Specimens. Then I shall explain a Process or two. And I shall conclude with a few Experiments. An Axiom, you know, is a thing that you accept without contradiction. For instance, if I were to say 'Here we are!', that would be accepted without any contradiction, and it's a nice sort of remark to begin a conversation with. So it would be an Axiom..."

L Carroll, from "The Professor's Lecture" in Sylvie and Bruno Concluded.

1.1 So Here We Are...

This thesis is intended as a contribution to the *Evolutionary Computation* (EC) community. In addition, it is hoped that it will be of more than passing interest to theoretical and evolutionary biologists, philosophers of evolution and the users of evolutionary models in general.

The reported work uses *genetic algorithms* (GAs) to characterise the trajectory of evolving populations over a *performance landscape* under the influence of selection and the genetic operators of mutation and crossover. By comparing the performance of an individual on some task with the performance of the other individuals making up the population, a *selective value* is assigned to each individual. Selection of individuals to act as parents of individuals in the next generation is then carried out on the basis of awarded selective values. Individual performances may be influenced by allowing *individual adaptation* to the selective environment. The transformational effect of such adaptive measures on the performance *surfaces* or *landscapes* over a population will be one of the major themes of this thesis. By considering evolution as acting on populations *as a whole*, I question the assumption (in both biology and EC) that evolution acts as a *global optimiser*. In its place is a view of evolution as a discoverer of *robust* solutions that are selectively buffered (*canalised*) against mutational or environmental change. This position is supported by means of a framework within which

several useful concepts for describing (any) evolutionary system are identified. This includes notions of within generation local search, both adaptive (*learning*) and maladaptive (*fault induction*). The effect of such *hybrid* algorithms (i.e. ones combining evolutionary and local search components) on the evolutionary dynamic is also considered in theory and experiment, using both the well known NK landscapes (Kauffman, 1993), and an original function, *The Tower of Babel.*

In natural environments, the selective environment of an individual is likely to *fluctuate* over the lifetime of the individual. Artificial Life (Alife) simulations in which populations of individuals evolve in some artificial world inhabited by predators and essential resources are often offered as models of such dynamic systems, but are difficult to consider analytically. By providing a selective environment that fluctuates over time in a well-specified way, it is possible to identify rather more closely the behaviour of a population evolving under such conditions. The performance landscapes over a population are perhaps more usefully represented by *seascapes* which change over relatively short periods of time. Sinusoidally fluctuating environments present one such easily characterised dynamic and these shall be discussed both experimentally, and with respect to theoretical biology analyses. One particularly interesting result is that the behaviour of an evolving population in such an environment is identical to the response of certain electronic filters used in signal processing. Whether this similarity results from the assumptions made in deriving the biological result, or signifies some deeper relationship between models of evolutionary systems and the theory of signal processing is posed as a question to philosophers of science.

The use of self-adaptive parameters adds a further dimension to the evolutionary dynamics of populations evolving in nonstationary environments. In the first instance, self-adaptive mutation rates may be used to 'tune' an evolutionary filter, or allow a population to track environments with varying rates of change. Secondly, by coevolving individual and cultural learning strategies in environments of different temporal grain (i.e. environments with different rates of change) a theoretically predicted trade-off between these strategies may be modeled.

1.2 Towards a General Model of Evolution

The pursuance of evolutionary computation to date has focused on its application as an optimisation tool in computer science and engineering, and as a potential simulation tool for theoretical biology. It is possible, however, that the evolutionary metaphor and the modeling techniques of EC reflects the (unconscious) striving for an overarching 'General Evolutionary Theory' or 'General Theory of (Evolvable) Adaptive Systems', as in the theoretical treatment of adaptive systems by John Holland (for example, (Holland, 1992), and its forerunner (Holland, 1962)) and (Atmar, 1994) on computer simulations of evolution; see also (Csanyi, 1981; Faber & Proops, 1991). Within the biological (and social) studies of evolution, there is a search for generality that will allow the evolutionary metaphor to be invoked at all levels of natural organisation (Vrba & Eldredge, 1984; Mani, 1991; Csanyi & Kampis, 1991), or the evolutionary dynamic to be modeled analytically (Williams, 1970). In economics, too, the evolutionary metaphor is being actively researched (Nelson, 1995; Saviotti & Metcalfe, 1991; Silverberg, 1988). Though the route to general theories has so far proved elusive (for example, the limited success of General Systems Theory and Cybernetics ((Pask, 1965) for a cybernetician's view of evolution)) the search is still on; and though a satisfactory general theory may be unattainable, cross-fertilisation of ideas between disciplines employing evolutionary metaphors should be regarded at least as a partial success of the generalist approach in the development of domain specific knowledge. The transfer of analytic *techniques* between disciplines may also allow for a reinterpretation of the mathematical models generated in one field in terms of those derived in another. The deconstruction of quantitative genetics analyses of evolution in terms of signal processing primitives is offered in chapter 6 of this thesis as an example of this approach.

This thesis may thus be seen as part of the search for general evolutionary principles, although it is presented with a biological emphasis since biology has to date been the dominant player in evolutionary thinking. It is likely, too, that many of the components of a general theory will be sourced initially from biology which provides a well developed mathematical theory of (biological) evolution. At some point, however, analytic tractability fails and numerical methods are required. The use of *evolutionary algorithms* (EAs) allows a further step to be taken, specifically the *simulation* of the theoretical models using real, rather than ideal, population distributions. EAs, of course, also have a theoretical component, although this has tended not to reflect the biological analyses. Consequently, I suggest that

EC and theoretical biology (particularly models of population genetics and artificial selection/breeding) do and will have much to offer each other.

However, it is at present difficult for computer scientists to employ mathematical results from evolutionary biology because as yet there are few specific pointers into that literature. For the biologists, there is little reason to read the EA literature since many of the simulations are hard to relate to accepted biological models. By explaining to the EC community (at least qualitatively) the assumptions behind, and results of, their population genetics models, the theoretical biologists would only benefit from the results of a now informed EC community. In a similar vein, (Miller, 1995) offers a set of guidelines for pursuing Alife as theoretical biology; (Noble, 1997) presents a more recent commentary on the status of Alife. My personal belief is that there is much that the philosophy of science too may have to offer in uniting these two approaches. One possible approach is through the analysis of the model structures (Lloyd, 1994) used in each camp.

As the evolutionary metaphor becomes applied outside biology, mechanisms not supported by natural systems, but nevertheless viable within *synthetic*, evolutionary systems, are likely to be identified. A typical example would be an abiological inheritance mechanism such as the direct transmission of acquired characteristics between generations through some form of *Lamarckian inheritance*[1]. To this extent, EC should not necessarily be constrained by the mechanisms of the dominant evolutionary theory (i.e. the biological one), although care must be taken not to divorce the generalist approach from accepted scientific methodology. Two other paths are available to the future development of EC (pure optimisation applications aside). Firstly, the limitation of EC models as models of biological mechanisms alone. Note that this will reduce to the general approach as models are explored through a relaxation of constraints. Secondly, one may restrict the argument to the evolution of systems *in logico*. This approach arises from arguments as to the otherness of computational life on the basis of the distinct environment it is likely to inhabit when compared to the natural environment of biological, carbon life (see for example (Ray, 1995) who develops a similar theme further). Again, it may be that in the limit this approach will identify *as a subset of computational models* the biological models and again the generalist goal will have been achieved.

---

[1] In the 'natural' world of introduced, genetically engineered systems, the inheritance of *designed* characteristics has been implemented from outside the system. The effects of genetically engineered captive organisms escaping into the wild and invading the wider evolutionary gene pool are at present a matter of debate.

Finally, returning to EC as an optimisation tool, there are good reasons here for investigating abiological mechanisms. For example, one might offer the assumption that the motivation behind a consideration of the inheritance of acquired characteristics (IAC) is to come up with robust evolutionary algorithms tuned to the environment and requirements of a particular evolving population of plastic, synthetic individuals.

1.3 Structure of the Thesis

In chapter 2 I present a model of evolution from the biology literature that provides the conceptual foundation of the models discussed herein. After introducing evolutionary algorithms in general, and genetic algorithms in particular, I briefly discuss the variety of mathematical models available in evolutionary biology. I then go on to show how each of the two disciplines may offer useful insights to the other in an interwoven review of the literature of each field as it relates to the introduction of variation into evolving populations in both static and dynamic environments.

In chapter 3, I develop a framework for describing evolutionary algorithms that is intended to clarify the terminological confusion identified in chapter 2 that permeates the literature of both evolutionary computation and evolutionary biology. A dominant theme of this thesis, the notion of a *valuation landscape*, is constructed in close conjunction with the values it represents and the search space structure over which it is imagined. A clarification of *valuation*, *evaluation* and *selective value* is offered, as well as that of *ruggedness* in both temporal and search-spatial terms. The characteristic trajectory of an evolving population from an initially random state to a quasi-species in a state of dynamic equilibrium is then presented, and discussed in relation to a measure of local ruggedness, the runtime operator correlation coefficient. The framework is completed by a discussion of the ways in which variation may be introduced into an evolving population.

In chapter 4, I employ the valuation landscape metaphor to discuss how within generation local search operators (both adaptive and maladaptive) may transform this structure. Previous explanations have suggested that local search 'smoothes' the landscape in some way. This notion is clarified and the corresponding 'sharpening' behaviour of maladaptive operators demonstrated. The applied style of selection is also shown to influence the form the surface transformation may take.

Chapter 5 presents an experimental counterpart to chapter 4. In particular, I use a genetic algorithm over the well known NK landscapes to show how the trajectory of an evolving population may be influenced by local search operators using a simple 'Darwinian' inheritance scheme and also the inheritance of acquired characteristics. A demonstration that evolution behaves as a robust optimiser, rather than a global optimiser, is also presented using the novel *Tower of Babel* evolvable evaluation function.

In chapter 6, I address the question of evolution in temporally fluctuating selective environments. Firstly, I present the observation that the mathematical biology models of populations tracking a sinusoidally fluctuating environment are identical to the Butterworth low pass filter response in the continuous case, and resembles the response of certain class of digital filter in the discrete case. Secondly, using a sinusoidally varying evaluation function (Cobb, 1990), I show how the former biological model sensibly predicts the steady state behaviour of a simple generational GA model, even when certain assumptions used in the derivation of the biological model are violated. The behaviour of a model employing a self-adaptive mutation rate is also studied using a sinusoidally fluctuating optimum whose frequency is time dependent.

In chapter 7, I introduce self-adaptive plasticity operators into the genetic algorithm model, again subject to a sinusoidally fluctuating evaluation function. The resulting behaviour of the evolving gene pool and the expressed phenotypes is again in accord with predictions from theoretical biology. By allowing traits acquired through adaptive plasticity (learning) to be transmitted to the next generation, a crude model of culture results. Theoretical results predict that in a fluctuating environment, there is a trade off between individual learning and cultural inheritance. By co-evolving the degree of adaptive plasticity and the rate at which traits acquired through this simple form of learning are transmitted to the next generation, this trade-off is experimentally demonstrated.

Finally, in chapter 8, I summarise the contributions of this thesis and offer a critique of its contents. Suggestions for the further development of the ideas introduced herein are also presented.

1.4 Published Work

Several parts of this thesis have appeared in published form, (although not all have undergone the process of peer review):

(Hirst, 1996d) - unreviewed; forms part of section 3.3.4.
(Hirst, 1997b) - forms parts of sections 3.2, 4.2, 4.4 and 4.5.
(Hirst, 1997c) - forms parts of sections 6.3 and 7.2-7.4.
(Hirst, 1997a) - forms part of sections 6.2 and 6.3.

Other published work not directly related to this thesis, although still within the domain of adaptive system design, includes:

- a precursor to the work on adaptive and maladaptive hybrid algorithms, (Hirst, 1996b)
- an early, but now outdated, review of *evolvable hardware* (EHW), (Hirst, 1996c);
- an unreviewed discussion document on adaptive processor design using reconfigurable hardware and evolutionary techniques, (Hirst, 1996a).

Chapter 2 - Evolutionary Models in Theory and Practice

Ummon's Sidetrack

A Zen student told Ummon: 'Brilliancy of Buddha illuminates the whole universe.' Before he finished the phrase Ummon asked: 'You are reciting another's poem, are you not?' 'Yes,' answered the student. 'You are sidetracked,' said Ummon.
Afterwards, another teacher, Shishin, asked his pupils: 'At what point did that student go off the track?'

In 'Zen Flesh, Zen Bones', compiled by P Reps, Pelican (1971).

2.1 Introduction

In this chapter, I present a broad overview of evolutionary models in general, and genetic algorithm (GA) models in particular. The scope is necessarily broad since I must introduce the traditional way of thinking about these models before counterpointing them with the alternative model that is described in chapter 3, explored in chapter 4 and illustrated by experiment in chapter 5.

In 2.2 and 2.3 I introduce evolutionary thinking and develop a typical evolutionary model drawn from the biological literature. In 2.4 I describe the various evolutionary algorithms (EAs) with reference to a canonical GA, and in 2.5 overview the mathematical structures used in biological and evolutionary computation analyses.

Mechanisms for supporting *adaptive evolution* in which variation is sourced at an adaptive rate are described in 2.6. Section 2.7 presents a case study of one particular adaptive GA, showing how tools of evolutionary biology may be usefully applied in the analysis of such a model. The effectiveness of various adaptive strategies applied in a fluctuating environment are reviewed in 2.8.

2.2 A Brief History of Modern Evolutionary Theory

Historically, there are two trends of thought as regards variation and heredity. The first was a dogmatic assertion of Special Creation, in which all variation was designed by God at the very genesis of life. The favoured inheritance mechanism which supported such a claim was known as *emboitement* (Bowler, 1973) which may be explained as follows. Consider the most recent individual in an ancestral line. Before it was born, it was represented as a tiny miniature of itself within its parent. The parent, containing the miniature of its unborn offspring, was likewise at one time carried in miniature form by its parent. Thus we go down the ancestral line, a nested set of Russian Dolls, each parent containing a miniature of its offspring, that miniature containing a miniature miniature (*sic*) and so on. An alternative view of inheritance favoured the introduction of variation through the inheritance of acquired characteristics (IAC), a position often referred to as Lamarckian inheritance[2]. In this case, pangenetic theories of inheritance were used to support the notion of some physical representation of traits which could be passed between individuals (Zirkle, 1945). In Darwin's offered theory of heredity, itself a reworking of the classical theory of pangenesis (Olby, 1963), he describes each part of an animal's body as producing a discrete 'gemmule' that acts as a seed for a similar part in the child, the full set of gemmules being collated by the reproductive organs and passed on in reproduction.

The classical explanations for the wide variety of life were eventually overthrown by Darwin's theory of evolution by natural selection. This theory, which presupposes the existence of living entities capable of reproduction, is actually a theory in two parts, one pertaining to inheritance, the other to selection itself[3]:

- *the theory of common descent* stresses the need for continuity between generations (like begets like) by the transmission of some hereditary material; speciation is viewed as a

---

[2] Lamarck favoured the idea of *use inheritance*. That is, an individual was capable of changing its form within its lifetime through *use* e.g. acquiring calluses on feet. By some unmentioned mechanism, those adaptive traits acquired through use were thought to be transmitted to the offspring of an individual. See, for example, (Mayr, 1972).

[3] For a recent discussion of the status of evolutionary theory, see (Lewis, 1981).

branching, rather than linear, process. This theory itself may be broken down into two further subtheories:

a) a theory of heredity concerning the *ideal* replication and the transmission of germline (genetic) information from parent to child;

b) a theory regarding the introduction of genetic variation;

- *the theory of natural selection (NS)*, which describes the competitive mechanism underlying phylogenetic branching, in which differential reproductive success arising from the expression of inherited variation tends to increase the proportion of well adapted individuals within a population.

This bipartite approach is reflected in several guises: for example, in population genetics, (Altenberg & Feldman, 1987) employs a transmission function and a selection operator, both of which are required to describe an evolving system; and an influential philosophical account of evolutionary theory by (Hull, 1980), on the controversy surrounding levels of selection, sees organisms as comprising *replicators* and *interactors*. The distinction is also implicit in the data/program duality of DNA, a feature predicted by von Neumann, prior to the discovery of DNA, in his theory of self-reproducing systems (von Neumann, 1961).

Despite the many successes of the Modern Synthesis of NS and mathematical models of genetics, there is increasing concern that the Neo-Darwinian position it underwrites is incomplete. On the one hand, claims have been made regarding the neutralist stance - that micro-evolutionary change is a consequence of genetic drift and selectively neutral mutations (distancing the two theories even more and relegating NS to a rather more minor role). On the other, what is called for is the incorporation of a theory that more closely identifies the relationship between the genetic complement of an individual and its expressed form (morphology, behaviour etc.) that is acted on by selection. In one particular view, the notion of self-organisation during morphogenesis is deemed necessary to explain the emergence of the ordered wholes which are the constituent parts of higher levels of organisation and which are targeted by selection (Kauffman, 1993). What this amounts to is the suggestion that the theory of common descent would be better represented by a theory of the reproduction of developed, morphological forms i.e. a third subtheory governing morphogenesis is required along with the two covering replication and variation. This position is propounded by what (Sterelny, Smith & Dickinson, 1996) terms the 'Developmental Systems Theorists' in their critique of that camp. This third subtheory would cover the way in which the transmitted germline material is related to the phenotypic variation acted on by NS. Self-organisation is

thus seen as explicating the programmatic nature of the genetic material by providing robust developmental pathways and employing the implicit information of self-assembling structures (see, for example, (Atlan & Koppel, 1990)); it may also be applied parsimoniously in accounts of homeostasis. The restricted class of systems exhibiting self-organisation also provides one way of rationally explaining both consistent developmental pathways and the repeated biological discovery of a limited number of morphogenetic forms.

There is also increasing interest on the effect of within lifetime behaviour on the evolutionary dynamic - for example, the recent resurgence of interest in the Baldwin Effect, which describes a Darwinian mechanism for the apparent inheritance of acquired characteristics ((Baldwin, 1896; Simpson, 1953), (Anderson, 1997; Turney, Whitley & Anderson 1996; Belew & Mitchell, 1996) for recent reviews in terms of EC). The role of adaptive behaviours, for example, may be incorporated into the original theories either as a subtheory of common descent as it relates to the sourcing of expressed variation; or as a subtheory of natural selection on the effect of adaptivity on competition.

Whether the concerns raised by the Developmental Systems Theorists or the 'lifetime behaviours' camps will force a shift in the perceived structure of evolutionary theory is unclear. For example, the issue may be resolved by recasting the bipartite view in terms of a theory of the transmission of variation, and a theory of selection between variants. Common descent is then seen as representing a hard constraint on the scope of the genetic transmission system, but other inheritance strategies, such as cultural inheritance (Boyd & Richerson, 1983; Feldman, Aoki & Kumm, 1996), may be accommodated.

2.3 Modeling Evolution

The fields of evolutionary biology and evolutionary computation are rife with terminology that is used inconsistently across, and even within, their respective disciplines ((Mahner & Kary, 1997; Byerly & Michod, 1991; Provine, 1986) for discussions on several confused concepts in the biological literature). The first task, then, is to characterise the evolutionary system described above with a set of concepts that shall be used throughout this thesis.

Figure 2-1 depicts the interactions that take place in an evolving system. Although taken from a biological text (Scheiner, 1993), it applies equally for EAs, illustrating at a coarse level the

intended similarity between the two processes. Individuals may be thought of in terms of two coupled components. Their genetic representation (*genotype*) and an expression of that genotype in a particular environment in morphological or behavioural terms (the *phenotype*). *Evolution* refers to the application of the genetic operators to the genetic complement of individuals selected according to their relative phenotypic adaptedness to the environment.

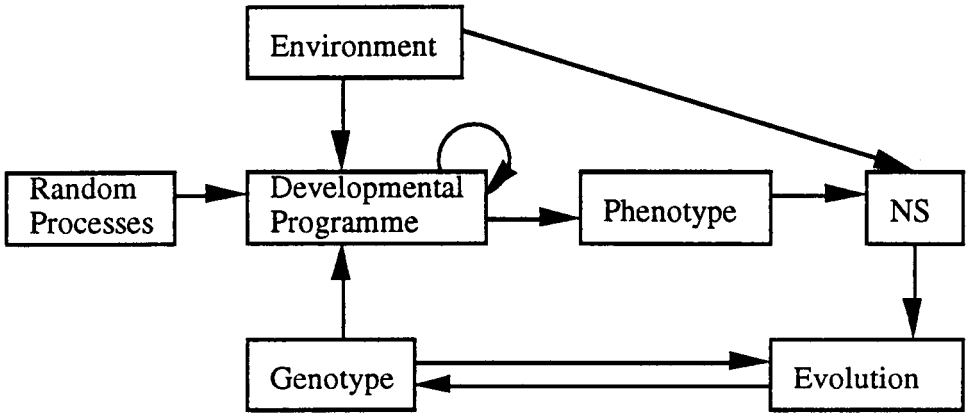

Figure 2-1 (taken from (Scheiner, 1993), his figure 1, p36, his caption): "Schematic of the relationship between genetics, development, environment, and evolution. The environment plays a dual role, affecting the developmental process and setting the fitness function."

Phenotypes are derived from genotypes by a *developmental map*, or developmental programme, which may or may not be influenced by the environment. That is, the environment may be involved in 'decoding' the selectively assessed phenotype from the genotype, as for example in the discovery or acquisition of adaptive behaviours through learning. Such genotype-environment interaction may result in what is known as *phenotypic plasticity*, defined by West-Eberhard as:

> "... the ability of a single genotype to produce more than one alternative form of morphology, physiological state and/or behaviour in response to environmental conditions" (West-Eberhard, 1989).

This covers not only maturation (phenotypic development, developmental plasticity) and learning (adaptive plasticity) (Hart, 1994), but also maladaptive plasticity, such as deliberately induced faults within the EC domain (3.3). Historically, the environment has been assumed to play no influential role other than in selection, but Scheiner, West-Eberhard and a growing number of other biological theorists have started to argue for its important

input into the developmental process (for example, (West-Eberhard, 1989; Scheiner, 1993; Ho & Saunders, 1984), and (Sterelny et al., 1996) for a critique of their position).

The developmental programme, then, mediates the expression of phenotypic plasticity and is responsible for determining the particular phenotype developed from any given genotype. The feedback loop presents a way of accounting for *epigenetic inheritance* (Holliday, 1987; Jablonka & Lamb, 1989). Variation in phenotypes frequently arises from mutational change at the genotypic representation level, figure 2-2a. Phenotypic variation in the absence of genotypic variation is also possible through plastic maturation, and may occur in one of two ways (Gomulkiewicz & Kirkpatrick, 1992): development may be switched between different developmental pathways at a single, crucial point in an individual's life history (usually quite early on), according to the environmental state at that point, and held constant thereafter (*non-labile traits*, figure 2-2b); or, alternative phenotypes may be acquired 'on-line' throughout the lifetime of an individual (*labile traits*, figure 2-2c; for example, homeostatic responses).

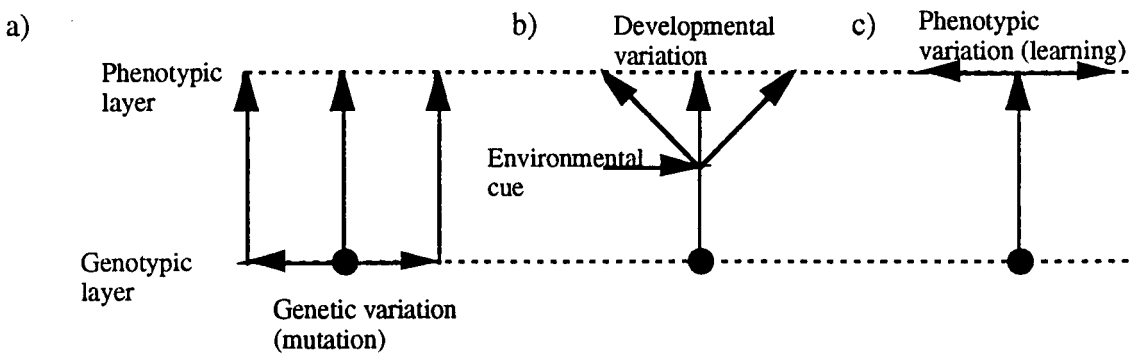

Figure 2-2: Introduction of variation: a) at the genotypic/representational level resulting from genetic mutations; b) during the developmental process as a result of environmental cues (e.g. reaction norms); c) at the phenotypic, fitness assessed level resulting from individual or cultural learning.

Non-labile traits are typically discussed in terms of *reaction norms* - "systematic and repeatable responses to some environmental cue" (Via 1993), p353. That is, genotypically identical individuals may consistently develop differently in different environments, the environment acting to switch between developmental pathways (i.e. reaction norms). This is in contrast to *developmental noise*, that "phenotypic variation apparently uncorrelated with environmental cues, occurring by errors of development or in response to random environmental variation" (p353). There is currently an ongoing debate in the biological

literature about the sense in which reaction norms and trait plasticity can be said to evolve (Via, Gomulkiewicz, De Jong, Scheiner, Schlicting & Van Tienderen,1995; Gomulkiewicz & Kirkpatrick, 1992; Schlicting & Pigliucci, 1995; Via, 1993; Wagner & Booth, 1997).

A third style of plasticity that I shall concentrate on extensively in the model I shall develop in chapter 3 is *learning*. Learning presents an opportunity for an *individual* to improve its selective value within the course of its lifetime by acquiring a rather more adapted phenotype than the one representing the direct expression of its inherited genotype. Traits that result from learning are distinguished from labile traits since they are taken to be acquired (even generated) in an active sense, rather than representing an automatic response to some environmental cue. However, the distinction is not a clear one and the path from non-labile, through labile to learned trait represents a continuous rather than a discrete classification[4].

In EC, learning may be modeled either trivially as a random process in which several alternative solutions are tried independently and the best selected therefrom (i.e. trial and error learning, a very weak form of reinforcement learning); or in a rather more directed way such as in the supervised learning of neural networks, or many of the machine learning algorithms.

In a given environment, some individuals within a population will be more or less adapted (that is, successful at surviving and hence producing offspring) than the others. In EC models, the relative adaptedness of an individual within a population is measured by a) developing an individual's *phenotype* from its *genotype* by some developmental (genotype-phenotype) map; b) applying an *evaluation* (or *objective*) function to some aspect of that *phenotype* (e.g. its structure or behaviour; note that the *e*valuation of a *phenotype* developed from some genotype gives the *valuation* of the inherited *genotype* (Aßelmeyer, Ebeling & Rosé 1996), discussed in more detail in 3.3.1); and c) transforming the result, according to the particular selection scheme applied, to a selective value (or *fitness propensity*) (selection function transformations are considered more thoroughly in 3.2.1). Ideal sampling of the

---

[4] This closely resembles problems in the theory of agents regarding agent capabilities (Kiss, 1992). A rule of thumb for distinguishing between non-labile and learned responses may be drawn from a control example - non-labile traits correspond to the selection of a single, fixed control law (with no environmental arguments) from a set of possibilities, following a single decision at a particular time; labile traits represent a control law that contains at least one environmental *variable*; learned traits are capable of adapting control law *coefficients* to effectively generate novel control laws.

population on the basis of selective values generates the selected breeding population with an appropriate distribution. The (*a posteriori*) *fitness* of an individual is then a measure of the actual number of times it was used in the breeding process; under ideal sampling and perfect, asexual inheritance, the fitness of an individual in the current generation gives its proportion of the next generation population[5,6]. In Scheiner's model, NS takes as arguments the environmental state and the phenotypes under selection, passing the resulting fitness propensity information (which is associated with the responsible genotypes) to the Evolutionary operators.

The 'Evolution' box of figure 2-1 encompasses the application of the genetic operators to 'genetic individuals' on the basis of their selective values revealed by applying NS to the particular phenotype the genes develop in a given environment. 'Evolution' in this sense is the sole source of genetic variation. Simple inheritance strategies (implicit in the bi-directional genotype-evolution link with no phenotypic input) are seen as a way of moderating the amount of (fitness) information that may be passed to following generations.

Cultural inheritance, the influence of previous phenotypes on the development of new phenotypes, may be included in Scheiner's model in a strong and a weak form as shown in figure 2-3. By providing the directed link, *A*, from the phenotype box to the environment, a strong form of cultural inheritance in which the *environment* is transformed by phenotypic behaviour is supported. Such an abstraction suggests that the way a population transforms the environment through behaviour may influence the evolutionary dynamic both through the role the environment has to play in development, and also via the influence it has on selection. In effect, then, populations co-evolve with the environmental legacy of their antecedents. A second, weaker form of cultural inheritance is modeled by link *B* from the phenotype box to the developmental programme. In this way, the phenotypic behaviour of

---

[5] The notion of fitness is one of the conceptual problems of interest to philosophers of biology, in that it is often applied in a way that reduces evolutionary theory to the level of a tautology (e.g. the fittest are those that are best adapted are those that are the fittest are those...) (Sober, 1984). (Byerly& Michod, 1991) offer a way out of this conundrum by isolating several different notions of 'fitness'. By distinguishing between *evaluation*, *selective value* and *fitness*, as I do here, the danger of circular argument is significantly reduced.

[6] (Stephens, Garcia Olmedo, Mora Vargas & Waelbroeck, 1997) go one step further in defining the *effective fitness* of an individual, which corresponds to its *expected* proportional representation in the next generation population, where crossover and mutation *are* allowed.

one generation may influence the development of the next, for example through education. Generally, when I refer to cultural evolution I shall be employing the *weak* sense of this term; this also represents the stance typically taken towards cultural evolution, (for example, as in (Boyd & Richerson, 1983)). The limiting case of inheritance between generations is the direct, germline transmission of one population's *acquired* phenotypic adaptations to the next ('Lamarckian inheritance'). This is represented in figure 2-3 by the direct link *C* from the phenotype box to the genotype one. Weak cultural inheritance and Lamarckian inheritance both represent instances of the inheritance of acquired characteristics (IAC). These models are explained in more detail in 3.3.2.

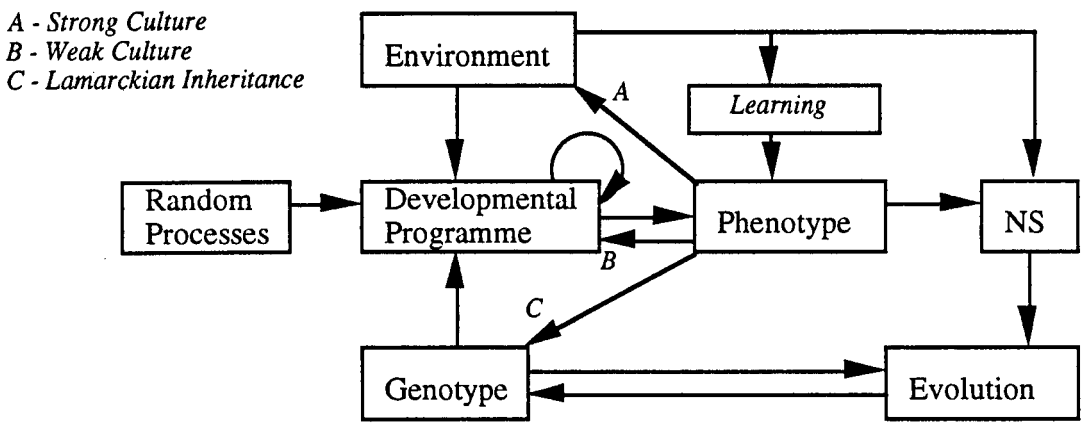

Figure 2-3: Augmented evolutionary model incorporating cultural and Lamarckian inheritance. B and C represent examples of IAC at appropriate levels of description.

If IAC is to be supported, the wider question of the *scope* of the transmission system is raised. Depending on who transmits what to whom, and under the assumption of non-overlapping generations, the following transmission schemes may be identified (Feldman & Leland, 1996):

- *vertical transmission*: in which a parent directly transmits information to its offspring; this mode is the assumed transmission model in most evolutionary discussions, as for example in simple Darwinian models of evolution, or Belew's *lineage model* (Belew, 1990);

- *horizontal transmission*: the transmission of information between (unrelated) members of the current population; this may occur at a phenotypic level (for example, through 'discussion' or disease) or at genetic level (for example, horizontal transmission in

bacteria, transmission via parasites or viral infection); horizontal transmission is not discussed further in the context of this thesis;

- *oblique transmission*: in which the members of one generation transmit information to unrelated members of the next, as is typically the case in cultural evolution. Two special cases are distinguished: frequency dependent transmission, where the characteristics of the most frequently encountered traits in a population are transmitted, and indirect transmission (for example, the *broadcast model* of (Belew, 1990), where a few specific individuals (such as a leader) are looked to for inspiration.

2.4 Introducing Evolutionary Algorithms (EAs)

Evolutionary algorithms (EAs) are the general counterpart to the algorithms employed in *genetic algorithms* (GAs), *genetic programming* (GP), *evolutionary programming* (EP) and *evolution strategies* (ES) - some gentle introductions to the field are given by (Mitchell & Forrest, 1994; Schwefel, 1994; Bäck & Schwefel, 1996). EAs are essentially a metaphor for Darwinian evolution by Natural Selection (NS): a population of individuals is somehow assessed and individuals are assigned a selective value on the basis of their performance in some environment with respect to the rest of the population; better adapted individuals are then probabilistically selected and used to 'breed' the next generation of individuals. Variation is introduced by random mutation, and crossover (mating between pairs of selected individuals, whereby each individual contributes part of its genetic makeup to the offspring). In the following section, I shall compare the different flavours of EA, and then go on to briefly comment on the ways in which EAs may be usefully applied.

2.4.1 Classifying Evolutionary Algorithms.

Genetic algorithms provide the most general approach to EAs, in that individuals are genotypically specified using a binary representation. For this reason, the thesis as a whole will be presented in terms of GAs. By introducing constraints on the way in which the bitstring codes for potential solutions, and constraining the operators that may be applied to manipulate bitstrings, the other classes of EA may be derived.

2.4.1.1 Genetic Algorithms.

Genetic algorithms (Holland, 1992; Goldberg, 1989; Whitley, 1994) manipulate low level (bitstring) representations that may be developed to realise some behaviour producing phenotype. Mutations are applied by randomly flipping bits at a fixed mutation rate and crossover should strictly be applied (at a fixed crossover rate) at random positions within a bitstring genotype, or 'chromosome', figure 2-4a. Generally, two parents are assumed for crossover, although multi-parent crossover is possible (for example, (Eiben, Raue & Ruttkay, 1994)). Asexual replication is achieved by the transmission of the parental genetic material to its offspring with only the mutation operator applied. In single point crossover, as in figure 2-4b, a single between loci crossing point is selected at random. In two point crossover, two between loci points are selected and the middle sections swapped. In uniform crossover, each bit is considered in turn, being independently sampled from either parent. Mutation and crossover may also be characterised in terms of the *neighbourhood* they induce, figure 2-4c; for example, a single bit/individual mutation operator will induce the *mutation neighbourhood* of an individual containing all the individuals exactly one Hamming unit away (see (Hirst, 1996d), and 3.2.2.3).

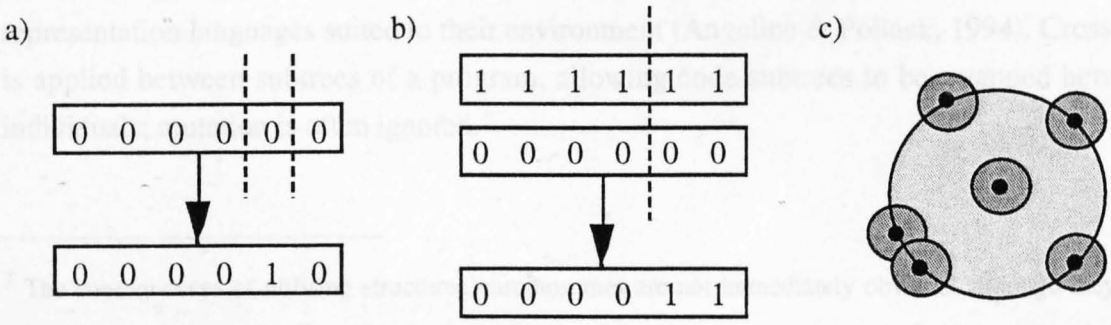

Figure 2-4: Simple genetic operators in a GA: a) bitwise mutation, and b) single point crossover of simple genotypic bitstrings; c) a 2-dimensional schematic representation of the effective mutation neighbourhood of each individual in a population and the crossover neighbourhood of that population.

The rates at which the genetic operators are applied are usually set as global parameters at the start of a run, although several adaptive strategies are possible (2.6). Often, chromosomes are structured using higher cardinality alphabets than the binary one, and packets of several bits are used to describe numerical phenotypic atoms, such as real or integer numbers. In such cases, crossover may be applied in a structured way *between* these 'numerical' gene packets. Similarly, structured forms of mutation may be used, as, for example, in a tri-state

representation system where the genotypic atom is a 2 bit packet that may be set by mutation to one of only three discrete states, rather than four, as 2 bits allows[7]. Other 'enhanced' operators have been developed to cope with a range of representations suited to particular problem domains (for example, ordered representations in the case of permutation problems (Davis, 1991)). Work has also been done on the adaptive evolution of parameter values (e.g. (Smith, 1997) for a recent review).

2.4.1.2 Genetic Programming.

A relative newcomer on the EA scene, *genetic programming* as championed by (Koza, 1992; Koza, 1994) uses relatively high level, often problem domain specific, hierarchical representations to evolve software programs that satisfy some behavioural (input-output) specification, howsoever defined[8]. These programs, most commonly represented by Lisp S-expressions or C linked lists, are comprised of primitive functions, (often simple logical or arithmetic functions), and terminals (such as inputs or constants), which are assigned randomly in the initial population. In line with GAs, evaluation proportional selection was favoured early on, although alternatives are now frequently applied. Work has been done on compressing subtrees into complex functions (modules) in an attempt to evolve representation languages suited to their environment (Angeline & Pollack, 1994). Crossover is applied between subtrees of a program, allowing code subtrees to be swapped between individuals; mutation is often ignored.

---

[7] The consequences of utilising structured chromosomes are not immediately obvious, although they may affect the search dynamics. For example, for loci structured as integers over the range 0...2, crossover between parents x.1.x and x.2.x will realise similar offspring. Using a binary chromosome and binary coding, there are four possible offspring since what is being crossed are x.0.1.x and x.1.0.x. If simple binary coding is used, realising a phenotypic expression over 0..3, one of these offspring must be 'repaired' to give a valid phenotype. Naive robust encodings such as {00 - **0**; 01 - **1**; 10, 11 - **2**} or {00, 11 - **2**; 01 - **0**; 10 - **1**} alleviate the need for repair, but each introduces a different bias (i.e. favoured allele) under both crossover and mutation.

[8] In this respect, GP as evolving deterministic programs resembles the intentions of the 'first wave' of Evolutionary Programming (Fogel, Owens & Walsh, 1966) which concentrated on the evolutionary manipulation of finite state machines. EP has since developed into a rather more numerical optimisation technique.

2.4.1.3 Evolutionary Programming and Evolution Strategies.

Evolutionary programming (Fogel, 1994) and *evolution strategies* (Schwefel, 1981; Bäck & Hoffmeister, 1994) both remove the genotype-phenotype distinction of GAs by working at the level of phenotypes. In each approach, consistency in gross level behaviour between generations was initially achieved by limiting evolutionary operations to mutation, although crossover was later introduced into the ES approach. 'Individuals' in the EP/ES sense represent mean phenotypic traits of populations, with ES and higher order variants of EP also coding independent trait variances on the genome. To this extent, mutation, rather than being a bitwise operation, represents either the perturbation of the inherited real valued trait according to a normal distribution with the evolved variance of that trait, the variance itself having first being mutated according to a given distribution (ES); or the setting of the trait value to one sampled from a distribution with mean set according to the initial gene value and variance calculated as a function of individual fitness (first order EP). A consequence of this self-adaptive approach is that the effective mutation rate evolves over time, in contrast to the fixed mutation rate frequently applied in GAs. Initial populations are sampled from a small region of the search space rather than the whole space, ES utilising a population chosen by mutating around a single individual, EP by sampling the initial population members from a uniform distribution over some predetermined range. One further distinction between EP and ES' is their employment of different selection strategies (EP uses tournament selection whereas ES employs elitist truncation selection (Bäck, 1994)).

2.4.2 Evolutionary Optimisation.

EAs may be applied as models of the evolutionary process, and as such represent a modeling/simulation environment for theoretical biology as discussed in 2.5.3. However, they are more frequently applied in the optimisation of complex parameterised problems although there is growing appreciation of the fact that "Genetic Algorithms Are NOT Function Optimizers" (de Jong, 1993), or at least that "genetic algorithms should not be thought of as *global* optimizers" (Vose, 1993b) (my italics). The experiments presented in chapter 5 address the sense in which evolution locates 'optimal' solutions more closely.

Even if GAs are not essentially *global* optimisers, this doesn't mean that they aren't capable of discovering useful solutions to *design problems* which require a special kind of *robust* optimisation:

- (Parmee, 1996) discusses the use of adaptive search methods in industrial design problems. GAs are shown to be good at finding 'high performance regions' where the valuation landscape is smooth, rather than rugged regions of the space where "solutions are susceptible to mild parameter perturbation which is likely to result in severe degradation."
- (Thompson, Harvey & Husbands, 1996; Parmee, 1996) both comment on the ability of the evolutionary search process discovering design solutions that are unavailable to traditional methods. For example, GAs may be freed of many of the simplifying constraints imposed to make the design process tractable (Thompson et al., 1996);
- (Thompson, 1995) discussed the implicit fault tolerance (or graceful degradation properties) of evolved problem solutions, specifically as it relates to the ability of an individual to cope with within lifetime faults that correspond to single mutations of the individual. Drawing on the work of Eigen, he showed how nearly converged populations evolve towards "optima that have surrounding regions of high fitness [evaluation]" rather than "isolated optima standing out alone amongst low fitness [evaluation] genotypes". In a follow up experiment over Kauffman's NK landscapes, (Thompson, 1996) showed how populations under weak selection tended to converge to regions of the search space in which the mean evaluation of the single bit mutants of a local optimum was higher than one would otherwise expect; that is, single mutations do not affect evaluation as adversely as might be expected.

What these results appear to suggest is that the evolutionary search process may be fruitfully thought of in terms of robust design rather than pure optimisation. Experiments presented in chapter 5 illustrate this point further.

By recasting GA evaluation functions in design terms, and taking away the emphasis on global optimisation, the ability of evolutionary search to discover robust solutions may be better exploited. For example, (Motta & Zdrahal, 1996) describe a class of design problems that appear to be ideally formulated for GA methods, although their review is limited to 'classical' problem solving techniques. *Parametric design problems* require the designer to find a set of *parameters* that satisfy certain *design requirements* according to a set of *constraints*, and minimise some *optimisation criterion*. A formal specification of the

parametric design problem is described which includes a set of *preferences* that should also be taken into account. The optimisation function formalises these preferences in a measure that may also take into account the financial cost of a design. For a *valid* design, all requirements must be achieved, and no constraints must be violated. A design is *complete* if all parameters are set within their prespecified ranges. The aim of the problem is to produce design *solutions* (complete *and* valid designs) that minimise the optimisation function. Where robust problem solutions are required, (i.e. solutions tolerant to parameter perturbations) I suggest this approach may be likened to a) a search for regions of the search space that encompass valid designs; b) optimisation of valid designs within those regions.

2.5 Mathematical Models of Evolutionary Algorithms and Evolutionary Biology

One might reasonably expect that analytic results from population genetics (the mathematical component of evolutionary biology) are applicable to EA models, but until recently there has been little transfer between the two disciplines[9]. In this section, I will briefly consider the form mathematical analyses have taken in each discipline and identify several cases in which a successful transfer of methods has been achieved. A particular example (based on (Hinton & Nowlan, 1987) and the large number of commentaries that followed) is considered in more detail in 2.7.

---

[9] I find it interesting that Goldberg's classic textbook on GAs, (Goldberg, 1989), was classified under the Dewey system as 006.3.1 (1. Combinatorial optimisation. 2. Algorithms. 3. Machine Learning) whereas a recent textbook, (Mitchell, 1996), is classified 575.101.13 (1. Genetics - Computer simulation. 2. Genetics - Mathematical models). What may be termed almost the *original* work, (Holland, 1992), initially published in 1975, is classified 564.5.01.5118 (1. Adaptation (Biology) - Mathematical models. 2. - Adaptive control systems - Mathematical Models). Another area of the library worth a visit are the 629.8s (control systems), where you'll find the early ECAL proceedings; however, the ALife workshop proceedings are classed under biological simulations (577.20). Where EAs are used as an optimisation technique within a particular domain, classification tends to be according to that domain. The importance of classification should not be under-rated if one subscribes to the view that serendipity plays a not insignificant role in basic research.

2.5.1 Biology.

Mathematical biological models of evolution come in several flavours, and are studied using a variety of techniques, both numerical and analytic. They do tend to have one feature in common, however, and that is the assumption that evolution is an optimisation process (Maynard Smith, 1978).

Population genetics studies the gene dynamics of a system (of fixed population size) as it evolves over time; evolutionarily stable strategies (ESS) offer a game theoretic approach towards predicting equilibrium proportions of genes or phenotypes ('strategies' employed in maximising some payoff) that are stable against invasion by alternative strategies; and the ecological approach of population dynamics uses growth equations (based on the Lotka-Volterra equation) to predict the growth of a population according to the fitness of its members and the carrying capacity of the environment. A very readable introduction to the whole area is (Roughgarden, 1979).

Population genetics itself may be further subdivided into single or two locus models, which assume either the diploid case with sexual reproduction or haploid asexual reproduction[10]; and quantitative genetics, which concentrates on the evolution of the distribution of expressed phenotypes, selection favouring 'individuals' close to some 'optimal' trait value (Mühlenbein, 1995). Many of the models require infinite populations, and in the latter case, panmixia (global mating). Populations structured in terms of age or spatial distribution may be catered for, albeit at the expense of simplicity. Single locus models were developed by considering simple, independent, additive traits, the equations coming in either deterministic or stochastic (Markov chain or diffusion) forms. Quantitative genetics focuses on the phenotypic level at which traits are expressed as continuous variables, rather than purely discrete instances as in the single locus models.

A different approach, offered by (Eigen, McCaskill & Schuster, 1988) (which represents an abridged view of (Eigen, McCaskill & Schuster, 1989)), describes the evolution of molecular sequences capable of imprecise replication. Optimisation is viewed in the sense of "global stabilization", specifically the discovery of a "*master sequence* together with its frequent mutants...(that is) a *quasi-species*" (p.6882). For our purposes, the identity of the master sequence may correspond either to the best adapted individual in the population (quasi-

[10] EAs tend to offer a hybrid of these two, namely haploid, sexual recombination.

species), or the 'consensus individual' ('centre of gravity' of the population) around which the population is centred. In a stationary selective environment with a single optimum, these definitions are likely to identify the same individual. However, under more complex scenarios, different individuals may be so described and in this case the master sequence should be identified in a way appropriate to the task at hand. Under certain conditions, the population may be held in a state of dynamic equilibrium as a result of the *mutation-selection balance*, which describes the interplay of these two processes (i.e. of mutation and selection)[11]. Alternatively, the quasi-species may roam rather more freely along (mutationally) connected *selectively neutral networks*.

The quasi-species view suggests that "the fitness of a given sequence is not solely determined by its own selective value. Rather, neighboring sequences contribute too, through mutations, and their influence becomes more important as the error rate [i.e. mutation rate] increases" (p.6886)[12]. Selective values are themselves given by the replication rate of the particular sequence. The quasi-species view thus casts doubt on the assumption of optimisation in some narrow sense, since:

> *"the target of selection is not the single advantageous copy, that is, an individual wild type [as was assumed in previous models]. Instead, it is the total quasi-species distribution that is the target, and this is not of some simple narrow-banded Poissonian type, but rather shows a fairly large nonsymmetrical dispersion with protrusions reaching far into the sequence space" (Eigen et al., 1989), p231*[13].

---

[11] For example, (Woodcock & Higgs, 1996) use such an explanation to describe why an asexual population forms a 'cloud belt' *below* the summit of an optimum in a selective environment with only a single optimum.

[12] This is an excellent, qualitative description of the *effective fitness* measure employed by (Stephens et al., 1997). It also ties in with the mutation-selection balance that holds the quasi-species together.

[13] Whether this should be taken as contributing to the debate on *group selection* is unclear. Charlesworth comments that such concerns (i.e. about group selection) may have stifled work on the evolution of recombination modelled in terms of its effects on the genetic variance *of a population* (Charlesworth, 1993).

2.5.2 Evolutionary Algorithms.

EAs differ from the more traditional models in computational biology in that rather than evaluating the temporal 'evolution' of a (set of) deterministic or stochastic equations that describe population behaviour, for example, the EA approach makes use of *actual* populations of individual entities. So rather than evolving (in the mathematical sense) a single quantity that corresponds to a probability distribution representing a feature of a population, the probability distribution arises from the distribution of traits in an *actual* population of discrete entities. One way EA models may be of benefit to theoretical biology, then, is to demonstrate that mathematically modeled probability distributions accurately reflect observable distributions of 'real' simulated populations, all other model parameters being equal. To date, the take up of EC models in computational biology has been slight (for example, a recent textbook on computational modeling (Brown & Rothery, 1993) makes no mention of any form of EC (and only briefly mentions neural network models, another form of soft computing that may in the future be of considerable worth to biological theorists)). Artificial life models of evolving ecologies, in which a population of individuals are introduced into a world where they may breed according to their success at performing some task in the world offer another form of simulation environment. In these models, population size is not fixed and very rich coevolutionary dynamics between an evolving population and its environment may be simulated.

The mathematical foundation of genetic algorithms is provided by the *Schema Theorem* (Holland, 1992; Goldberg, 1989). Although frequently criticised, (for example, (Vose, 1993a; Mühlenbein, 1995)), it remains a topic of current research for Goldberg's research group, and elsewhere (for example, (Stephens & Waelbroeck, 1996)). The theorem describes the proportional growth of a particular schema in an evolving population. *Schema* are essentially string templates defined over 0, 1, #, where the hash represents a don't care state. So for example, the schema 0#10# would match the strings 00100, 00101, 01100 and 01101. The Schema Theorem gives a lower bound on the proportion of individuals in a population undergoing proportional selection bearing a particular schema, given the proportion and mean fitness of individuals carrying the schema relative to the mean fitness of the whole population, and the disruptive effects on the schema of crossover and mutation. The theorem thus describes how schema with above average fitness increase their representation in the population over time; consequently, the mean fitness of the population should increase over time. This does not mean that an optimal schema will maintain a universal representation within a population, however, *even when the population has*

supposedly converged to an optimum containing that schema. This is a consequence of the mutation-selection balance identified above, and described in terms of the Schema Theorem by (Jennison & Sheehan, 1995), with reference to a simple evaluation function allowing only a single beneficial schema[14].

A second approach to predicting the behaviour of GAs are the exact, Markovian models, reviewed and further developed in (Vose, 1993b), for example, and extended by (Whitley, 1995) to incorporate individuals capable of learning.

The amount of variation in a gene pool has important consequences for the volume of the search space an evolving population can in general 'see'. Any particular population may be characterised by a *population defining schema*, (PDS) or *population schema*, which is the highest order schema (highest number of fixed positions) capable of describing every individual in the population; (in practice, the proportion of a particular allele at a particular locus taken over the population as a whole may have to exceed a threshold value (of a few per cent) before it usefully qualifies for the defining schema). The amount of variation may be measured as the *width* of the gene pool, which is simply the defining length of a genotype minus the order of the defining schema of the gene pool. So for example, for a population of individuals length $L = 4$ comprising {1101, 1001, 1011, 1000}, the PDS is $g_{\text{def}} = 1***$, order 1, giving a population width, W, of 4-1 =3. To exhaustively sample every possible individual that satisfies the PDS, the population size N must satisfy:

$$2^W \leq N \tag{2.1a}$$

$$\text{hence} \quad W \leq \frac{\log N}{\log 2} \tag{2.1b}$$

although this is not a sufficient condition since several individuals may be identical. Also note that a population of size 2 containing 2 complementary individuals is enough to generate *any* population defining schema. The probability that all (binary) alleles are present in a randomly generated population of size N and genotype length L is given by (Rowe & East, 1995) as:

---

[14] In particular, the representation of the schema within the equilibrium population was reduced for increasing mutation rate, and decreasing selective advantage of the schema.. Cf. the relationship between *selective advantage* and *copying fidelity* (i.e. mutation rate) on the *error threshold* above which the mutation-selection balance fails to maintain a quasi-species' integrity (Eigen et al., 1989)

$$p(\text{all alleles present for all binary loci}) = \left(\frac{2!S(N,2)}{2^N}\right)^L \tag{2.1c}$$

where $S(i, j)$ is the Stirling number of the second kind.

The number of possible populations is given by (Nix & Vose, 1991) as[15]:

$$\binom{N+2^L-1}{N} = \binom{N+2^L-1}{2^L-1} \tag{2.1d}$$

Whereas the population defining schema gives an idea as to the volume of search space that may in principle be sampled by the next generation, a second measure is required to provide information as to the 'centre of gravity' of the gene pool; the *consensus individual*, defined as the constructed individual for which each bit, b, is set at the majority allele taken over the whole of the population, suits this purpose admirably. In turn, this suggests an alternative measure of the population width, specifically the mean Hamming distance of each individual from the consensus individual. Whilst identifying the consensus individual, the bitwise convergence may be calculated by averaging the majority allele proportions over the length of the genotypic bitstring.

Now, the genetic operators act on a (selected) sample of the current population, which in turn will give rise to a *selected population defining schema* (SPDS). It is easy to see that if the width of this schema is narrower than that of the current population schema, selective sampling has reduced the amount of (hopefully less beneficial) variation in the population. An increase in the order of the population defining schema is also what tends to happen when a population begins to converge.

The PDS provides a simple way of describing the variation available within a population, although a more exact approach is to measure the proportion of each allele at each locus giving rise to a *bitwise-probabilistic population defining schema* (BP-PDS, and hence the corresponding structure for the selected population (BP-SPDS)). Although such a representation fails to capture the distribution of higher order schemata, it still provides the

---

[15] Recall that $\binom{x}{y} = \frac{x!}{y!(x-y)!}$.

basis for GA derivatives, such as the PBIL scheme of (Baluja, 1994), and the 'Compact Genetic Algorithm' of (Harik, Lobo & Goldberg, 1997). For binary loci, only a single probability is required at each locus and it specifies, by definition, the proportion of 1s in the population at that locus. So for example, for the population {1101, 1001, 1011, 1000}, the BP-SPDS is 1.0/0.25/0.25/0.75. The BP-SPDS provides the basis for the simulated crossover of (Syswerda, 1993) in which offspring are generated by deriving individuals as instances of that structure. The gene pool recombination scheme of (Mühlenbein & Voigt, 1995), derives *two* instances of the BP-SPDS, and these are then subjected to crossover in the normal way. In effect, this latter scheme is the limiting case of multi-parent recombination with uniform crossover, in which every selected individual gets the opportunity to contribute to any particular child.

2.5.3 Applying the One to the Other.

Generally speaking, there has been little transfer of mathematical techniques between the evolutionary biology and computation communities. One notable exception is the work of (Mühlenbein, 1995), who uses mathematical analyses of artificial and natural selection to develop his EA models. Quantitative genetics models that describe the evolution of population mean traits bear some resemblance to the ES approach. They may also be used to predict the mean behaviour of populations in a GA model; for example, in chapter 6, I demonstrate how quantitative genetics analyses of models of evolution in a sinusoidally varying selective environment predict the behaviour of population mean expression in a GA.

(Altenberg, 1995b) suggests the use of certain measurement functions that extract fitness and schema information from a population so that application of Price's Covariance and Selection Theorem (from the biology literature) becomes possible. Price's Theorem relates the expected value of some measure of the offspring from parental pairs in the current population to the population mean of that measure in the next generation by using the covariance between the actual parental and expected offspring fitnesses. By focusing on the distribution of fitnesses and schemata in a population, Altenberg is able to rederive the Schema Theorem from Price's Theorem, and also present an extension which demonstrates how schema processing is not in general useful, except when recombination is used as an operator.

Random genetic drift - the fixation of alleles through sampling error - is a major feature of evolution in GAs since they often utilise small populations. Drift in (Hinton & Nowlan,

1987), an important experiment reviewed in 2.7 below, is explained using diffusion equations (adapted for the haploid conditions used in that experiment) by (Harvey, 1993), who also cites a finite Markov chain analysis of drift by (Goldberg & Segrest, 1987). (Asoh & Muhlenbein, 1994) discuss drift arising from crossover in the absence of selection and mutation. In terms of wider model dynamics, an asexual version of (Hinton & Nowlan, 1987) is analysed in terms of classical population genetics by (Fontanari & Meir, 1990).

The information flow should not be limited to exploiting mathematical models from population genetics in order to address foundational issues concerning EAs, however. For example, (Collins & Jefferson, 1992) used a GA model to successfully extend an important analytic model of sexual selection by relaxing in turn several assumptions necessary for the tractability of the analysis. GAs are also particularly appropriate for work on finite populations, which is often analytically intractable, although some amount of theoretical, (van Nimwegen, Crutchfield & Mitchell, 1996), and numerical, (Woodcock & Higgs, 1996), work has been done in the biology. Artificial life models, such as the Latent Energy Environments (LEE) of (Menczer & Belew, 1996), provide a metric in terms of population size that may be compared to ecologically minded growth equations. In addition, the LEE model often allows the carrying capacity (i.e. the ideal equilibrium population size) to be calculated, and hence the overall adaptedness of evolved populations to be assessed. Another advantage of the simulated environment is that gene frequencies are available for study in addition to the population size metric. In addition to arguing for a synthesis of population biology and EC techniques, (Taylor, Jefferson, Turner & Goldman, 1988) demonstrate an early Alife modelling environment in which they simulate three ecological models, including the well known coevolution of 'rabbits and foxes'.

2.6 Adaptively Introducing Genotypic Variation

In EAs, genotypic variation may be introduced adaptively through the use of appropriately styled genetic operators. (Coyne & Paton, 1994) distinguishes between *exogenous* and *endogenous* adaptation. In exogenous adaptation, rates are modified according to current fitness measures; with the endogenous approach, an individual codes for its own mutation rate. The interested reader is directed to one of the recent reviews of the area which consider both adaptive crossover and adaptive mutation strategies (Smith & Fogarty, 1997; Hinterding, Michaelewicz & Eiben, 1997; Angeline, 1995), although an overview in the style of (Smith & Fogarty, 1997) is offered here. Three defining characteristics are suggested:

what is adapted (for example, the mutation rate, the crossover rate, or both)? what is the scope of the adapted component (e.g. the mutation rate at the level of specific genes, individuals or the population as a whole)? and what drives the adaptation (typically, self-adaptation, a deterministic temporal policy, or a response to a population performance measure)? Table 2-1 summarises the strategies used by various authors contemplating determinate static and dynamic optimisation problems under these categories. For the case of mutation rate, the most effective strategy appears to be the application of a high rate during the early stages of evolution (given an initially random population), with progressive 'cooling' of the mutation rate over time.

Author	*What?*		*Where?*			*How?*		
	Mutation	Xover	Gene	Ind'l	Pop	Self-adapt.	Temporal	Rating
(Bäck, 1992b)	√		√	√		√		
(Bäck & Schütz, 1996)	√			√		√	√	
(Cobb, 1990)	√				√			√
(Coyne & Paton, 1994)	√		√		√	√		√
(Davis, 1989)	√	√			√			√
(Fogarty, 1989)	√				√		√	
(Hart, 1995)	√				√			√
(Hirst, 1997c)	√			√		√		
(Julstrom, 1995)	√	√			√			√
(Spears, 1995)		√		√		√		
(Smith & Fogarty, 1996)	√	√	√			√		
(Srinivas & Patnaik, 1994)	√	√		√				√
(Tuson & Ross, 1996a)[16]	√	√		√		√		

Table 2-1: Summary table of the use of adaptive genetic operators in GAs.

In addition, the evolution of mutation rates has been considered in coevolution (Maley, 1997) and Alife models (Bedau & Seymour, 1995). Theoretical models of *optimal* mutation rates have been offered in both the EC literature (for example, (Bäck, 1992a; Mühlenbein, 1995)),

[16] In a separate work, (Tuson & Ross, 1996b) investigate Corne's COBRA scheme. This approach differs fundamentally from the self-adaptive strategies in that a ranked 'population' of operators is maintained, the rank position setting the probability with which a given operator is to be applied. Reranking is done on the extent to which offspring generated by an application of a given operator are better adapted than their parents. What this scheme amounts to is the co-evolution of a population of individuals and a population of operators.

and in theoretical biology (Gillespie, 1981; Eigen et al., 1989). Where the environment is non-stationary, the consensus appears to be that the optimal mutation rate is an increasing function of the rate at which the environment changes, a conclusion that was reached experimentally by (Cobb, 1990) in her GA investigation of sinusoidally fluctuating environments.

2.6.1 Self-Adaptive Operator Settings.

The adaptive mechanism for introducing genotypic variation that I shall employ throughout this thesis will be a *self-adaptive mutation rate*. A key component of evolution strategies (Bäck & Hoffmeister, 1994), self-adaptive mutation rates were first employed in GAs by (Bäck, 1992b). On each individual was a gene that specified the per bit mutation rate to be applied to *every* bit within that individual, figure 2-5a. Results on a standard test suite showed that an individual level self-adaptive mutation rate was viable, although problems with the binary representation made fine tuning of the evolved rate difficult (Bäck & Schutz, 1996). Theoretical considerations also suggest that allowing the evolved rate to apply to itself may result in instability (Weber, 1996). This may be countered by utilising a background mutation rate that maintains variation at a fixed rate within the rate gene, figure 2-5b. (Altenberg & Feldman, 1987) consider modifiers that are transmitted perfectly between generations in an attempt to identify evolutionarily stable values. The application of an externally set mutation rate to the evolved rate gene was followed in a later work, (Bäck & Schutz, 1996), where a time varying environment (specifically, a switching MaxOnes-MaxZeroes evaluation function[17]) was used to demonstrate the adaptive capabilities of the strategy. Comparison between the self-adaptive approach, a deterministic, time scheduled mutation rate and a fixed ('optimal') mutation rate on a range of static optimisation problems resulted in the deterministic schedule outperforming the self-adaptive regime, which in turn fared better than the fixed mutation rate model. An investigation of the evolution of self-

---

[17]MaxOnes (also known as the Bitcounting, Counting Ones, OneMax, or Degree of Match (to **1**, a target sequence of all 1s) function): an evaluation function much studied because of its amenability to analysis (e.g. (Bäck, 1992a; Bäck & Schutz, 1996) in their analysis of optimal mutation rates, (Mühlenbein & Schlierkamp-Voosen, 1995) following a quantitative genetics style of analysis). For an individual of length L bits, its evaluation is given by the normalised sum of 1s it contains $E(g) = \frac{1}{L}\sum_{i=0}^{L-1} g_i$ where g_i is the i'th locus of individual g.

adaptive rates in environments that change with varying periods has not, however, been reported. In chapter 6, I address this apparent gap in the literature.

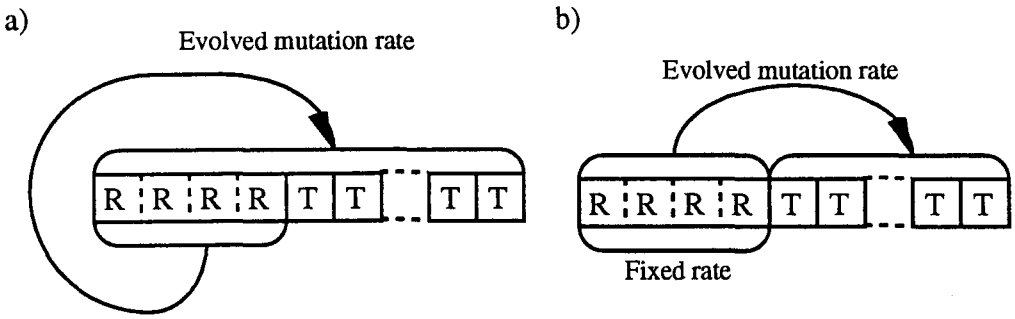

Figure 2-5: The self-adaptive evolution of mutation rates: a) applied over all bits; b) fixed, background rate applied to the rate gene, which specifies the mutation rate over all genes other than itself.

2.6.2 The Evolution of Modifier Genes.

In the biological literature, the self-adaptive rate genes are known as *modifier* or *regulatory genes* (Altenberg & Feldman, 1987; Charlesworth, 1993; Weber, 1996; Kondrashov, 1995). Modifier genes are genes that mediate the expression of transmission of other genes. Since a modifier has no direct phenotypic expression, it is selectively silent in terms of its effect on generation specific population means; that is, it induces a higher dimensional neutral network. The evolution of modifiers that affect the transmission of genetic material between generations, such as a mutation rate modifier gene, and the sense in which they are subject to selection, is qualitatively discussed for asexual organisms in (Weber, 1996). Several quantitative accounts of the evolution of transmission modifiers have also been developed (for example, (Altenberg & Feldman, 1987; Kondrashov, 1995; Charlesworth, 1993)).

In particular, (Altenberg & Feldman, 1987) offers a mathematical treatment for the diploid case that suggests the modifiers are subjected to selection under the *marginal fitnesses* that arise from the differential transmission of their associated expressed genes[18] - "we assume that the modifier alleles are perfectly transmitted regardless of the processes acting on the selected haplotypes [an individual gamete's genotype], so that the only force acting on the modifier locus is the *induced selection* resulting from its effect on the transmission of the

[18] The *marginal fitness* of a modifier is the mean selective value of all the individuals that bear it.

other loci" p562, my italics. Since modifiers are only indirectly acted on by selection, they effectively hitchhike on the back of adapted individuals.
More recently, and encouragingly from the EC community, (Stephens et al., 1997) has described the evolution of modifiers in terms of the *effective fitness* of individuals, which gives for each individual in the current generation its expected proportion of the population in the next. The effect of the modifier is *implicitly* accounted for under this measure.

The landscapes induced by these two models (i.e. the surfaces of marginal and effective fitness) are discussed further in 4.3.

2.6.3 Non-random Mutation.

Generally, mutation induced genetic variation will be *random.* That is, it is produced without reference to the 'preferences' of the selective environment. Recent work in experimental biology, however, has generated a lively debate in that literature about the possibility of *directed* or *adaptive mutation* (Symonds, 1994; Sniegowski & Lenski, 1995). Originally presented as a case for genetic loci which seemed to mutate to useful alleles in the presence of harsh selection conditions, *directed mutations* would seem to suggest a form of 'supervised genetic learning' in which only appropriate mutations are encouraged. More recent opinion tends to favour the presence of randomly hypermutable sites in certain critical genes (*adaptive mutation*). Whilst the former argument would be seen to support the Lamarckian position of environmentally acquired variation, the weaker adaptive mutation stance conforms to the neo-Darwinian requirement of randomly introduced variation.

2.7 The Effects of Adaptive Phenotypic Variation

Under most developmental maps, genetic variation will realise corresponding *phenotypic* variants, which are then acted on by selection. However, the introduction of adaptive phenotypic variation by *learning* may have a significant impact on the evolutionary dynamic as a consequence of the *Baldwin Effect* (Baldwin, 1896; Simpson, 1953; Anderson, 1997). Offered as "a sequential process in which acquired characteristics are replaced by genetic characters" (Simpson, 1953), the Baldwin Effect has been by and large ignored by evolutionary biologists. However, since its introduction into the EC literature by (Hinton &

Nowlan, 1987), it has been the focus of much interest within that domain, recently culminating in a complete reader of the area covering both old and new work, (Belew & Mitchell, 1996), and a special issue of the journal Evolutionary Computation, (Turney et al., 1996).

There are two main components in Baldwin's argument:

- *learning enhances survival* - by being better adapted, an individual better survives, and produces more offspring than a maladapted individual;

- *genetic assimilation of acquired traits* - that there are phylogenetic consequences of such behaviour is now argued as follows: through inheritance, variants of the parental population are born into the offspring population. If phenotypic variation similar to that acquired in the parental population arises from random genetic variation in the offspring population, and without the need for costly learning (i.e. the trait is expressed directly through a simple, deterministic developmental process rather than stumbled across by the rather more flexible, although expensive, learning process), then the child will be rather more adapted from birth, and more likely to prosper than even an individual who comes to learn the adapted trait during the course of its lifetime. Originally acquired characteristics thus apparently become heritable traits.

If adaptive plasticity is governed by a modifier gene that regulates the *expression* of other genes, (i.e. affects the *selection* component of the selection-transmission model of (Altenberg & Feldman, 1987), rather than the *transmission* component as discussed previously with respect to the genetic operators) the modifier may influence developmental plasticity in one of two ways: deterministically, or probabilistically. I shall illustrate this distinction with two concrete examples.

- In the first case, imagine a modifier which flags whether within lifetime search to a locally optimal phenotypic state is to be pursued, as in the memetic algorithms of (Radcliffe & Surry, 1994). All individuals within the basin of attraction of a given local optimum will attain that optimum. There is thus no way in which selection can act so as to distinguish between individuals closer to, or further away from the optimum, as long as they are all in its basin of attraction. That is, there is no variance in marginal fitness between individuals carrying the set flag in the local basin of attraction. The only way to distinguish between individuals is by applying some sort of cost function. Reasoned

argument by (Mayley, 1996a; Mayley, 1996b), and support from a series of GA experiments he conducted, suggests that for genetic assimilation to occur, exactly this sort of cost to fitness (as evaluation) must be associated with learning. Once an adapted phenotypic state has been reliably discovered by adaptive individuals, then the adaptive trait may be selected against in favour of individuals that directly inherit the adapted state because the learning process is itself costly.

(Mayley, 1996a) discusses the implementation of learning costs with respect to how an individual is evaluated. In *posthumous assessment*, an individual is evaluated following learning; a cost may then be exacted on the amount of learning used in attaining that particular pre-costed evaluation. A second approach is to evaluate an individual continually (*continual assessment*) - the evaluations following each learning trial are summed, and an average evaluation produced when learning is complete. Note that the grain of the continual evaluation assessment function may be varied; for example, in a coarse grain assessment, the final evaluation may be the mean of two evaluations measured before and after learning. A rather more comprehensive assessment of costs that may be incurred is given by (Mayley, 1996b). In his analytic treatment of learning, (Anderson, 1995) discusses cost in terms of the "loss of reproductive effort".

- For the probabilistic case, let the modifier flag whether probabilistic learning is to take place or not. In this case, imagine that each individual is allowed a certain number of learning trials, each of which corresponds to flipping a single bit of its inherited state. Selection is based on the best evaluation returned from the set of learning trials. Now, for two similar individuals, if the number of learning trials is low, it is likely that the individuals will sample different acquired states and as a result return different evaluations. A variance of marginal fitnesses results and the modifier gene *is* subject to 'induced selection' without the need for an explicit cost. To further increase the selection pressure against redundant plasticity, a simple explicit cost that is proportional to the number of learning trials may be applied.

Hinton and Nowlan's key paper demonstrating the Baldwin Effect in a GA, (Hinton & Nowlan, 1987), implicitly incorporated both mechanisms. Their experiment was designed as follows: for a population size of 1000, under proportional selection with crossover rate 1.0 and no mutation, each individual was represented by a string of 20 loci, each of which could take on the value **0**, **1**, or **?** with initial probabilities 0.25, 0.25, 0.5 respectively. The task was to match a target sequence of all **1**s. For individuals containing one or more **?**s, random (probabilistic) 'learning' was allowed. A learning trial comprised of setting all the **?**s in a

particular individual to **1** or **0**, each with probability 1/2, with a maximum of 1000 learning trials per individual[19]. Learning stopped immediately if the target sequence was attained. Individual evaluations were subjected to an explicit learning cost:

$$f_{HN} = 1 + \frac{19(1000 - n)}{1000} \tag{2.2}$$

where n was the number of learning trials required to find the target sequence.

Individuals containing one or more **0**s necessarily achieved the minimum evaluation of 1. Analysis of this evaluation function by (Fontanari, 1990; Harvey, 1993) gave the expected *fitness* of an individual bearing q ?s and 20-q **1**s as:

$$f_{HN,e}(q) = 20 - \frac{19\left(1 - r^{1000}\right)}{1000(1 - r)} \tag{2.3a}$$

where r is the probability of failing to match the target sequence on any particular trial:

$$r = 1 - \left(\frac{1}{2}\right)^q \tag{2.3b}$$

The cost of learning may be tuned by introducing an explicit cost factor. So for example:

$$f'_{HN}(q) = u(q).f_{HN}(q) \tag{2.4a}$$

where:

$$u(q) = \frac{1000}{1000 + 2^{q+k}} \tag{2.4b}$$

---

[19] If in the first generation each individual is allocated the maximum number of trials, then $1000.1000 \approx 10^{20}$ trials are required which corresponds closely to the size of the whole search space. According to (Harvey, 1993), the probability of *not* finding the sequence given a random population is ~0.572. In terms of just locating the optimum, an exhaustive search utilising a similar number of trials is rather more likely to succeed.

and k is used to tune the *explicit* cost of learning[20]. Equation (2.4a) is plotted in figure 2-6. The resulting curve shows how the introduction of learning is capable of transforming the surface of expected evaluations over individuals away from its original needle-in-a-haystack appearance. This 'smoothing' of the 'fitness landscape' is a central metaphor in of the influence of learning on evolution, and I shall clarify and develop it throughout this thesis, theoretically in chapter 4, experimentally in chapter 5. (Whitley, Gordon & Mathias, 1994b) showed exactly how a partially deceptive evaluation function may be smoothed by learning. In addition, they considered a Lamarckian inheritance strategy in which states acquired by learning in one generation could be directly transmitted to the next.

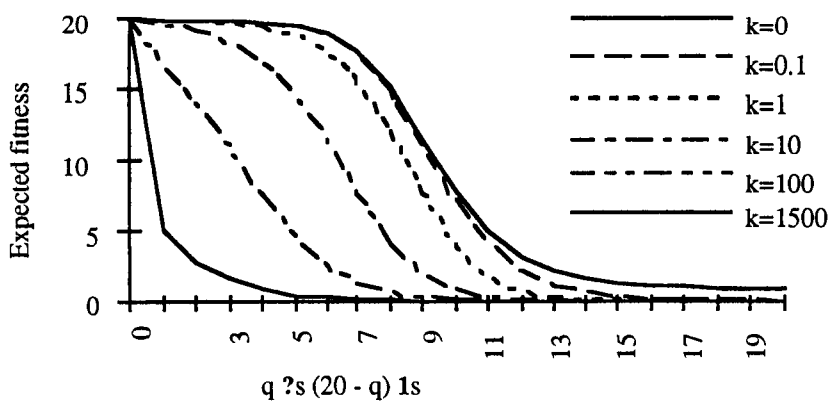

Figure 2-6: Expected fitness for the explicitly costed evaluation function of (Hinton & Nowlan, 1987), based on the analysis of (Fontanari & Meir, 1990; Harvey, 1993).

The behaviour of Hinton and Nowlan's original computational model has been well discussed in (Belew, 1989; Belew, 1990; Harvey, 1993) as well as in the original paper. In brief, in the absence of mutation the proportion of **0**s in the population quickly falls to zero (around generation 10) whereas the proportion of **1**s and **?**s tends over several hundred generations to a limit dependent on drift and founder effects, though typically in the ratio 85%/15% give or take the odd 10%.

---

[20] A related family of curves may be generated by modifying the maximum allowable number of learning trials, T, which gives an expected evaluation of $f_{HN'e}(q) = 20 - \frac{19(1-r^T)}{T(1-r)}$. Note, however, that in this case the minimum evaluation limits at 1, whereas in the case of increasing cost the asymptotic minimum value tends to zero.

Using a derived model, (Turney, 1996) investigates further appropriate time dependent schedules for the amount of plasticity applied. As with the evolution of mutation rates, good strategies appeared to require large amounts of learning in the early stages of search from a random population, with a progressive reduction over time. Again, though, the behaviour of such strategies in a well defined temporally fluctuating selective environment has not been studied. I address this point in chapter 7, where I implement a GA using a self-adaptive learning rate and a nonstationary evaluation function.

If mutation is introduced into the Hinton and Nowlan model, the **0** allele maintains a token representation in the population[21]. (Belew, 1990) plots a series of population mean evaluation curves for various mutation rates. Low mutation rates (10^{-5}, 10^{-3} per bit) afford convergence at a high mean evaluation, a rate of 0.01 per bit convergence at a significantly lower level, and for a rate of 0.2 per bit the population mean evaluation never exceeds its initial minimum value.

Complementing the experimental approach, (Fontanari & Meir, 1990) offer a classical genetics analysis of an asexual variant of the experiment using an infinite population and given in terms of the proportion, p, of **1**s in the population. Without learning or mutation (i.e. only selection tempers the growth rate of particular genotypes) there is a stable equilibrium at $p = 1$ and an unstable one at $p = 0$. The time to fixation of the stable equilibrium is dependent on the initial value of p, although once p reaches 0.6 it rapidly attains unity. For the original initial proportions (0.25 **1/0**; 0.5 **?**), time to equilibrium is of the order of 10^8 generations. When mutation is introduced at a low rate (in fact, greater than $\sim 4.35 \times 10^{-4}$), the final equilibrium value of p becomes less than one, although fixation time is reduced, and the unstable equilibrium at $p = 0$ is replaced by a stable one at $p \approx 0.5$; only this equilibrium remains for increasing mutation rate. The diffusion analysis of (Harvey, 1993), which does *not* rely on an infinite population, reveals that for any *particular* locus, a

---

[21] Naive codings of the original H&N experiment may assign to each loci 2 bits to represent the three alleles, rather than using single trits (and structured genetic operators) for each locus, with possible consequences for bias as discussed in a previous footnote (note 7, p. 19). (Belew, 1990) encountered such a "curious result" using an encoding complementary to the first suggested in that footnote in which mutation of the plastic allele realised either the other plastic allele or a **0**, hence increasing the selection pressure *against* **?**s where mutation was allowed. Using a population of initially 50% each of **1**s and **?**s, and no mutation, by the second generation a significant proportion of **0**s had appeared - crossover between 01- **1** and 10 - **?** alleles produced the 00s responsible.

50/50 balance between **1** and **0** alleles is attained for a mutation rate greater than 0.5 times the reciprocal of the population size, and fixation at either one value or the other for lower rates of mutation (with almost equal likelihood in the case of weak selection). With learning and no mutation, both authors derive the expression given above. With the introduction of mutation into the learning regime, the ratio of **1**s to **?**s decreases for increasing rate and the equilibrium representation of **0**s is non-zero.

In addition to introducing mutation, (Belew, 1990) extended the original model by weighting the probability with which an individual learns the correct allele setting. In his *lineage* model, individuals which attain the target sequence are able to influence the plastic development of their offspring through a parameter CA, the 'cultural advantage'. Culturally advantaged offspring attain the correct setting ($0.5 + CA$) of the time. Experiments with CA set to 0.1 show how the introduction of culture reduces the degree to which genetic assimilation of the target state occurs - cultural inheritance reduces the costs associated with incorrect learning and the selection pressure against **?**s is reduced. As a consequence, early convergence of **?**s may result in a lower asymptotic mean fitness than for populations not receiving cultural benefits, although this final level is reached in rather fewer generations. In a second *broadcast* model, a cultural advantage is afforded to B randomly chosen members of the next generation. For small values of B, this mode of transmission may perform badly since in initial generations at least it is likely that if only a few random individuals are selected, they will carry **0**s and hence be unable to attain the target state. For larger values of B, (Belew takes $B = 16$) performance improves over the lineage model. Belew also argues that as B increases, the 'fidelity' of the broadcast message may be affected for the worse, and he models this by reducing CA (although it is still positive). Nevertheless, this assumes that erroneous information is not transmitted, a process that may be modeled by setting $CA < 0$, giving worse than evens odds that the correct allele will be set. The behaviour of such a system is investigated in 5.3.2, where I present a model that *does* transmit maladaptive information between generations.

The use of adaptive operators by a fixed *proportion* of the whole population, rather than every individual, has been considered in static environments for the case of learning by (Hart, 1994), and for IAC by (Houck, Joines, Kay & Wilson, 1997). In each case, *partial* application of the operator was found to be the best approach.

2.8 Evolution with Fixed and Adaptive Strategies in Fluctuating Environments

Given the ability of adaptive plasticity to generate phenotypic variation within the lifetime of an individual, it would seem reasonable to expect that such strategies should have been investigated with respect to *temporally varying* evaluation functions. Whilst adaptive plasticity has been investigated using static evaluation functions (see above), and in Alife environments using artificial neural networks (Ackley & Littman, 1991; French & Messinger, 1994; Nolfi, Elman & Parisi, 1994; Nolfi & Parisi, 1995; Parisi & Nolfi, 1996; Belew, McInerney & Schraudolph, 1991; Menczer, 1994), there has been little work of any kind on well-defined, dynamic evaluation functions.

The way in which a well defined environment fluctuates may be classed in several ways:

- (Cobb, 1990) distinguishes several possible dynamics; for example, periods of stasis may be interrupted by cataclysmic change (discontinuous nonstationarity), or the environment may be in a continuous state of flux (consider an optimum state described by square and sine waves respectively). In addition, the effects of the change - the relationship between environmental states before and after a state change - may either be random (a Markovian Switching Environment, (MSE)) or dependent on either previous states or some other parameter, such as a time counter (State Dependent Nonstationary Environment (SDNE));

- In the population genetics literature, four classes of temporally varying environment are commonly identified: constant (stationary) optima, then directional (typically moving with a fixed velocity, such as $p_{opt}(t+1)=p_{opt}(t)+1$), cyclic (all of which represent SDNEs); or random change (an MSE) (Lande & Shannon, 1996; Charlesworth, 1993);

- To more completely characterise a fluctuating environment, the range of states the optima pass through, and over what period, must also be specified. Typically, a system may be bistable, cycle through a bounded range of values or pass through an infinite array of states without repetition (Feldman et al., 1996). For the 'randomly varying' environments, the autocorrelation of the optima over time gives a useful measure of the (un)predictability of future states.

Although characterised in terms of the dynamics of some optimum sequence, one should bear in the mind the likely dynamics of the evaluations received by any arbitrary individual over time. For example, imagine a population set the task of tracking a sinusoidally varying optimum. An *error* value representing the distance between the individual's expressed phenotype and the current optimum is used to evaluate each individual. With reference to figure 2-7, note how the change in evaluation over the individual with phenotype 0 has a period twice that of the optimum; the error for the individual with the phenotype 1 is a sinusoid 180 degrees out of phase with the optimum.

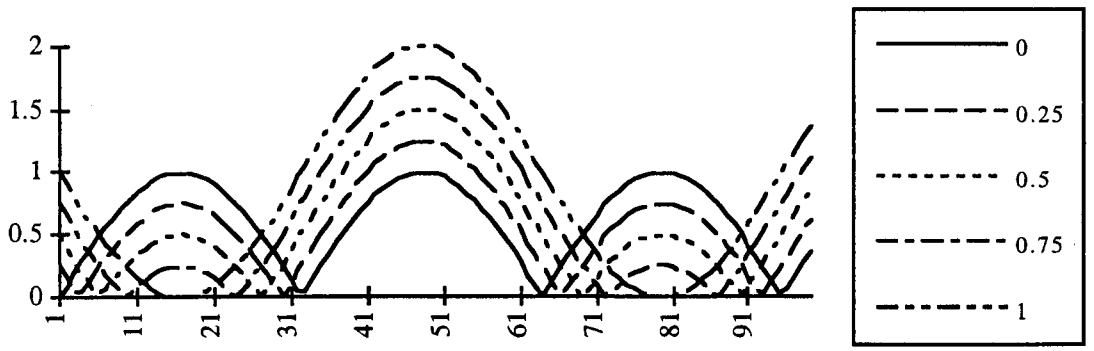

Figure 2-7: The change in error (i.e. evaluation in a minimisation *problem), calculated as the absolute distance from an optimum that fluctuates sinusoidally between +1 and -1 with period 63 generations, for 5 individuals (reals in the range 0..1).*

Predicting later work, (Cobb, 1990) distinguished two ways in which a population may cope with a dynamic environment: either by increasing its memory of previous events in the hope they will reoccur; or alternatively, by introducing variation in the population so as to cover more potential solutions. In his study of a dynamic environment switching between 2 states, (Goldberg & Smith, 1987) uses a diploid system which is capable of storing information about *both* possible states, and (Dasgupta & McGregor, 1992; Dasgupta, 1994) uses a switched hierarchy (gene activation model) to concurrently evolve multiple possible solutions on each individual. (Cobb, 1990; Grefenstette, 1992; Cobb & Grefenstette, 1993; Coyne & Paton, 1994) take the alternative approach of introducing literal genetic variation through a variety of mechanisms including hypermutation and the immigration of random individuals. (Bäck & Schütz, 1996) demonstrated the viability of a self-adaptive mutation rate in bistable environment but did little more. The common feature of all these methods is their concentration on the ability to maintain an appropriate repertoire of expressible *phenotypes* through *genotypic* variation. That is, the 'width' of the gene pool (2.5.2) is large enough to

source the various phenotypic variants required to cope with the range of presented environments. (Vavak, Fogarty & Jukes, 1996) utilise an adaptive developmental map to model a degree of phenotypic local search. Finally, (Behera & Nanjundiah, 1996) , employ plasticity in the form of simple learning to increase the amount of 'virtual' phenotypic variation in the population. The setting and the various strategies reported within these papers is given in table 2-2.

Author	Dominant Source of Variation	Evaluation Function
(Behera & Nanjundiah, 1996)	Adaptive plasticity	Bistable unitation
(Cobb, 1990)	Adaptive hypermutation	Sinusoidal 32 bit real
(Cobb & Grefenstette, 1993)	Adaptive hypermutation Random immigrants	Transforming one predefined hilly landscape into another
(Coyne & Paton, 1994)	Self-adaptive, switched hypermutation Adaptive hypermutation	Random knapsack
(Dasgupta & McGregor, 1992)	Homeotic (switched) genes	Bistable knapsack
(Dasgupta, 1994)	Homeotic (switched) genes	Sinusoidal 32 bit real
(Goldberg & Smith, 1987)	Diploidy	Bistable knapsack
(Grefenstette, 1992)	Adaptive hypermutation	Transforming one predefined hilly landscape into another
(Vavak et al., 1996)	Adaptive maturation	Smoothly changing real

Table 2-2: Summary table of EC papers reporting dynamic evaluation functions.

What little interest there has been in the EC literature for well defined, temporally varying evaluation functions has tended to concentrate on the evaluation means of the best individual in the evolving populations. Since one of the motivating factors for such studies is to see whether GAs are suitable for online control where a degree of 'value tracking' is required, this focusing on performance scores and on how well the algorithms perform is not wholly surprising: *global optimisation at every instant* is, after all, the ultimate, if unattainable, goal.

However, from an evolutionary biology point of view, it is typically the population dynamics that are of interest. Using a quantitative genetics approach, (Lande & Shannon, 1996) and (Charlesworth, 1993) offer expressions for the relationship between the population mean phenotype and the optimum in a range of fluctuating environments (stationary, constant velocity, random, sinusoidal). In particular, in a moving environment, the population mean tracks, but lags, the optimum. In a sinusoidally varying environment, the population mean additionally represents an attenuated form of the optimum. In chapter 6, I show how closer

inspection of such results may offer some important insight into the ontology of the evolutionary process, and also demonstrate the extent to which the theory predicts the behaviour of a GA model.

(Anderson, 1995) furthers the analytic approach by accounting for the notion of plasticity by a reduction in selection strength. By introducing learning, in a steadily moving environment, the population mean lags further than in the case without learning. Increasing the mutation rate has only a 'modest' effect on this lag. Through a numerical simulation, using costed learning, he is also able to show that increasing amounts of learning are supported for increasingly dynamic environments. I provide further results on the evolution of plastic individuals in a sinusoidally varying environment in chapter 7.

The biological literature has yet more to offer the EC community however. An increasing body of analytic work is concerned with the trade-off between the various ways, and the various *rates*, in which phenotypic variation is introduced into a population (Boyd & Richerson, 1983; Boyd & Richerson, 1988; Feldman et al., 1996; Lachmann & Jablonka, 1996; Jablonka, Oborny, Molnar, Kisdi, Hofbauer & Czaran, 1995). Typically, three inheritance strategies are compared (Boyd & Richerson, 1983):

- *simple genetic transmission* (the offspring inherits the parent's initial phenotype);
- *cultural transmission* (or *social learning*; the offspring inherits the parent's final phenotype, and may go on to 'learn' another);
- *pure individual learning* (the offspring's initial phenotype is derived from a probability distribution of phenotypes representing the outcome of an effective learning rule within the current environment).

Essentially, what is at stake is an understanding of the suitability of mutation only, individual learning and social learning (or cultural inheritance) strategies in environments of varying temporal grain. All the theoretical models describe a qualitatively similar result, so I shall concentrate on reviewing one of the more recent models.

(Feldman et al., 1996) investigated the coincident evolution of both phenotypic and genotypic qualities by considering *phenogenotypes*. Modeling asexual, haploid individuals in a between generation, fluctuating environment with a tunable period, interspersed with periods

of stasis (no environmental change), two modes of learning are discussed - individual learning, which always finds the optimal setting for a trait in the current environment, though subject to a fixed cost, and social learning (also (independently) costed) in which an individual inherits the behaviour of a random individual from parental generation (i.e. oblique transmission). This gives rise to three possible phenogenotypes - IL (individual learners), SLC and SLW (social learners correct and wrong respectively). Two classes of environmental variation are considered:

- in an *infinite state environment*, (no single environmental state reoccurs), where the environment changed state between every generation, there is a single equilibrium corresponding to a converged population of individual learners. For constant environment, two stable equilibria (in terms of population composition) are possible, depending on the relation between social and learning costs. Where the individual learning cost is the greater (as one might reasonably expect) then social learners will take over the population. Conversely, where individual learning is the cheaper of the two in terms of fitness cost, and given a small proportion of social learners in the initial population, the population will converge solely to individual learners. Also considered were situations where after each environmental change, l - 1 generations of stasis followed. For the alternating case $l = 2$ (one period of change, followed by one of stasis, followed by one of change,...) convergence to a population of individual learners is the only stable equilibrium. As l increases, then stable equilibria of individual and social learners becomes possible if the ratio of costs is favourable. Generally, for fixed appropriate cost settings and increasing l, the proportion of individual learners at equilibrium decreases. For a given l, all other things being fixed, increasing the individual cost or decreasing the social cost *reduces* the equilibrium mean frequency of individual learners. Extending the intermittently fluctuating environment model to a more general (aperiodic) one in which environmental change occurs with a probability u, stasis otherwise, numerical models show that for small u (loosely, large l, or significant autocorrelation) the proportion of individual learners will be low, whereas for higher values of u (i.e. increasingly dynamic environments) the frequency of individual learners in the equilibrium populations will be high;

- in a *two state model*, (only two possible environmental states), the situation is now rather more complex, since social learners who learn from a previously 'incorrect' individual will display the optimal phenotype SLC in the current generation. For $l = 1$, (environmental change every generation), or $l = 2$, there are no stable equilibrium

populations containing social learners. For cases of $l > 2$, the equilibrium populations are polymorphous, containing both social and individual learners. Numerical methods reveal that evolutionarily stable populations can be identified for a parameter, L, that regulates the probability with which an individual will display individual (as opposed to social) learning. Generally, for increasing l, L decreases, all other things being equal. A relative increase in the social learning cost (always less than the individual learning cost) increases the value of l for which the polymorphous equilibria are stable.

These results are similar to the conclusions drawn earlier by (Boyd & Richerson, 1983). However, Boyd and Richerson also considered a simple mutation only strategy. This was shown to be the best adapted transmission strategy in slow moving environments, where the population could evolve to an optimum under its own steam and without the need for costly learning of any sort. These conclusions are crudely summarised by figure 2-8. In chapter 7, I present a series of GA experiments in which some degree of experimental support for the theoretical models described above is provided.

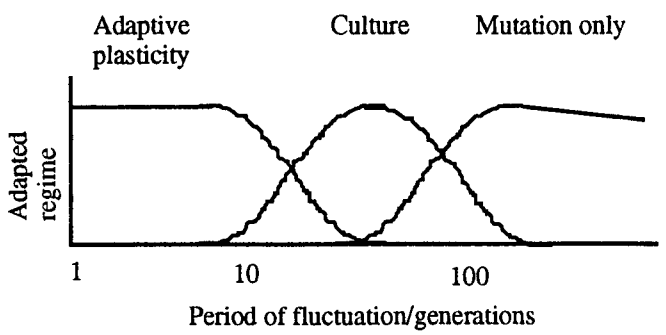

Figure 2-8: The appropriate choice of strategy for a fluctuating environment of a given period.

Finally, (Levins, 1968), (abridged in (Roughgarden, 1979)) offers a broad analysis of evolution in dynamic environments in terms of game theoretic evolutionarily stable strategies (ESS). Two discrete environments are considered, with an identical range of phenotypes in each, along with the corresponding range of fitnesses. For each phenotype, its fitness in the second environment may be plotted against its fitness in the first to give a fitness set[22]. Where the environments are relatively alike (i.e. the fitness curves over the range of

[22](Maynard Smith, 1978), p35, goes one further by subdividing the fitness set into a phenotype set for each environment, and the fitnesses then associated with the individual phenotypes. In this way, he is able to account for frequency dependent selection.

phenotypes are not wholly dissimilar), a *convex* fitness set, figure 2-9a, results; for widely differing environments, a *concave* curve, figure 2-9b, is the outcome.

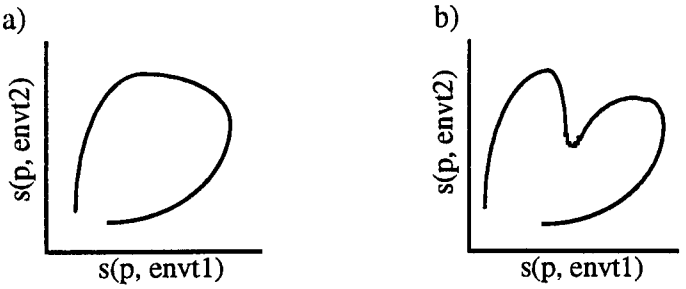

Figure 2-9: stable strategies for variable phenotypic expression across two distinct environments; the axes represent selective values of continuous traits in the respective environment: a) convex fitness set; b) concave fitness set (after (Levins, 1968)).

By identifying the optimal individual in each environment, two further points may be plotted corresponding to the extreme strategies in which each solution is 'monomorphous' in turn. Now, to identify the ESS, an optimisation criterion must be identified. In fine grain environments (rapid fluctuations) classical theory holds that the arithmetic mean fitness ($c_1w_1 + c_2w_2$) is optimised, where w_i is the population mean fitness of phenotype i, and c_j is the proportion of time spent in environment j; for coarse grain environments, the geometric mean fitness $w_1^{c_1}w_2^{c_2}$ is optimised. It turns out that for convex fitness sets (that is, similar environments) the equilibrium population will be converged on an intermediate phenotype (i.e. one that copes equally with each environment). For concave fitness sets, however, two outcomes are possible: firstly, in the fine grained environment, the population converges on the optimal phenotype in the most frequently occurring environment; secondly, for the coarse grained environments, a polymorphous population supporting both optimal phenotypes is to be expected.

2.9 Where now?

In 2.3 I described and extended an exemplar model of evolution from the biological literature and used it to identify the various key features of an evolutionary system. Evolutionary algorithms and quantitative genetics models were introduced in 2.4 within this wider context and in relation to each other. In both sections, however, a considerable amount of

terminological and interpretive confusion was identified. The framework presented in chapter 3 is intended to address this issue and provide a canonical model of evolution. Consequences of the model are described theoretically in chapter 4 and by experiment in chapter 5.

The ways in which variation may be sourced adaptively at both phenotypic and genotypic levels has been reviewed, and a theoretical trade-off identified between mutation, plasticity and cultural inheritance according to the rate of environmental change. EC work to date has concentrated on the evolution of genetic operator application rates in static and coevolutionary (ecological) models, but there has been little principled study in well-defined nonstationary environments. In chapters 6 and 7, I use one such test environment to rectify this gap in the literature.

Chapter 3 - A Framework for Describing Evolutionary Algorithms

"The universe (which others call the Library) is composed of an indefinite and perhaps infinite number of [connected] hexagonal galleries...[It] is a sphere whose exact centre is any one of its hexagons and whose circumference is inaccessible...All the books are made up of the same elements...[and] there are no two identical books. From these two incontrovertible premises [one] deduced that the Library is total and that its shelves register all the possible combinations of the twenty odd orthographical symbols (a number which, though extremely vast, is not infinite)."

JL Borges, from "The Library of Babel" in Labyrinths.

The grand, old Duke of York,
He had ten thousand men,
He marched them up to the top of the hill,
And he marched them down again.
And when they were up they were up,
And when they were down they were down,
And when they were only half way up,
They were neither up nor down.

Traditional.

3.1 Introduction

The landscape metaphor has achieved an almost universal status in the evolutionary computation (EC) community as an aid to understanding problem structure and the operation of evolutionary algorithms (EAs) over those structures, although I shall argue here and in chapter 4 that it is often applied inappropriately.

A good understanding of the nature of the valuation landscape - phenotypic evaluations viewed over the genotypic search space - is central to this thesis, although there is much that can be gleaned from simply observing the action of selection on phenotypes. For example, the success of quantitative genetics, or the conceptualised motion of an evolving population of phenotypes through an *evaluation landscape* which depicts received evaluations over a *phenotypic* search space. Section 3.2 develops the landscape metaphor solely as it relates to *evaluations* and the apparent (phenotypic) targets of selection. The selection and evaluation functions that give rise to performance landscapes are described in section 3.2.1. The ontology and microstructure of the possible search spaces underlying these surfaces will then

be discussed (3.2.2) and a clarification of the notion of landscape *ruggedness* offered (3.2.3).

In order to describe the evolutionary dynamic through an appropriately structured *genotypic* search space, however, a clear notion of the valuation landscape is essential. In section 3.3, I explore the notion of development from genotypic representations to phenotypic expressions: section 3.3.1 elaborates on the subtle distinction between valuation and evaluation, a distinction that will be later exploited in chapter 4; in section 3.3.2, I describe the range of possible developmental maps, including those which incorporate within generation local search; and in section 3.3.3, I discuss the logical status of these various maps. Finally, in section 3.4, I use the search space idea, in conjunction with the landscape metaphor, to describe the evolutionary dynamics of individuals within an evolving population in both static and dynamic selective environments.

3.2 Essential Components of Evolutionary Systems

In this part, I shall introduce a framework that may be used to describe the evolution of a population through some search space according to some evaluation function and under a given selective regime. Knowledge of four functional components of the system are essential if a characterisation of the evolutionary dynamics of a population is to be described: the operator induced structure of the genotypic search space, the nature of the developmental map, the behaviour of the evaluation function and the type of selection.

This approach - of understanding the evolutionary process as a whole - will provide the context for making an informed choice of operators and algorithm parameters, given a problem to be tackled by evolutionary methods. In addition, it presents a single, consistent metaphor for a mind's-eye visualisation of the evolutionary dynamic of a population: search over a *valuation landscape* which represents the received evaluations of individual genotypes, visualised over a *genotypic* search space. In maximisation problems, the search process is geared to climbing valuation 'hills' (i.e. attaining population mean valuation); in minimisation problems, the idea is to descend into the deepest 'valleys'.

3.2.1 Introducing the Notion of *Landscapes*.

Virtually everyone in the EC community is familiar with the notion of a *fitness landscape*, conceptualised as a surface of individual fitnesses plotted over binary genotypes one Hamming unit apart. Originally proposed by Sewall Wright (Wright, 1988) there is historically a certain degree of confusion as to the actual nature of the space over which the landscape lies (Provine, 1986), notwithstanding the conceptual problems surrounding the notion of fitness itself (Byerly & Michod, 1991). The height of the surface at a particular point represented either the fitness of the individual specified by that point, or the mean fitness of the population. The former representation of the search space presented a gross simplification of the actual, high dimensional space, typically by assuming that neighbouring individuals were separated by a single mutation. The latter view required a search space describing all possible gene frequencies, in which each point gave the gene frequencies over a population. As Wright points out, the device is a "pictorial representation" invented to illustrate his Shifting Balance theory of evolution, that is "useless for mathematical purposes". However, it does present an intuitive view of the relative fitness values between genetic neighbours and is frequently used today by the EC community as a measure of adaptedness over a real or imaginary space.

A simpler view of landscapes has recently been suggested by (Gavrilets, 1997), in the guise of *holey landscapes*. Rather than representing a continuous range of fitness values, the holey landscape simply denotes whether or not an individual is viable or not. Whilst microevolutionary dynamics may be explained by appealing to 'local evolution' and the scaling of locally optimal fitness peaks, macroevolutionary trends (such as speciation) are better understood with reference to effectively neutral evolution over the holey landscape. This view may be utilised in understanding constrained optimisation problems in EC. For example, where hard constraints exist, individuals that violate these constraints are by definition unsatisfactory (i.e. not viable).

Typically, there are three assumptions critical to the sensible application of the (Wrightian) landscape metaphor that are often ignored.

- First, the landscape idea is typically used to help visualise the motion of a population as it evolves over a stationary *evaluation* landscape, where individual evaluations are taken to be independent of each other. However, one should also bear in mind the *selective*

surface which represents the evaluation landscape over a population transformed by the particular selection function being used.

- Second is the identity of the search space - Provine (Provine, 1986) points out that Wright is not consistent in his view of what the search space underlying the landscape corresponds to. For example, at some times it is taken to be the space of all individual genotypes, at others to be the gene frequencies in the population.

- The third assumption relates to the underlying structure of the hyperplane over which the surface is plotted, that is, what connects neighbouring points? Wright himself argued that the multidimensionality of the population space "should be explicitly assumed to underlie the two-dimensional one of the diagram and that its origin should be at whatever peak was under consideration" ((Wright, 1988) p120). In 3.2.2.3, I offer a pragmatic simplification to a visualisation of the whole search space, which concentrates on the microstructure of the search space in terms of operator induced search space neighbourhoods (SSNs). The view is further complicated by considering developmental (genotype-phenotype) mappings and crossover. Evaluation is usually applied to phenotypes, but the evolutionary operators of recombination act at the genotypic level. One must be clear to distinguish, then, between the evaluation surface plotted over the phenotypic representation space, and the resulting *valuation* surface defined over the genotypic space (3.3.1).

In the rest of this section, I shall address each of these points in turn. In chapter 4, I shall further develop the distinction between stationary evaluation and selective surfaces, showing how they may be transformed through the choice of selection operator or the introduction of some 'learning' operator that may or may not act with some cost to an individual's evaluation. In addition, the transformation of population structure through the inheritance of acquired characteristics (IAC) will be discussed.

3.2.1.1 Landscapes in Evolutionary Computation.

In the field of EC, there has been considerable confusion between the role of the objective (*evaluation*) function to be optimised and the fitness function that represents the evaluation function transformed through selection and sampling. The two are often assumed to be one and the same, as for example in Goldberg's widely used textbook (Goldberg, 1989), but as I shall demonstrate, this is not the case. In biology too, there are certain 'philosophical

problems' with the notion of fitness (e.g. (Byerly & Michod, 1991; Sober, 1984)), and it may be that a clearer understanding of fitness in evolutionary computation terms may shed some light on the biological position.

In the simple model of the 'within generation' components of a simple genetic algorithm, figure 3-1, the space of individual genotypes, **G**, maps into an evaluation space, E, directly (as in the case of the MaxOnes (bitcounting) evaluation function[23]); or according to their phenotypic expression, **P**, derived through some development function[24]. An example of the latter case is the MaxInt evaluation function, where the aim is to maximise the integer value of an individual phenotype that is suitably coded for by the genotypic bitstring. The aim of the exercise is to optimise the evaluation function according to some criterion. What this typically amounts to is minimising the error between expressed phenotypes and some optimal, target phenotype.

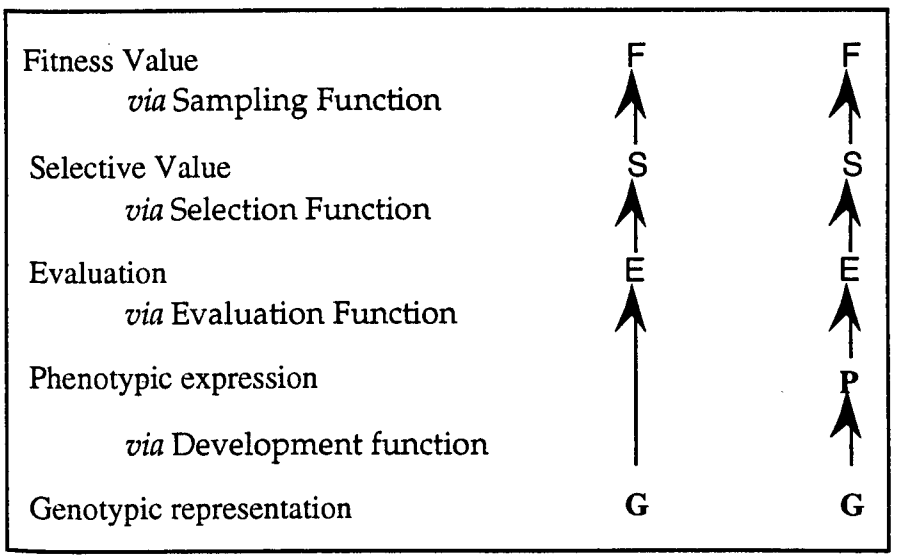

Figure 3-1: Within generation functions of an EA utilising a direct encoding (e.g. for bitcounting evaluations) or simple developmental mapping from genetic representation to evaluated phenotype.

---

[23] See section 2.6.1.

[24] A further decomposition of *phenotype* into *morphology* and *behaviour* (i.e. the relationship between an individual's morphology and its environment) is essential for discussing many Alife models, but this distinction lies beyond the scope of the simple models described in this thesis.

In simple optimisation models, all the determinants of selective value are known (i.e. the evaluation and selection functions and the optimisation criteria). The *selective value*, S, returned by a selection function supplied with an *evaluation*, may itself be thought of as a 'fitness propensity', or 'prescriptive fitness', since it defines the probabilistic distribution of selected parents. The 'actual' or 'real fitness' (fitness, F) of an individual is then a simple function of the *actual* number of times it is sampled for breeding. The *effective fitness* (Stephens et al., 1997) is the expected proportional representation of the individual in the next generation population. This usage is in contradistinction to the typical conception of a fixed landscape described by a static evaluation surface overlaying a fixed structure Hamming graph, since:

- surfaces of selective value are defined on the basis of the evaluation of an individual relative to the evaluations of other population members;

- surfaces corresponding to the *expected evaluation* of individuals in the next generation reflect the current population structure and the current genetic operator set (Stephens et al., 1997); this approach underwrites the description of population dependent neighbourhoods described below (3.2.2.3).

Note that selective value and fitness are equivalent if an ideal sampling function is employed, although typically two sorts of noise will be introduced during sampling: one due to finite population size (i.e. required proportions are not supported by the finite population size), and the other due to sampling error.

Generally, and again in contrast to traditional usage, the fitness function represents the evaluation function combined with a selection function and a sampling algorithm, rather than the evaluation function alone. All these measures are instances of the general class of *performance measures* and all may be visualised as landscapes or surfaces over an appropriate space. The approach also suggests the possibility of an additional performance measure, specifically the mean evaluation of the *selected* population, in contrast to the mean evaluation of the population as a whole. This may be of interest in that it allows one to measure the gross correlation between the mean evaluation of selected parents and the mean evaluation of the population they give rise to through recombination. Such measures frequently offer insights into the effectiveness of given operator sets for particular evaluation functions (see for example, (Manderick, de Weger & Spiessens 1991)).

(Packard, 1988) made a related distinction between *a priori* (evaluation) and *a posteriori* (fitness) functions. The former relate to explicit, predefined objective functions that the user wants to optimise; the latter are rather more akin to biological measures of fitness, describing as they do observed population frequencies of evolved individuals in an artificial environment such as one might find in an artificial life simulation. Note that for the Alife case, the observed populations distributions are generated through the survival of evolved entities who must often sufficiently satisfy some *a priori* selection requirement with an implicit evaluate-and-select mechanism, such as "if (energy obtained from food consumed - energy expended in work) < 0 then die". In this particular example the condition part of the statement specifies an evaluation function, the action part executes a selective process. (Lund & Parisi, 1994) comments similarly on a set of experiments in a simple Alife setting in which the amount of energy extracted from particular food elements is allowed to evolve. This corresponds to the evolution of individual evaluation functions, since "the internal mechanisms for processing ingested elements...are more or less the same thing as the fitness formula"[25]. It is also worth pointing out that for this example, the availability of 'food' is partially responsible for determining the strength of selection, and is in turn dependent on the behaviour of the

---

[25] In Alife models where there is an implicit evaluation function embedded in all the agents that represents some function to be optimised, it is possible to evolve this function (Lund & Parisi, 1994; Lund, 1994; Lund, 1995). This stands in contrast to the use of a single, fixed internal evaluation function applied equally to all individuals. In addition, a halfway house between these two approaches is available - evolving an internal model of the imposed evaluation function, and using the internal function to mediate behaviour. In a review article on decision making in animals, (McFarland, 1977) supposes that behaviour sequences are characterised by environmentally defined *cost functions*, whose optimisation represents an adapted sequence of behaviours by the responsible individual. Cost functions (i.e. *evaluation* functions) may also be thought of as defining niches - the choice of which one of several coexistent cost functions to optimise thus represents niche selection in a habitat supporting several niches. However, individuals do not select behaviours on the basis of the (directly unknowable) cost function. Rather, each individual behaves so as to optimise an *objective function* which "incorporates the optimality criterion [which defines the niche] represented as a set of trade-off relationships between the state and behaviour of the animal and their associated costs and benefits" and is further "envisaged as a property of the individual animal, presumably genetically determined, and possibly modifiable by learning. Thus objective functions are expected to differ from one individual to another, although they may be similar in related individuals" (p. 17). NS is thus responsible for evolving individual objective functions which act as a model for the 'evaluation awarding' cost function and govern selection of appropriate behaviours. The "evaluation net" in the ERL architecture of (Ackley & Littman, 1991) effectively performs this task.

population members. As a consequence, the selective environment may exhibit a complicated dynamic. However, it may be possible to characterise the behaviour of certain ecological factors, such as predator-prey population sizes predicted by the Lotka-Volterra equation, for example.

3.2.1.2 Characterising Selection.

The selection strategies used in EAs fall into one of two classes - proportionate based selection or ordinal (rank based) selection (Miller & Goldberg, 1996) (or (Blickle & Thiele, 1995) for a thorough review). *Proportionate* schemes select individuals in proportion to their evaluation, whereas *rank based* schemes order the population in terms of evaluation, and select individuals according to their rank within the population. More formally:

- In a straightforward *proportionate based selection* scheme (PBS), the selective value of an individual *is* equivalent to its evaluation, normalised by the population mean evaluation:

$$S = \frac{E}{\bar{E}} \tag{3.1}$$

  In order to alleviate the problem of very low or very high selective values resulting from evaluations that deviate slightly or significantly from the mean, this expression may be scaled in several ways (Goldberg, 1989).

- In *rank based selection* (RBS), the selective values of individuals are defined in relation to the other individuals making up the current population through a ranking operator. For example, in linear rank selection individuals are ranked in evaluation order and then selected with probability:

$$p_i = \frac{1}{N}\left(\eta^- + \left(\eta^+ - \eta^-\right)\frac{i-1}{N-1}\right) : i \in \{1,\ldots,N\} \text{ (Blickle \& Thiele, 1995)} \tag{3.2}$$

  where p_i gives the probability of selection (i.e. the selective value) of the ith ranked individual, N is the population size and the rank maximum $\eta^+ = 2 - \eta^- : \eta^- \geq 0$ parameterises selection strength. Selection probabilities for the range of linear rank selection strengths and a population of fixed size 50 is given in figure 3-2. For rank minimum (η^-) values approaching 1.0, the difference in selection probability between the highest and lowest ranked population members is minimal (i.e. weak selection, see below). Under RBS, individuals with a similar valuation may be ranked randomly with

respect to each other and then either assigned different selective values or the mean selective value over the similarly evaluated individuals. In either case, the (time) average effect is to reduce the variation of *expected* selective values, i.e. to flatten the landscape over the similarly evaluated individuals. Consequently, the selective surface may no longer be described by a deterministic linear function of an individual's rank position in the population. This is demonstrated by Example 3.1.

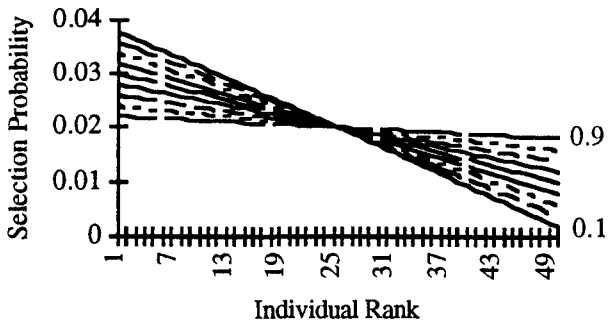

Figure 3-2: Linear rank selection probabilities over rank minima 0.1..0.9 for fixed population size 50.

Example 3.1. If rank selection is applied over a population containing similarly evaluated individuals, then linear rank selection becomes *noisy*. For example, imagine the trivial case of the two bit MaxOnes function, where three possible evaluations are possible, based on the Hamming distance H from the sequence 11. For a population of size N, let a individuals contain two 1s, b individuals carry a single 1, and c individuals carry no 1s (such that $a + b + c = N$). For an individual in a given Hamming band, it's *expected* selective value under rank selection is now given by table 3-1.

Under the assumptions of proportional selection and ideal sampling, it is possible to save the use of the word *fitness* as referring to individual *evaluations*, since under such circumstances fitness is directly proportional to evaluation. However, this is not strictly true for RBS, where rather than reporting *fitness* measures, best and mean *evaluations* are typically presented.

H(11)	Mean expected selective value	Variance
0	$P_a = \frac{\sum_{i=1}^{a} p_i}{a}$	$v_a = \frac{a\sum_{i=1}^{a} p_i^2 - \left(\sum_{i=1}^{a} p_i\right)^2}{a^2}$
1	$P_b = \frac{\sum_{i=1}^{b} p_{a+i}}{b}$	$v_b = \frac{b\sum_{i=1}^{b} p_{a+i}^2 - \left(\sum_{i=1}^{b} p_{a+i}\right)^2}{b^2}$
2	$P_c = \frac{\sum_{i=1}^{c} p_{a+b+i}}{c}$	$v_c = \frac{c\sum_{i=1}^{c} p_{a+b+i}^2 - \left(\sum_{i=1}^{c} p_{a+b+i}\right)^2}{c^2}$

Table 3-1: Expected selective values for a population evaluated according to a 2 bit MaxOnes evaluation function, and containing similar individuals, under RBS.

Changes in evaluation (such as those arising from learning) affect these two types of selection in different ways. According to (Miller & Goldberg, 1996):

> "A selection scheme is said to be scale invariant if multiplying the individuals' fitness by a constant does not change the selection pressure. A selection scheme is said to be translation invariant if adding a constant to every individuals' fitness does not change the selection pressure. Proportionate selection methods are normally scale invariant, but translation variant. Ordinal-based selection schemes are translation and scale invariant."

Selection itself may be characterised in many ways. In particular, selection schemes are often classed as being either *hard* or *soft*, *strong* or *weak*, distinctions I shall apply as follows:

- In *hard* selection, only a fraction of the population may be selected for breeding - for example, in truncation selection with low threshold *t*, where the best *t%* of the population are selected with equal probability. In a *soft* selection scheme, all the individuals in a population at least get a chance to reproduce - for example, in a binary tournament selection scheme, breeding pool individuals are selected by repeatedly picking two individuals at random from the current population and selecting the fittest. Note that both

tournament and truncation selection are parameterised to allow the tuning of the breadth of selection - increasingly hard selection may be applied by increasing tournament size or reducing truncation threshold, for example.

- *Strong* and *weak* selection relate to the *selection pressure* that may be applied against weaker members of the population with reference to the fittest members of the population. The selection scheme used in the experimental work reported in this thesis is parameterised to allow for the tuning of selection strength. As I shall demonstrate in chapter 5, selection strength greatly affects the robustness of individuals discovered through evolutionary search.

3.2.1.3 Understanding the Evaluation Function.

For the *a priori* evaluation functions identified above, a distinction may be made between static functions, which are fixed over time, or dynamic (nonstationary) ones which are a function of time (typically, in optimisation problems, the optimal value is a function of time). Since most of the EA research on *a priori* evaluation functions has related to static cases, the use of the landscape metaphor to characterise the 'absolute worth' of individuals throughout the search space has been widely used. This stands in stark contrast to the surface of selective value, and the *a posteriori* fitness surfaces of Alife which are likely to vary in height *over the same individual* depending on the make up of the current population or the current environmental state. However, it is possible to employ globally imposed evaluation functions that do change over time.

The particular forms of non-stationarity I shall be interested in relate to evaluation functions whose optima may vary in time[26]. That is not to say that such evaluation functions are necessarily noisy[27]. The easiest way of describing fluctuating evaluation functions is in terms

---

[26] (Littman & Ackley, 1991) offer an alternative view of *constant utility non-stationary environments*, in which the development function changes over time but the evaluation function remains constant. The evolvable fitness function of (Lund & Parisi, 1994) corresponds to evolution of the parameters of the evaluation function itself.

[27] A noisy evaluation is one for which the awarded evaluation of an individual is given by the accurate evaluation plus a random noise component.

of the previously mentioned view of minimising the error between expressed phenotypes and some optimal phenotype. Fluctuations arise through the motion of the optima through the phenotypic search space and evolution should track one or other of these optima. In terms of evaluation landscapes, the surface may be likened to an ocean surface, undulating in time as well as over the search space. The faster the rate of change of the optima, the 'choppier' the surface of the 'seascape' (see also sections 2.8 previously, and 3.4 below).

In section 2.8, I reviewed the various dynamics an evaluation function could support. To briefly recap, target optima may be stationary, directional, cyclic or randomly changing; and they may vary between two possible states, or within a finitely bounded or an infinite number of states.

In chapters 6 and 7, I shall present experiments using sinusoidally fluctuating environments in which the aim is to track an optimal sequence that traces a sine wave of fixed period and amplitude through a phenotypic space of bounded real numbers over time. Whether or not the environmental change (i.e. the change in the optimum) is significant in evolutionary terms depends on the amplitude and period of the cycle compared to the available range of phenotypic variation and the generation time; in this way, a slow changing environmental variable may play the role of a parameter in an individual. For the purposes of this thesis, I shall tend to concentrate on a time frame in which the environmental period ranges from a single generation[28] to tens of generations, such as the Intermediate Length Cycle of (Lachmann & Jablonka, 1996), which is loosely defined as "a cycle length that is longer than the generation time of the organism, but not long enough to allow adaptation through the fixation of classical mutations". In the systems under investigation, equilibrium under the genetic operators is typically attained in tens to low hundreds of generations.

3.2.2 Defining the Search Space.

An appropriate description of the search space must consider two things: the nature of the atomic points of the space; and the relationship between neighbouring points.

---

[28] It may be that small within generation fluctuations may be loosely considered as noise. Where the environment changes significantly within a generation, careful thought must be applied to when or how frequently evaluation measures are made, especially if within lifetime learning is supported.

3.2.2.1 The Atomicity of Possible Search Spaces.

In what follows, I shall distinguish a range of search spaces over which one may visualise various performance surfaces (e.g. evaluation landscape, the surface of selective value etc.). It is through these search spaces that a population moves as it evolves, guided in some way by the action of selection. It should become obvious that the atomic points characterising each space constrain the choice of performance surface that one may visualise over that space. Four spaces are distinguished: the *genic frequency* space, the *individual genotype* space, the *genotypic population* space and the *individual phenotype* space. Throughout the rest of this section, I assume a fixed population size, containing N individuals, each with a binary genotype of length L bits, and the direct, independent evaluation of individuals (i.e. for the genotypic/genic spaces, the developmental map is taken to be the identity). Table 3-2 summarises the nature of each genic/genotypic space described below.

- in the *genic frequency* space, each dimension corresponds to the proportion in the population of binary alleles set to **1** at each locus. This sort of space is typically used for plotting population mean evaluations; the evaluation of individuals and the mean evaluation of particular populations is only recoverable if genic (individual locus) evaluations are additive. Mean selective value under RBS has no sensible application over this space.

- the *individual genotype* space, in which each node corresponds to a unique genotypically defined individual. This space is suitable for identifying the independent evaluation of every possible genotypic individual. Given a performance surface over this space, it is possible to generate the corresponding mean for a given set of gene frequencies or a given population. The reverse mapping (from means to individual evaluations) is not possible unless there is additivity of individual genic evaluations. Selective values arising from RBS are not representable over this space since a population is not representable. However, it is possible to address this by augmenting the space in the following way: associate with each genotype an integer quantity with the range 0...N, with the constraint that for a population of size N at any instant the sum of these values taken over the whole space equals N. Setting this value represents the number of individuals in a population bearing that genotype. There is then a unique mapping between the genotypic population space and this augmented individual genotype space, and a many-one map from the augmented space to the genic frequency space. In addition, the mean evaluation of a population may be derived from individual genotypic evaluations.

- the *genotypic population* space, where each point represents a particular, unique population (Vose, 1993b). This search space will happily underlie a surface representing population mean evaluations. If individual evaluations are independent, the evaluation of an individual is given by the mean evaluation of a population (size N) containing N similar individuals[29]. For an infinite population, the genotypic population space maps surjectively onto the space of genic frequencies. It is possible to associate with each population point the individual genotypic space, within which are identified the individual genotypes in the given population. For each population, the selective value of each individual under RBS may now be represented over the identified individuals within the embedded space.

- the *individual phenotype* space, where each point represents a distinct, phenotypically defined individual. It is impossible to find a *general* expression for the volume of the phenotypic space, as this depends on the developmental mapping from the genotypic space. However, where the mapping is known it *is* possible to calculate the size of this space. For example, to visualise a surface over individuals each comprising a sequence of X Gray coded integers in the range $0 \ldots (2^Y - 1)$, one requires a search space containing 2^{XY} points in order to characterise each individual. For a Z point permutation problem, (e.g. the Z-city Traveling Salesman Problem), there are $Z!$ possible permutations (well formed phenotypic individuals), which is far smaller than the combinatorial search space fixed by the $0 \ldots (Z - 1)$ range of each of the Z genes, which contains Z^Z points.

Generally, I shall assume the individual genotype or phenotype search spaces to underlie discussed landscapes, the search spaces being denoted **G** and **P** respectively. A summary of certain properties of the spaces for the non-developmental case, (strictly, the case where the development function is the identity) is given in Table 3-2. Note that an additional time dimension is required if the temporal evolution of the population is to be displayed.

---

[29] If the population structure affects individual evaluations, then the mapping from population to individual spaces may still hold, although it is likely that the reverse mapping will not. Note that in such a case, as for non-additive genic evaluations and the genic space, the mean population evaluation is not necessarily a useful measure.

3.2.2.2 The Underlying Structure of Landscapes.

There has been much interest of late on the structures that form the basis of an evolutionary search space; that is, on the relationship between neighbouring points of the search space. A view that is increasingly important is of a search space represented by a graph whose nodes represent individuals, or pairs of individuals, and whose edges represent operator defined connections (e.g. (Culberson, 1994; Gitchoff & Wagner, 1996; Jones, 1994; Happel & Stadler, 1995)). Neighbouring individuals (or populations, in the case of (Vose, 1993b)) are then connected by single applications of the genetic operators. This view makes concrete the notion of *neighbourhoods*, (3.2.2.3 below) since the neighbours of an individual are all connected to the individual by a single, operator induced edge.

Search Space Nodes	Volume of search space	Surfaces	
		Evaluation	RBS
Genic frequencies (A)	L real dimensions	Population mean (1)	×
Genotypic (individual genotypes) (B)	2^L	Individual (2)	×
Genotypic Populations (C)	$\binom{N+2^L-1}{N}$	Population mean (3)	N/A
Genotypic Populations including explicit individuals (D)	$2^L.\binom{N+2^L-1}{N}$	Population mean (4); Individual (5)	N/A √

Table 3-2: Example search spaces that may underlie a 'performance landscape' in an algorithm without *development.*

Typically, the search space over which a landscape is visualised is the individual genotype space. This genotypic space is usually structured as a Hamming graph, in which undirected edges connect binary represented individuals to neighbours a single Hamming unit away. Under this view, a global perspective is offered over what is essentially a topological map of the whole search space (i.e. locally connected, uniformly weighted edges connecting single bit neighbours). The Hamming 1 structure is attractive because a) it offers a readily understood neighbourhood relation; b) it is fixed. It is often the case that this structure is identified with a search space structure induced by a mutation operator, under the assumption

that, to all intents and purposes, mutation is most likely to result in a single bit flip per individual (of course, whether this is true or not depends on the setting of the operator; however, the rate is often set at 1/L, and where an exact operator is applied (the mutation equivalent of stochastic universal sampling) this assumption will hold). The operator neighbourhood approach tries to capture the sense of operator defined moves over a fixed Hamming graph structure by calculating the probability that a single operator application will generate offspring a given Hamming distance away from an individual. Where crossover between individuals is supported, the distribution of neighbouring points becomes rather less clear. However, for populations of size two, (Culberson, 1994) demonstrated an isomorphism between crossover and mutation on binary strings by comparing graphs structured by single bit mutation and crossover between complementary individuals. (Gitchoff & Wagner, 1996) provided a similar result through comparing the simple mutation graph with a crossover hypergraph for which nodes corresponded to the single individual nodes of the mutation graph, but whose (hyper)edges were complex.

Even in the asexual case, the Hamming one graph may not appropriately represent points connected by mutation. For example, in a contrived example, where an exactly even bit flip mutation operator is used (flips exactly 2, 4, etc. bits per individual), the binary space is partitioned into two halves, one containing strings with an even number of 1s, the other containing strings with an odd number of ones. Using such an operator, if the initial population contains individuals bearing only an even (or only an odd) number of ones, *in principle* only half the search space may be visited.

3.2.2.3 Landscape Microstructure - Operator Neighbourhoods in the Genotypic Domain.

In my review of GAs (2.4.1.1), I mentioned how the genetic operators of crossover and mutation may be described in terms of the neighbourhoods they induce. In this section, I describe the operator induction of these neighbourhoods in more detail. As well as providing a visually appealing way of defining operators, knowledge of these neighbourhoods also allows us to visualise the fine grained structure of the search space in the vicinity of a population, and hence characterise the ruggedness of the performance landscape over this region of the search space.

Definitions.

A note on notation: the Hamming distance, i , between two individuals, g and g' is denoted $H(g', g) = i$; the Hamming band i around an individual (i.e. the mutant class, i) is given by $H(i|g)$ and is strictly the set of individuals exactly i bit flips away from the individual g. Populations are denoted using **bold** characters: let us define $\mathbf{g}$ for the set of individuals in the current population, size N, with members $g_0 \ldots g_{N-1}$, $\mathbf{g}'$ denoting the next (directly descendant) generation. Each individual is assumed to carry L bits used for evaluation purposes.

The Mutation Neighbourhood.

Ideally, the search space structure for an asexual (mutation only) search should be presented as a fully connected graph, with undirected edges weighted according to the probability of going from any one individual, to any other, under a single application of the current mutation operator. For the mutation only case, with a single, fixed, universally applied mutation rate, it is then straightforward to define a (2^L, 2^L)-matrix, M, giving the probabilities of going from any one individual to any other under mutation. For a search space, A, over points $a_0, \ldots, a_{2^L-1}$, elements of the population vector, $p = (p_0, \ldots, p_{2^L-1})$, give the probability, p_n, that individual a_n is in the current population. The mutation matrix is now defined by letting $M_{i,j}$ be the probability that a_j will mutate into a_i. In the absence of selection, the probabilities of finding individuals in the next generation is simply $p = Mp$. Selection may be accounted for by a diagonal matrix, S, such that $S_{k,k}$ is the selective value of a_k within the current population. Note that the diagonal elements of S are calculated (according to the specified selection regime) as a function of the current evaluations of the current population members. Mutating a 'breeding' population arrived at through the application of selection is now given by MSp. Whilst the mutation-operator matrix is likely to be symmetrical, (the likelihood of generating a_i from a_j is the same as a_j from a_i), when selection is taken into account the likelihood is that the probability of selecting a_i and mutating it to a_j will not be the same as selecting a_j and mutating it to a_i, since the selective values are likely to be different.

If necessary, the search space distance between individuals may given in terms of the number of times an operator must be applied to generate one individual from another. The result of such an analysis is a fully connected, individual-centric graph with weighted edges connecting individuals representing the number of operator applications required to get from one individual to another with a certain probability. Such an approach is not as alien as it may

seem when first compared to the Hamming 1 structure. For example, visualising the Hamming n neighbours going away from a single individual is easy. However, if the neighbourhood relations between these neighbours is to be visualised at the same time, the visualisation breaks down. Using any visualisation tool, where search spaces are large there may be a problem with visualising any measure over the whole space to a reasonable degree of resolution. In such a case, it is likely that at any one time a very much smaller subspace will be the focus of attention, again bringing us back to a local rather than global perspective.

In the neighbourhood approach, I try to combine the simplicity of a Hamming structured space with the actual 'moves' through the search space induced by the application of real operators. A mutation operator that flips each bit independently with probability r_m, induces the whole mutation evaluation neighbourhood $N\_mg(g)$. $N\_mg(g)$ acts as a marker for a) all the individuals that may be reached from g by a single operator application and b) the probability with which each of those possible offspring will be generated.

The probability p_m, such that the offspring, g', of an asexually reproducing individual, g, will lie within the Hamming band i around the parent, $g' \in H(i \mid g)$, (i.e. the parent receives exactly i mutations during asexual reproduction):

$$p_m(g' \mid g) = \binom{L}{i} r_m^i (1 - r_m)^{L-i} \tag{3.3a}$$

where $$\sum_{i=0}^{L} p_m(g' \mid g) = 1 \tag{3.3b}$$

and $$\sum_{i=0}^{L} \binom{L}{i} = 2^L \tag{3.3c}$$

i.e. the mutant offspring are binomially distributed over the neighbourhood.

Knowledge of (3.3) means that the links on a fully connected graph may be weighted according to the probability that the link will be traversed. Note that where a modifier gene for mutation is supported, identical individuals over the L bits used for evaluation but carrying different mutation rates will have different offspring distributions (replace r_m by $r_m'(g)$, which corresponds to the self-adaptive mutation rate of individual g).

It is straightforward to calculate the likelihood that any particular individual, g', will result from the application of the mutation operator, m, to an individual, g:

$$p_m(g' \mid g) = r_m^H (1 - r_m)^{L-H} \tag{3.4a}$$

$$\text{where} \quad H(g', g) = H \tag{3.4b}$$

The effective extent of the operator may be defined relative to some appropriately chosen threshold. For example, we may define the 95% likelihood neighbourhood, j_{95}, as follows:

$$j_{95} = \min \lfloor j \rfloor : \sum_{i=0}^{j} p_m(g' \mid g) \geq 0.95 \tag{3.5a}$$

That is to say, the lowest integer value of j for which at least 95% of the offspring of an individual are no further than j Hamming units away from the individual. For low mutation rates, j will be small. This represents a justification for the use of the Hamming 1 graph as the 'default structure' for a landscape.

An alternative formulation allows us to set the mutation rate such that 95% (say) of the offspring of an individual are no further than j Hamming units away from that individual:

$$r_{m,95} = \min r_m : \sum_{i=0}^{j} p_m(g' \mid g) \geq 0.95 \tag{3.5b}$$

Such a view of landscapes replaces the simple fixed Hamming 1 structure with a family of individual centred, weighted graphs, one for each population member(!), where weights represent the likelihood of reaching any given individual from the one. This focuses attention on the behaviour of elements within the current population in terms of population structures that are likely to result in the next generation. An understanding of the local stability of populations is the intention, rather than longer term forecasting.

I shall now consider the size of the exact nth genotypic *mutation* neighbourhood of an individual g, $N\_mg(n, g)$. The neighbour of an individual in this case is an individual into which the first may be transformed, or from which the second is derived, by the application of an exactly-n bit flip mutation operator. For the particular individual concerned, its genotypic mutation neighbourhood is the set of all mutation neighbours of the individual,

plus the individual itself (that is, it is inclusive). Unless otherwise stated, the mutation neighbourhood is assumed to be the *single* (exactly one) mutation neighbourhood. So for an individual 000, its mutation neighbourhood comprises the individuals 000, 001, 010, and 100. It is not hard to calculate the size of the (single) mutation neighbourhood: for individuals length L, the size of the single mutation neighbourhood of an individual, g, denoted $|N\_mg(1, g)|$, is given by 1 (the original individual) + L single mutant neighbours; hence $|N\_mg(1, g)| = L+ 1$. Usually, it will make sense to normalise neighbourhood sizes with respect to the size of the search space, thus giving an impression of the proportion of the whole search space in principle observable by a searching population. For an *exactly n bit* mutation operator, which generates individuals in the nth *mutant class* (Eigen et al., 1989), there are $\binom{L}{n}$ mutant neighbours of an individual (not including itself). Consequently, for an *exactly-n bit mutation* operator which is taken to allow $0 \ldots n$ bit changes, the size of the neighbourhood is given by:

$$|N\_mg(n,g)| = \sum_{i=0}^{n} \binom{L}{i} \tag{3.6}$$

with a limit $n = L$ in which case $|N\_mg|_{\max} = 2^L$.

That is, the nth mutation neighbourhood includes the n-1...0th mutation neighbourhoods. Additionally, the single mutation neighbourhood of the single mutation neighbourhood (*sic*) of an individual is equivalent to the double mutation neighbourhood of that individual.

The mean valuation of a Hamming band around the parent individual is given by:

$$\bar{e}(H(i \mid g)) = \frac{\sum_{h \in H(i|g)} h}{\binom{L}{i}} \tag{3.7}$$

By combining (3.3a) and (3.7, an approximation for the mean valuation over the whole of the neighbourhood may now be obtained (i.e. the expected valuation of the offspring of an individual):

$$\hat{e}(N\_mg(g)) \approx \sum_{i=0}^{L} p_m(g' \mid g)\bar{e}(H(i \mid g)) \tag{3.8a}$$

The exact mean valuation over the neighbourhood is generally given by:

$$\hat{e}(N\_mg(g)) = \sum_{i=0}^{2^L-1} p_m(g_i \mid g)e(g_i) = \sum_{a \in A} p_m(a \mid g)e(a) \tag{3.8b}$$

(Recall A is the search space). Note that by changing the mutation rate, the expected valuation of the offspring will be altered. For two 'selectively identical' individuals carrying *different* self-adaptive mutation rates, the expected valuation of their offspring is thus likely to be different.

An exact expression gives the expected valuation of the next generation as (s is the selective value):

$$\hat{e}(\mathbf{g}') = \sum_{g \in \mathbf{g}} s(g) \sum_{a \in A} p_m(a \mid g)e(a) \tag{3.9}$$

It is interesting to note that for an asexual regime with an *exactly n* bit mutation operator, there are different consequences for n odd or even:

- for *n odd*, (i.e. a mutation operator that flips exactly 1, 3, 5, etc. bits), and $n < L$, every point in the search space may be visited given the repeated application of the operator to any individual;
- for *n even*, (i.e. one that flips exactly 0, 2, 4, etc. bits), and $n < L$, the underlying graph is partitioned into two halves - one containing all genotypes with an even number of 1s, the other containing all genotypes with an odd number of 1s. For example, for exactly two bit mutation, the mutation operator applied to any pair of bits will produce its complement, maintaining the odd/even aspect of the 0 or 1 bitcount (00 <-> 11, 10 <-> 01). If the initial population contains individuals bearing only an even (or only an odd) number of 1s, then the search can in principle only visit half the number of points in the search space as a whole;
- for *n even or odd*, and $n = L$, the graph is partitioned into $2^{(L-1)}$ distinct subgraphs, corresponding to complementary pairs of individuals.

In the case where the population schema is exhaustively sampled by the population, then the population is effectively a point source at that instant, with mutation neighbourhood:

$$\left|N\_mg(1, \mathbf{g}_{\text{def}})\right| = 2^W(L - W + 1) \tag{3.10}$$

Note that the mutation operator is able to generate individuals in the next generation that lie *outside* the population schema that defines the parental generation.

One might also speculate on the 'reverse mutation neighbourhood' of an individual. That is, the set of individuals that can generate the one with a single application of the operator. Note that this supports *asymmetry* in the mutation induced search space structure, as for example in the case where individuals carry different self-adaptive mutation rates. The notion of reverse neighbourhoods also provides a way of visualising the construction of the effective fitness of an individual.

The Crossover Neighbourhood.

The second genetic operator to consider is *crossover*. Typically, one of three styles of crossover (single point, double point, or uniform) is implemented, each generating its own associated neighbourhood, although other variants are possible. Crossover neighbourhoods are denoted $N\_xg(style, \{indl1, indl2,...\})$; where only a single argument is given, context will make it clear whether this refers to the style of crossover, or a set of individuals:

- the *single point crossover* neighbourhood, $N\_xg(1, \{g_i: 0<i\leq n\})$, is taken as covering the set of all possible single point crossover operations (including null crossover) between the parents, g_i. For two distinct parental individuals, g_1 and g_2, differing by H > 0 bit positions, (i.e. separated by Hamming distance, H >0), the size of the single point crossover neighbourhood is:

$$|N\_xg(1, \{g_1, g_2\})| = 2H \qquad (3.11a)$$

  So for $g_1 = 000$ and $g_2 = 011$, H=2 and $|N\_xg(1)| = 4$, namely: 000, 011, 010, 001. An upper bound is given for the case of complementary parents, where H=L and hence $|N\_xg(1)|_{max} = 2L$

- the *two point crossover* neighbourhood, $N\_xg(2, \{g_i: 0 < i \leq n\})$, is taken as covering the set of all possible two point crossover operations (including null crossover) between the pair. For two distinct individuals, g_1 and g_2, H bits apart, the size of the crossover neighbourhood is:

$$|N\_xg(2, \{g_1, g_2\})| = 2 + H(H-1) \text{ (Gitchoff \& Wagner, 1996)} \qquad (3.11b)$$

An upper bound is given for the case of complementary parents, where H=N and hence $|N\_xg(2)|_{max} = 2 + N(N-1)$.

- the size of the *uniform crossover* (Syswerda, 1989) neighbourhood, $N\_xg(U, \{g_i: 0 < i \leq n\})$, is given by:

$$|N\_xg(U, \{g_1, g_2\})| = 2^H \tag{3.11c}$$

with upper bound $|N\_xg(U)|_{max} = 2^L$. We thus see how in principle the uniform crossover of two complementary individuals allows offspring to be sampled from anywhere in the search space. In cases H>2, uniform crossover allows samples to be taken from a greater proportion of the search space than single point crossover allows. In a population with u unconverged loci, the neighbourhood size is thus $N\_xg(U) = 2^u$. That is, for uniform crossover, the size of $N\_xg(U)$ for the population as a whole may be calculated from the selected population width.

- in *bit based simulated crossover* (Syswerda, 1993), *for the population*, $\boldsymbol{g}$, the size of the crossover neighbourhood, $N\_xg(S, \boldsymbol{g})$, is effectively given by:

$$|N\_xg(S, \{g_1, g_2\})| = 2^W \tag{3.11d}$$

where W is the width of the BP-SPDS[30], with upper limit $|N\_xg(, S)|_{max} = 2^L$. Knowing the BP-SPDS also allows the likelihood of generating any *particular* individual in the neighbourhood to be calculated.

It is interesting to note that for a population $\boldsymbol{g} = \{g_1, g_2\}$, $H(g_1, g_2)$<=3, that $N\_xg(, \boldsymbol{g})$ is a proper subset of $N\_mg(, \boldsymbol{g})$ for the modes of crossover that I have considered. That is, where parents are no more than three Hamming units apart, single mutations of both individuals searches over at least the same volume of the search space as crossover. In general, for single point crossover, H>3, (2H-6) individuals[31] from the crossover neighbourhood will fall outside the union of the single mutation neighbourhoods of the parents. For uniform crossover, the number is $(2^H - 2(H+1))$; that is $2^H - 2H - 2$. In each case, the union of parental nth mutation neighbourhoods will cover at least the crossover neighbourhoods for parents $(2+n)$ Hamming units apart. In any crossover scheme, crossover

---

[30] Bitwise probabilistic selected population defining schema, section 2.5.2.

[31] That is, 2H - 2 parents - 2.2 mutants corresponding to crossover 1 bit in at either end for each parent.

is redundant unless the parental schema is of at least width 2 (i.e. at least two *s). The greater the distance between *s, the more likely that offspring different to the parents will be generated.

When addressing the problem of the distribution of offspring over the crossover neighbourhood, one must additionally consider a) the likelihood of any two (or more) individuals being picked as parents; b) the probability that any given individual will result from crossover between the parents. Note that the landscape structure in this case (i.e. the transition matrix between individuals) is rather more sensitive to current population structure. A description of the crossover structure may be generated as follows: for each individual, the probability (given it has been picked as a parent) of generating any child will be the sum over each other possible parent of the probabilities that the child will result from the crossover operation. An additional component will account for the probability that the individual is chosen as a parent. The work involved in the calculation may be halved by assuming that (for two parents) two offspring result, complementary at each complementary parental locus. This approach is of a similar style to Stadler's notion of P-structures (Stadler & Wagner, 1996), which map parental couples onto possible offspring. In turn, it allows for the definition of a matrix that describes the likelihood of mapping from each individual to each of its possible offspring, given all possible parents. Composing this matrix with the mutation matrix then gives a generational, or recombination, operator matrix.

In contrast to the relative simplicity of analysing mutation operators, crossover offers far more difficulties, typically as a result of having to account for the relative distribution of like and unlike alleles in parents in single and two point crossover. However, calculating the probability distribution of offspring using uniform and bit-based simulated crossover (BBSC) is possible since all loci assort independently.

When considering uniform crossover, the visualisation trick is to imagine a Hamming line of length W (the number of unfixed alleles) connecting the two parents. For two parents, g_1 and g_2, such that $H(g_1, g_2) = H$, and one of the two offspring, g', the probability that the child will be a distance i from one particular parent is given by:

$$p_U(\{g_1, g_2\}, g' \in H(i \mid g_1)) = \binom{H}{i}\left(\frac{1}{2}\right)^H \tag{3.12a}$$

$$\text{for} \qquad 0 \le i \le \lfloor H/2 \rfloor \tag{3.12b}$$

Note the symmetry in the solution - for $i > \lfloor H/2 \rfloor$:

$$p_U(\{g_1, g_2\}, g' \in H(i \mid g_1)) = p_U(\{g_1, g_2\}, g' \in H((i - \lfloor H/2 \rfloor) \mid g_2)) \qquad (3.12c)$$

Also note that for two offspring:

$$H(g_1, g_2) = H(g_1', g_2') \qquad (3.13a)$$

and again a symmetry relation holds:

$$p_U(\{g_1, g_2\}, g_1' \in H(i \mid g_1)) = p_U(\{g_1, g_2\}, g_2' \in H((i) \mid g_2)) \qquad (3.13b)$$

Using BBSC, offspring are generated according to the probability distribution of each allele at each locus in the selected population, and so the distribution of offspring over the genotypic search space is trivial.

Unlike mutation, crossover is only able to generate instances of individuals that are covered by the defining schema of the parental population. As has been mentioned, if single offspring are produced, then the population width may be reduced under crossover, although initially present variation may be preserved eternally if *direct replacement* schemes such as that of (Rowe & East, 1995) are implemented. Selective sampling similarly serves to reduce population width by sampling only actual *instances* of the parental population schema. The structure of landscapes over crossover neighbourhoods has been considered in some detail by (Stadler & Wagner, 1996).

One other aspect of crossover that the selected population defining schema and the neighbourhood approach suggests for small populations is the different effect on gene pool variation that results from crossover producing a single offspring, as opposed to two complementary children. For example, for two parents 0010 and 0111, the parentally defined schema is 0*1*. Where two complementary offspring are produced, as for example in (Culberson, 1994), this defining schema will be transmitted to the next generation gene pool and the Hamming distance between the (unmutated) offspring will be the same as that between the parents. If, however, there is only a single offspring from this mating, then variation is lost from the 1st and 3rd loci as far as this contribution to the next generation goes. For such a short string, it is likely that other breeding pairs will make up this loss. However, if the genotypic length is large, and the population relatively small, such losses are

unlikely to be recovered except through mutation. A line of reasoning similar to that used to justify Muller's Ratchet may possibly be useful in taking this intuition further.

The Recombination Neighbourhood and the Relative Roles of the Genetic Operators in Converging Populations.

As populations converge, there are two things to bear in mind - the evolution of the population defining schema; and the changing importance of crossover and mutation, at least in terms of the evolution of the *size* of their respective neighbourhoods. Having introduced the mutation and crossover neighbourhoods, it is now possible to use them to investigate the relative roles of crossover and mutation in the search process. I shall term the volume of the search space G in principle searchable by a population, $\boldsymbol{g}$, of individuals acted on by the genetic operators of mutation and crossover as the genotypic search space, or recombination, neighbourhood of the population, $N\_rg(\boldsymbol{g})$.

The recombination neighbourhood (and offspring distribution) represents the neighbourhood induced by all the available genetic operators to the selected population; so for example, in a typical case, the recombination neighbourhood represents the 'neighbourhood composition' of the crossover and mutation neighbourhoods of each (pair of) individuals. The resulting offspring distribution is then a summation of the function composition of the mutation and crossover offspring distributions of each (pair of) individuals in the current population.

Now, as the population, **g**, converges under selection towards an individual g, $N\_xg(\boldsymbol{g}) \rightarrow g$, and $N\_rg(\boldsymbol{g}) \rightarrow N\_mg(g)$. For selected population width W and individual length L, assuming that crossover can in principle exhaustively sample the current selected population defining schema in producing the next generation, then $2^W(L - W)$ individuals outside that population schema may be seen by an exactly one mutation operator (that is, for each individual covered by the population defining schema (of which there are 2^W), each has $(L-W)$ single mutant neighbours).

Also note that for 2 mating individuals distance H apart, crossover will always act to produce offspring such that the distance from either parent to child is less than or equal to H. In the case of mutation, and two individuals H apart, a single bit mutation applied to either one of them will serve to reduce the distance between them H/L of the time; so for example, for $H < L/2$ and the asexual case we would expect the mutant offspring of one individual to be further away from the other individual than its own parent is more than half of the time.

Landscapes, neutrality and modifier genes.

I discussed earlier (2.6.2) how the addition of selectively neutral modifier genes (such as a mutation rate modifier) may be described in terms of a variety of different landscapes.

In the first case, taking the search space to be the space over individuals ($L + |r_m|$) bits long, where $|r_m|$ is the size in bits of the mutation rate modifier, the valuation landscape is seen to incorporate significant neutrality (i.e. individuals are connected through neutral modifiers).

The evolutionary behaviour of the modifier may be considered by constructing a second landscape that plots a surface of 'marginal evaluations' over a search space of modifier genes (cf. (Altenberg & Feldman, 1987):

$$\bar{e}(r_m') = \frac{\sum_{g \in \mathbf{g}} e(g)\delta(g)}{\sum_{g \in \mathbf{g}} \delta(g)} \tag{3.14}$$

where $\delta(g) = 1$ if $r_m'(g) = r_m'$, $\delta(g) = 0$ otherwise.

Finally, the effective valuation landscape (cf. (Stephens et al., 1997)) describes the mean expected valuation of the offspring of each individual, equation (3.8), again over the whole search space (i..e. individuals ($L + |r_m|$) bits long)[32]. On this landscape, there are no obviously distinguishable neutral networks.

Learning Neighbourhoods.

As with other operators, adaptive local search (learning) operators may be described in terms of the neighbourhoods they induce. It has been suggested that if beneficial consequences of learning, such as the Baldwin Effect, are to be exploited, recombination and learning neighbourhoods should overlap each other considerably (Mayley, 1996b). This is

[32] Note that it is also possible to consider the 'effective marginal valuation' landscape by replacing the individual valuation term in equation (3.14) by the mean expected valuation of the offspring of each individual (equation (3.8)).

easily achieved if moves through the search space induced by learning are defined in terms of the genetic operators.

Where IAC is supported, the learning neighbourhood effectively describes *another* genetic operator, since moves through the search space initiated by learning become immortalised as changes of the genotypic representation of an individual.

3.2.3 The Ruggedness of Landscapes.

Having identified the structure of the search space underlying a landscape, it now becomes possible to characterise the *ruggedness* of the landscape. Intuitively, a landscape may be more or less rugged, although like 'complexity', ruggedness means many things to many people. Colloquially, "a *smooth* landscape is one in which neighbouring points in the space have nearly the same fitness...A *maximally rugged* landscape is one in which the fitness values are entirely uncorrelated. Knowing the fitness at one point would then carry no information about the fitness of neighbouring points" (Kauffman, 1993). Typically, ruggedness is referred to as a property of a landscape as a whole that relates in some way to the distribution of fitness values (evaluations) over a structured search space. The degree of ruggedness is often taken to be an indicator of how effective an evolving population is likely to be in locating (global) optima. As such, measures of ruggedness have been used as predictors of the effectiveness of GAs at finding optima on various landscapes. This approach reflects the widespread and unspoken assumption that evolution acts as a global optimiser (de Jong, 1993). In this thesis, I take the view that evolution acts as a robust optimiser that locates evolutionarily stable solutions, or quasi species (Eigen et al., 1988), and consider ruggedness in this light. In order to do this, I shall couch the notion of ruggedness in terms of the relative evaluations of (in the fact, the evaluation correlation between) parents and their offspring in an evolving population. By so doing, not only are the genetic operators that define the abstract, 'out of time' neighbourhood relations of the search space taken into account, but also the population structure and selective environment that determine the selective landscape perceived by an evolving population. Landscape structure is also reflected in the degree of correlation of an operator (Lund, 1994; Manderick et al., 1991; Stadler & Wagner, 1996), which is a measure of the degree of similarity of the evaluation of a parent individual in one generation and its offspring in the next (assuming a fixed evaluation function). (Altenberg, 1995b; Altenberg, 1997a) refines this approach by defining the *evolvability* of an operator as the likelihood that the operator will produce offspring

superior to their parents. In terms of the neighbourhoods described above, operator correlation is 'visually' demonstrated by imagining the range of evaluations over the operator neighbourhood. Each of these measures makes the same point - the degree of ruggedness (correlation, evolvability) *depends on the particular choice of operators and the given evaluation function*[33].

3.2.3.1 Categorising Historical Measures of Ruggedness.

Measures of ruggedness may be loosely classified in terms of the landscape properties they reflect and the uses to which they are applied:

- measures based on the distribution of local optima: for example, fitness of local optima, lengths of adaptive walks to optima, and number of local optima (Weinberger (1991); Kauffman, 1993);

- the correlation between evaluations of neighbouring points in the search space: for example, the autocorrelation of a random walk through the landscape (Weinberger, 1990), and a derived measure, the correlation length (Manderick et al., 1991; Hordijk (1997); the operator correlation coefficient (Manderick et al., 1991), and the correlation function (Stadler, 1995). Note that each describes a single measure for the whole landscape, calculated in the absence of selection;

- the spread of fitness evaluations of single mutation neighbours of a local optimum, often approximated as the individuals a single Hamming unit away from it (Kauffman, 1993);

- an indicator of the likelihood that an operator set will effectively find a (global) optimum - for example, fitness distance correlation (Jones & Forrest, 1995), evolvability (Altenberg, 1995b).

The first three properties represent varying degrees of scope in their application. The number of local optima is a measure taken over the whole of the landscape; the correlation measures

---

[33] In fact, a stronger claim is possible: for cases where the surface ruggedness is non-homogeneous over the search space, the exact location of the population within the search space may influence the ruggedness over it.

are derived from evaluations of individuals encountered on a long, genetic operator induced walk through the landscape - these are often interpreted as characteristic of the whole landscape under the assumption of isotropy; and the distribution of fitness values in the vicinity of a local optimum offered as a measure of the local ruggedness about a distinct point in the search space.

The properties also differ in other important ways. For example, the optima related measures belie an assumption about the behaviour of populations evolving under selection (i.e. that locating optima is the 'goal' of evolving populations under selection). In contrast, the correlation measures are typically applied to random walks (i.e. walks in the absence of selection) under the assumption of statistical stationarity of evaluations encountered along the walk. Another difference reflects the different assumptions about the perceived use of the measures: globally derived, 'one-number-per-landscape' measures suggest a desire to predict the effectiveness of a GA in locating 'good' optima, or the likelihood of identifying the global optimum - this is particularly evident in the case of Fitness Distance Correlation (Jones & Forrest, 1995), for example, where the correlation is measured between evaluations and distances *from the global optimum.* Measures derived from sampling walks through the landscape refine this view to one in which the effectiveness of the chosen operator set is predicted (at least in cases where the walk is based on operator induced steps); and locally applied measures provide information about the evaluation sensitivity of individuals to perturbations in particular regions of the landscape.

3.2.3.2 A Measure of Local Ruggedness - The Runtime Operator Correlation Coefficient.

Now, operator correlation and evolvability both relate to the ruggedness of a landscape viewed over a search space defined in terms of operators applied to individuals. However, it is also possible to consider the ruggedness of a landscape viewed over *populations.* Recalling 3.2.2.1 on search space atoms, and the work of (Vose, 1993b) on Markov chain models of evolution through the space of possible populations, it is possible to describe a surface of population mean evaluation over the genotypic population space. Neighbouring populations are those that may be reached by application of the genetic operators to a particular population. On a graph based view of the search space, the edges may be weighted by the likelihood that a neighbouring population will form the next population. These weights will be a function of the operator parameters and the effects of selection. In effect, the neighbourhood of a population may be interpreted in terms of a many-worlds model, where

each of the neighbouring populations is potentially the population that will form the next generation. Since each of the edges from any given population are likely to be weighted differently, the landscape should be transformed over each neighbouring state by a factor related to the quantity required to normalise each edge weight. The ruggedness of the population mean landscape is then reflected by the spread of mean values over the many possible directly descendant populations. The notion of the degree of evolvability *of the population* may also be entertained, along with the idea of a recombination operator correlation coefficient derived from the relationship between mean evaluations of sample populations and their possible descendant populations.

The advantage of considering ruggedness over a population search space, as opposed to the individual space is that it takes into account the rates at which genetic operators are applied. On the individual view, the correlation of the operator does not consider how often that operator may be used - it only comments on the effectiveness of it if it is used[34]. Since different rates of operator application affect the structure of the population search space, they will affect the ruggedness of the landscape. This result, intuitive though it is, is not supported by the view of evaluations over a search space of individuals. However, by considering the range of evaluations over the recombination neighbourhood of the population, one can get an idea of how 'stable' the population is likely to be, since by 'looking into the future' information about the evaluations of individuals likely to populate the next generation may be obtained. If the surface is rugged over the recombination neighbourhood, it suggests that the population will be in a state of considerable flux.

Finally, consider the case where the evaluation function is dynamic over time. Ruggedness of the surface over a fixed individual now becomes meaningful in a temporal sense. As operator correlation may be defined across generations, temporal ruggedness may be viewed in a related way. That is, a temporally rugged landscape is one for which the value of the surface over an individual in one generation is significantly different to the value over an identical individual in the next generation. This may be likened to a 'choppy sea', and hence a seascape, rather than landscape, metaphor, as mentioned previously (section 3.2.1.3; recall also figure 2-7, where the rate at which, and extent to which, evaluations of individuals changed was dependent on the location of the individual within the search space). Temporally smooth landscapes, on the other hand, appear to change only slowly.

---

[34] The exception being if one imagines the evolutionary landscape to be the effective fitness of (Stephens, 1997).

For the purposes of this thesis, I shall describe *local* landscape ruggedness in terms of a correlation measure derived from the operator correlation coefficient initially presented in (Manderick et al., 1991), the *run-time operator correlation coefficient, opcc-r*. The operator correlation coefficient describes the similarity of parent and offspring evaluations given the application of a genetic operator. The measure may be applied to a population sampled randomly over the whole search space to give a single landscape measure, or to individuals within an actual population (either under or in the absence of selection). For asexual reproduction, the operator correlation coefficient is defined with respect to the covariance of evaluations between parents, g_i, and their offspring, g_i':

$$\rho_{op} = \frac{\mathrm{cov}(e(\mathbf{g}), e(\mathbf{g}'))}{\sqrt{\mathrm{var}(e(\mathbf{g}))\,\mathrm{var}(e(\mathbf{g}'))}} \tag{3.15}$$

For sexual populations, the midparent evaluation, $e(\mathbf{g}_{mp})$, equal to the mean evaluation of an individual's parents, should be used in place of $e(\mathbf{g})$. If the operator correlation is low, then the landscape is rugged in the sense that parents and offspring have dissimilar evaluations. However, selection may still be able to hold selected members a population on a sharp peak. It then follows that the operator correlation by itself does not necessarily reflect the ruggedness *qua* stability of the population, although it does reflect the ruggedness *qua* variance of offspring evaluations (individual sensitivity). Although applied by (Manderick et al., 1991) to individuals randomly sampled from the search space and in the absence of selection, the measure may be applied to members of a given population (including one converged to a single point) either in the presence or absence of selection. In particular, I define the run-time operator correlation coefficient to be applied to selected parents and their offspring: Note that this measure differs from the correlation coefficient of (Manderick et al., 1991) in its application, since it includes selection as part of the operator (as well as being applied to the current, rather than a random, population):

$$\rho_{opcc-r} = \frac{\mathrm{cov}\left(e\left(\mathbf{g}^s_{mp}\right), e(\mathbf{g}')\right)}{\sqrt{\mathrm{var}\left(e\left(\mathbf{g}^s_{mp}\right)\right)\mathrm{var}(e(\mathbf{g}'))}} \tag{3.16}$$

The run-time operator correlation coefficient may be used to provide an estimate of the overall landscape ruggedness if it is quoted for the evaluations of a random population and a first generation, descendant population (for example, the opcc-r between generations 0 and 1 of a GA run).

3.3 Transforming Landscapes Through Development: The Genotype-Phenotype Map

The *developmental map* is responsible for relating genotypes and the phenotypes that usually serve as the basis for their evaluation (Wagner & Altenberg, 1996). Although in a GA any evaluation function may be defined directly over the genotypic bitstring, convenience dictates that they are usually described over some coarser level of description. Typically, the map used in EAs is one-one (e.g. for evaluation functions defined over the integers, the developmental map will be binary or Gray coding), although this isn't *necessarily* so. Simple, but nevertheless interesting, evaluation functions may be defined directly over the genotypic space (i.e. bitstrings), such as the so called *unitation* evaluation functions which act only on the number of 1s in a string (e.g. the MaxOnes, or Parity evaluation functions). Where an evaluation function is constructed over a phenotypic search space, neighbourhood relations *in the phenotypic space* should be defined relative to genotypic neighbours (which are usually conveniently taken to be Hamming 1 neighbours as discussed earlier (3.2.2.2)). If an apparently sensible neighbour pairing in the phenotypic space is not reflected by a neighbour relation between genotypic representations of the phenotypic neighbours, this suggests either the developmental map does not have a simple interpretation, or conversely the Hamming graph structure does not accurately reflect the neighbour relations induced by the genetic operators.

3.3.1 The Valuation Landscape.

The opportunity to clarify further the notion of evaluation landscapes now presents itself. Specifically, a decision must be taken as to whether the landscape refers to the surface over the operator structured individual phenotypic or genotypic search spaces. Following (Aßelmeyer et al., 1996), I suggest that the objective values associated directly with the genotype of an individual are termed *valuations*, whereas the *evaluation* is the evaluation of the particular phenotype. So for example, in the case of MaxInt, where the evaluation function is defined over phenotypic integers, different *valuation* landscapes may be produced by different developmental maps such as binary or Gray coding. However, the evaluation landscape over the phenotypic integers will remain the same (figure 3-3). Where the developmental map is the identity, as I assumed in section 3.1.2.1, then the valuation and evaluation surfaces are identical.

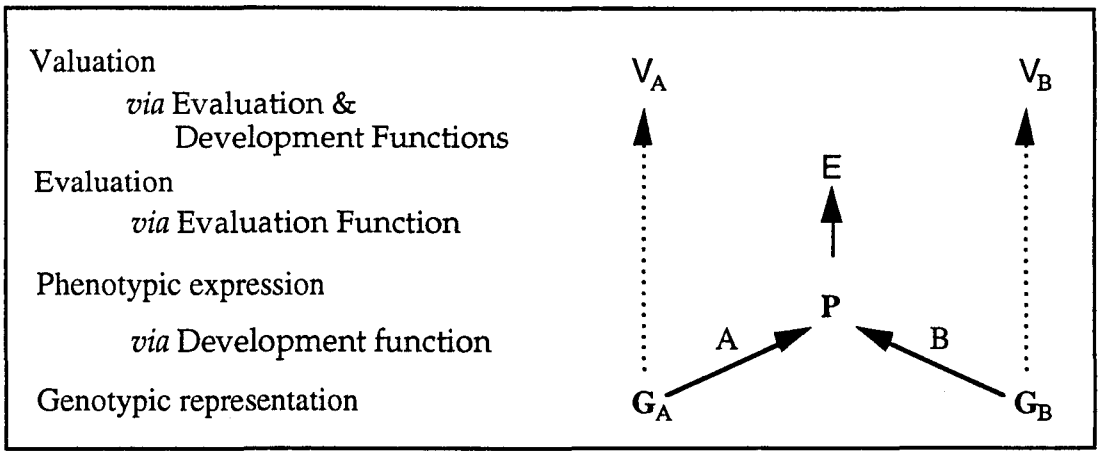

Figure 3-3: Different developmental maps (A, B), may map genotypically different individuals into the same phenotypic expression, realising different valuation landscapes but the same evaluation landscape.

Rather more generally, for the developmental map, *d*, which maps individual genotypes, $g \in \mathbf{G}$ onto phenotypes $p \in \mathbf{P}$:

$$d: g \rightarrow p \tag{3.17a}$$

and evaluation function E:

$$E: p \rightarrow \Re \tag{3.17b}$$

the valuation, V, gives the mapping:

$$V: g \rightarrow \Re \tag{3.17c}$$

where the valuation function represents the function composition of the developmental map and the evaluation function:

$$V(g) = E(d(g)) \tag{3.17d}$$

For the case of the identity development function, $g = d_I(g)$, the valuation and evaluation of an individual are synonymous:

$$V(g) = E(d_I(g)) = E(g) \tag{3.17e}$$

The way the developmental map can derive *different* valuation surfaces for a given evaluation function is absolutely crucial to the central argument of this thesis, since by employing different development functions or genetic operators we can tune the ruggedness of the

valuation landscape which is the landscape ultimately navigated by a genotypically represented population. The ways in which adaptive developmental maps may be implemented by adaptive, within generation local search operators (*learning* operators) will be discussed in more detail in section 3.3.2. The consequent effect on the induced valuation landscapes will be described conceptually in chapter 4 and demonstrated experimentally in chapters 5 and 7. In accord with (Hart, Kammeyer & Belew, 1994), the development function may be subdivided into a maturation component and a learning component:

$$d(g) = l(m(g)) \tag{3.18a}$$

Consequently, as a result of introducing learning, the valuation function before learning:

$$\mathcal{V} = \mathcal{E} \circ m \tag{3.18b}$$

is different to the valuation function following learning defined in terms of moves through the phenotypic search space:

$$\mathcal{V}_{lp} = \mathcal{E} \circ l \circ m \tag{3.18c}$$

or a learning operator defined in terms of moves through the genotypic search space:

$$\mathcal{V}_{\mathrm{lg}} = \mathcal{E} \circ m \circ l \tag{3.18d}$$

The particular form this transformation takes is discussed further in chapter 4.

Now, in section 3.2.2.2 I implied that there was no unique way in which the genotypic search space should be structured, although typically a Hamming one graph was used. If the landscape metaphor is to be consistently applied, then the ruggedness of the valuation landscape will depend not only on the developmental map, but also on the structure of the underlying search space. Figure 3-3 demonstrates this point if developmental maps are both assumed to be the identity, but the two genotypic search spaces are differently structured (e.g. one uses asexual individuals with exactly two bit mutation, the other single bit mutation with crossover). Alternatively, we may consider the induced neighbourhoods of an individual for two differently defined mutation operators, for example. In this case, we may consider one of the search spaces to represent a structurally transformed version of the other. If the maturation function is fixed, there is no learning, and each phenotypic individual receives a unique evaluation, then the valuation surfaces over the two *different* mutation neighbourhoods will also be different.

3.3.2 Hybrid EAs Utilising Within Generation Local Search

An attractive feature of GAs using diverse populations is their ability to find good solutions in a search space whilst still sampling widely from that space. There is, however, no guarantee that any particular individual is locally, let alone globally, optimal. One way around this problem is to introduce plasticity. In 2.3, I suggested there were two ways in which phenotypic plasticity could be introduced into the development process - through reaction norms (i.e. plasticity through maturation) and through 'learning' (within generation adaptive local search). In this section, I shall demonstrate the various ways in which within generation local search may be embedded within an EA, and the ways in which it may interact with inheritance. Note that to be effective, local optimisation need not necessarily be carried out by every individual and may prove more advantageous if only pursued by a portion of the population (Hart, 1994). A formal treatment of genetic algorithms employing universal local optimisation, occasionally known as 'memetic algorithms' is given in (Radcliffe & Surry, 1994).

In the simplest GA, with no local search, evaluation is based on the bitwise assessment of each individual according to the objective function to generate the 'direct evaluation' of the individual, figure 3-4a. The selection function is passed the evaluation of each individual in the population, and the fitness propensity of each individual is calculated. Sampling on the basis of these propensities generates the breeding population which is used by the recombination (crossover and mutation) operators to generate the next generation proper.

One of the simplest models of adaptive local search is trial and error learning. This treats learning as a process of selection, instantiating to all intents and purposes a form of 'within lifetime' evolution: for a given 'problem task', which an individual may learn to 'solve', the individual generates a 'population' of candidate solutions. Each of these is assessed *with respect to the task*, the most successful being adopted as the most appropriate 'solution' (i.e. the 'thing' that is learned). Whilst there are other learning models that may be employed, they share several common properties, two of which are of particular interest: firstly, that there is a certain degree of plasticity between the inherited genotype and the 'genotype' whose evaluation is used in selection, figures 3-4b and c; secondly, that the learning process employs feedback from some objective measure of performance on the task to be learned (Hart et al., 1994). In an EA, this 'learning task' may *or may not* correspond to the 'evolutionary task' specified by the evaluation function (Nolfi et al., 1994). From this it

follows that there may exist several performance measures (each with its own attendant *performance surface*, or *landscape*) that may sensibly be applied to an individual, only one of which is used as its evolutionary evaluation function. In particular, where the 'learning task' and the 'evolutionary task' are complementary (i.e. improving performance on the 'learning task' results in a *degradation* of performance on the 'evolutionary task'), local search is shown to be maladaptive. I distinguish this special case as a model of *fault induction* rather than *learning*.

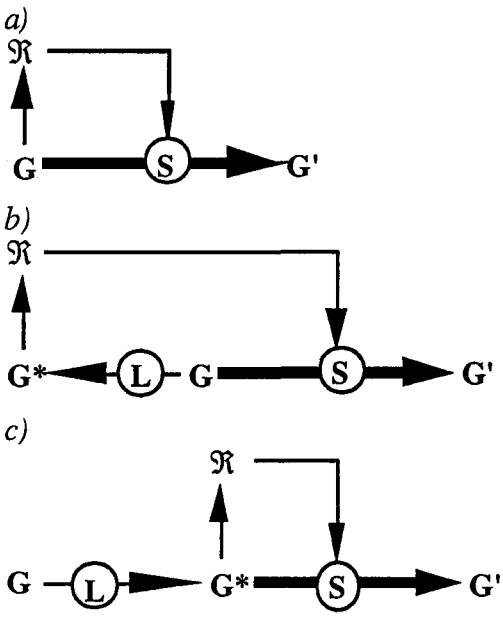

$\forall g \in \mathbf{G}$: $\mathfrak{R}$ = evaluate(g);
select(**G'**)_from(**G**)_on_basis_of($\mathfrak{R}$);

$\forall g \in \mathbf{G}$: $g^* = \text{learn}(g)$: $g^* \in \mathbf{G}^*$;
$\forall g^* \in \mathbf{G}^*$: $\mathfrak{R} = \text{evaluate}(g^*)$;
select(**G'**)_from(**G**)_on_basis_of($\mathfrak{R}$);

$\forall g \in \mathbf{G}$: $g^* = \text{learn}(g)$: $g^* \in \mathbf{G}^*$;
$\forall g^* \in \mathbf{G}^*$: $\mathfrak{R} = \text{evaluate}(g^*)$;
select(**G'**)_from(**G***)_on_basis_of($\mathfrak{R}$);

*Figure 3-4: Learning and inheritance in a non-developmental EA. The current population of genotypic individuals is denoted **G**, and the selected individuals from which the next generation is derived, **G'**; **L** represents within generation local search, and **S** selection acting according to supplied evaluation values; (note the diagrams should be read from the bottom up): a) simple genetic algorithm; b) simple genetic learning with local search and simple inheritance; c) simple genetic learning with local search and IAC; where local search is adaptive, this corresponds to a form of adaptive mutation.*

In a Darwinian system, utilising what I shall term *simple inheritance*, a selected individual passes to the breeding gene pool (from which the next generation is produced) the genotype it inherited. This means that for an individual that employs learning, the evaluation associated with its genotype may actually be the 'direct' evaluation of a different genotype. I shall refer to this set up as a simple, plastic GA, figure 3-4b. Now, it is possible for the genotype which is passed into the gene pool to be the genotype whose *direct* evaluation corresponds to the

evaluation of the learned individual, figure 3-4c. In this case there is thus a mapping from an inherited genotype, to its 'fittest neighbour' under learning, this neighbouring genotype being the one then passed to the gene pool; in other words, the acquired trait is being passed on - the inheritance of acquired characteristics (IAC). This mechanism essentially describes an adaptive mutation process.

As I have mentioned, the GA bitstring may be decoded in a structured way to give a rather more complex phenotype (e.g. binary to reals in a given range is a simple mapping). Additional functions may then be applied to further transform the decoded genes. This introduces a *maturational* mapping, figure 3-5a, from **G** into **P** rather than the identity of **G** and **P** (Hart, 1994b). The evaluation function is now seen to act on the phenotypic expression of a particular genotype. The subject of *automatically* finding appropriate maturation functions to 'facilitate' evolutionary search by producing a smooth *valuation* surface (i.e. automatically solving the *representation problem* in EAs) is likely to become a significant area of research (see, for example, (Altenberg, 1995a; Wagner & Altenberg, 1996) for recent work on the *evolution of evolvability* through the evolution of the genotype-phenotype map).

Again, within generation local search may be introduced although now at the phenotypic level, figure 3-5b. If the model is described purely at the phenotypic level, with parents passing on the 'benefit' of their own learning to their offspring, a form of weak cultural inheritance is seen to occur, figure 3-5c. If the two level genotypic/phenotypic model is maintained, it may be possible to map an acquired phenotype back down to its genotypic representation and pass this on by inheritance, figure 3-5d. In addition, such a model allows for a *dual inheritance* strategy in which the choice of phenotypes acquired through local search may be influenced by those discovered in earlier generations, figure 3-5e.

*Figure 3-5 (overleaf): Learning and inheritance in a simple genetic developmental model. The current population of genotypic individuals is denoted **G**, and the selected individuals from which the next generation is derived, **G'**; **M** represents maturation, **L** represents within generation local search, and **S** selection acting according to supplied evaluation values; (note the diagrams should again be read from the bottom up): a) maturation and simple inheritance; b) maturation and local search with simple inheritance; c) direct cultural inheritance; d) maturation and local search with IAC; e) dual inheritance model - maturation and local search under the bias of the acquired phenotypes of previous generations.*

a)

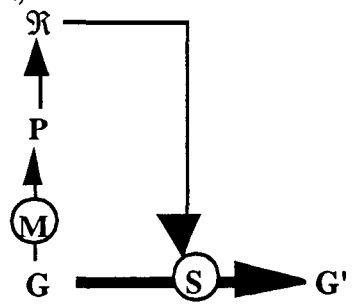

$\forall g \in \mathbf{G}$: p = develop(g): $p \in \mathbf{P}$;
$\forall p \in \mathbf{P}$: $\Re$ = evaluate(p);
select(**G'**)_from(**G**)_on_basis_of($\Re$);

b)

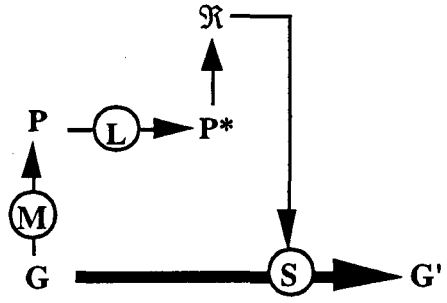

$\forall g \in \mathbf{G}$: p = develop(g): $p \in \mathbf{P}$;
$\forall p \in \mathbf{P}$: p* = learn(p): $p^* \in \mathbf{P}^*$;
$\forall p^* \in \mathbf{P}^*$: $\Re$ = evaluate(p*);
select(**G'**)_from(**G**)_on_basis_of($\Re$);

c)

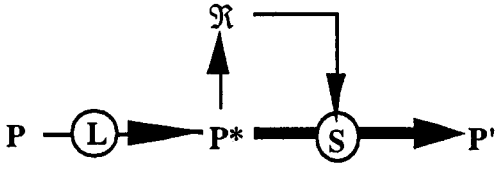

$\forall p \in \mathbf{P}$: p* = learn(p): $p^* \in \mathbf{P}^*$;
$\forall p^* \in \mathbf{P}^*$: $\Re$ = evaluate(p*);
select(**P'**)_from(**P***)_on_basis_of($\Re$);

d)

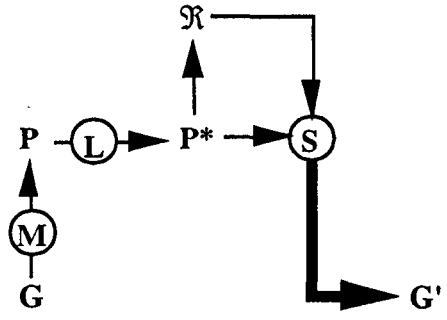

$\forall g \in \mathbf{G}$: p = develop(g): $p \in \mathbf{P}$;
$\forall p \in \mathbf{P}$: p* = learn(p): $p^* \in \mathbf{P}^*$;
$\forall p^* \in \mathbf{P}^*$: $\Re$ = evaluate(p*);
select(**G'**)_from(**P***)_on_basis_of($\Re$);

e)

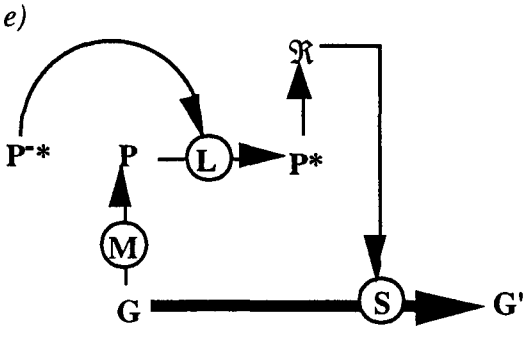

$\forall g \in \mathbf{G}$: p = develop(g): $p \in \mathbf{P}$;
$\forall p \in \mathbf{P}$: p* = learn(p, p^{-*}): $p^* \in \mathbf{P}^*$;
$\forall p^* \in \mathbf{P}^*$: $\Re$ = evaluate(p*);
select(**G'**)_from(**P***)_on_basis_of($\Re$);

3.3.3 The Logical Status of the Developmental, Genotype-Phenotype Map

To consider in slightly more formal terms the nature of the developmental map, and specifically under what conditions it facilitates IAC, let us consider figures 3-6 (originally inspired by (Hart et al., 1994), who was rather more interested in using maturation to focus search onto particular regions of the phenotypic search space). The domain and range of figure 3-6 represent the space of all possible genotypic and phenotypic individuals.

For the bijective map, figure 3-6a, each and every phenotype is represented by a unique genotype. This straightforwardly allows for IAC, since the inverse developmental map is well-behaved and allows us generate the genotypic representation of any phenotype.

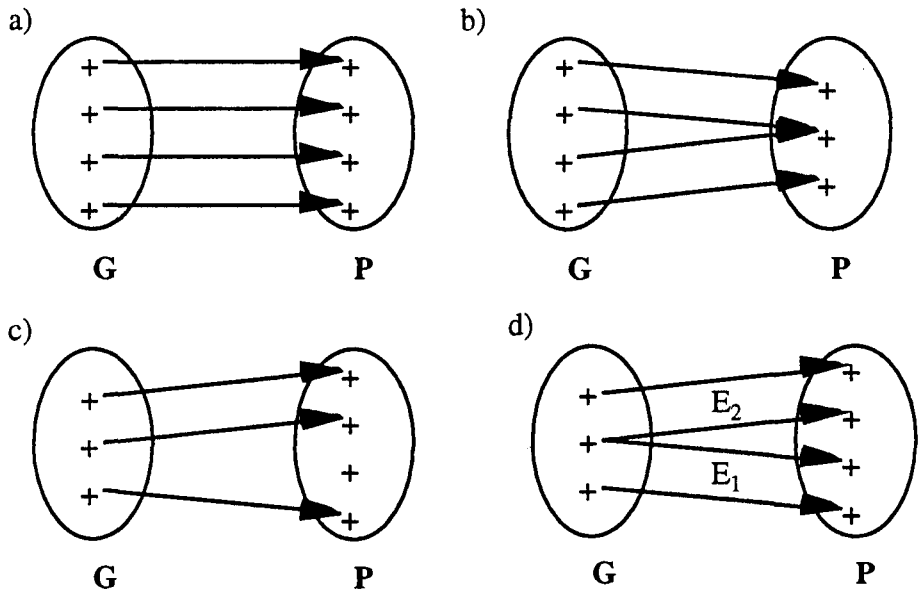

Figure 3-6: Genotype-phenotype mapping functions: a) bijective (one-one onto); b) surjective (onto); c) injective (one-one); d) the existence of an environmentally switched reaction norm.

In the case of figure 3-6b, the surjective map, every phenotype is represented by at least one genotype. In this situation, IAC will be possible, although there exists the possibility of any particular phenotype being represented by several distinct genotypes. Such a mapping may be usefully employed to introduce genetic variation into a population.

In the injective map of figure 3-6c, there may exist phenotypes for which there is no genotypic correlate but which may be acquired through learning. In such a case, for IAC to occur, then some form of 'repair' from unrepresentable phenotypic individual to nearest representable neighbour must be supported (Hart et al., 1994).

If a reaction norm exists, figure 3-6d, the question arises as to how the environmental detector and the alternative expressions are coded for if a phenotype that is typically environmentally switched is acquired through some within generation process.

There are often, however, rather more complex developmental maps being used which may not be so well behaved. For example, at a low level of representation, figure 3-7, where the mapping between genic (i.e. the space of individual genes) and individual phenotypic trait spaces is considered, it is possible to identify rather more closely the range of relationships possible between individual genotypes and phenotypes. This a necessary consideration since the most efficient implementations of IAC should only require that those traits that change during the parent's lifetime are reverse encoded, rather than having to redescribe the whole genotype each generation.

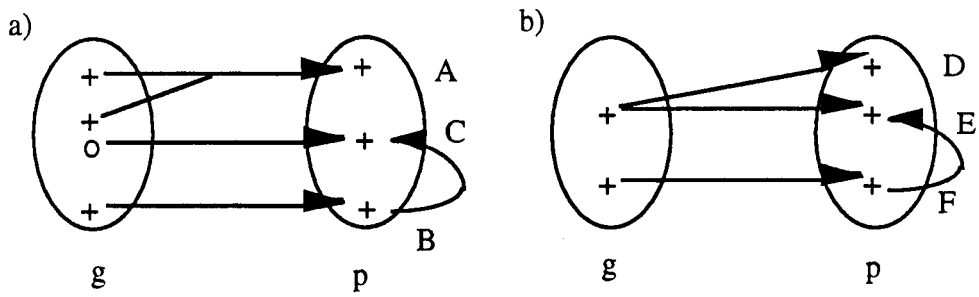

Figure 3-7: Mapping between individual gene and phenotypic trait spaces: a) a polygenic (one trait, many genes) mapping (+ and o represent complementary alleles of a single gene); b) a pleiotropic (one gene, many traits) mapping.

In a *polygenic* trait, figure 3-7a, many genes contribute to the expression of trait A. If A is acquired, then two representational genes must be derived. If a trait is acquired that requires the reverse transcription of a gene that is in conflict with a pre-existent and co-existent *genetically coded for* trait, such as trait C, this may account for the non-existence of a genotypic representation of a given phenotype. In a pleiotropic map, figure 3-7b, where a single gene is expressed in two phenotypic traits, IAC may be prohibited since the acquisition of trait such as E may mean that the reverse encoding of the trait may, under development, realise an additional, 'unacquired' (an unwanted) 'side-effect' trait (D).

Where direct IAC (the reverse transcription of acquired traits) is required, but not supported, it may be possible to augment the inheritance system so that direct transmission of traits at the higher phenotypic level is allowed. That is, a dual inheritance mechanism is introduced.

Finally, a form of inheritance in which a parent 'designs' the representation of its intended offspring may be ultimately possible. Such a model may be based on the *emboitement* style of heredity - the first parent may contain within it a recipe for the construction of its offspring. By taking into account environmental conditions, for example, the parent could tune the design of its offspring to one it assumes will be better adapted than the bare design it carries. This strategy may be thought of as a rather extreme version of the number of supported reaction norms, in which all possible phenotypes are coded for, and from which an environmental switch may effectively decide. Alternatively, one may regard it as a form of 'bequest of characters I would have liked to have acquired', or 'giving my child the education I never had'.

3.3.4 Landscape Microstructure - Operator Neighbourhoods in the Phenotypic Domain.

By introducing a developmental map, the structure of the phenotypic search space may be hard to visualise, especially in the case of complex developmental mappings. Unless otherwise stated, I shall assume that the phenotypic space is structured similarly to the genotypic space (that is, neighbours in the genotypic space are neighbours in the phenotypic space). A Hamming 1 relation is usually assumed for convenience in the genotypic space (3.2.2.2), although genotypic neighbourhood relations that accurately reflect real genetic operators may be more appropriate. If a *simple* developmental function is used, then phenotypic neighbours should interpretable as such in a sensible way (e.g. for a phenotypic search space over the integers, numerically close individuals would reflect an intuitive neighbourhood relation). However, this is not necessarily an appropriate (or even achievable) view. For example, where local search at the phenotypic level is supported, the phenotypic search space structure is perhaps better defined in terms of the neighbourhoods induced by those local search operators.

Where the genotypic neighbourhoods are to be respected (that is, phenotypic neighbours are by definition genotypic neighbours), the corresponding neighbourhoods in the phenotypic domain may be obtained straightforwardly. For example, for the case of mutation, and a

bijective maturation function, the *phenotypic* mutation neighbourhood, N_mp, is given by $N\_mp = d(N\_mg)$. Where the maturation mapping is one-many, for example, it becomes appropriate to talk in terms of a maturation neighbourhood, $N\_dp(g)$, of a genotypic individual which contains only and all those phenotypic expressions of the particular genotype under the maturational mapping being applied. To account for 'reaction norms' the development function should include an environmental argument, $d(e, g)$.

Despite being phenotypic mutation neighbours, this does not mean that two such individuals are neighbours in an intuitive sense at the purely phenotypic level. For example, single mutations at the genotypic level may induce gross distortions at the phenotypic level, whereas genetically distant individuals may bear significant phenotypic similarities. This is all to do with the particular form of the development function used (see for example (Hart et al., 1994) and 3.3.3 previously). In addition, given two individuals that may be interpreted, by virtue of similarity, as neighbours at the phenotypic level, there may exist several (or no) possible neighbour (genetic operator induced) pairings at the genotypic level.

3.4 The Three Phases of Evolutionary Optimisation

Imagine a static evaluation landscape (e.g. a stationary evaluation surface over a Hamming graph) with an initially random, fixed size population dotted around it. Reproduction is asexual, with some given (low) mutation rate so that the Hamming graph represents a reasonable approximation to neighbouring individuals in terms of operator application and structuring of the landscape. If the trajectory through time of an evolving population is plotted within the search space, three phases of the search process may be identified: *exploration*, *exploitation* and *equilibrium.*

- Phase 1, exploration: the initially random population samples over the whole of the search space; this is attested to by the high values adopted by self-adaptive mutation rates, for example, during the early generations of a GA run. Over the course of the first few generations, the initially dispersed population cloud identifies regions of the search space with relatively mean high evaluation. The population begins to condense in these regions as regions of low relative evaluation are selected against, lineages in regions of higher mean evaluation being favoured; the runtime operator correlation remains constant at its initial value;

- Phase 2, exploitation: as time goes by, and phase 1 leads into phase 2, embryonic quasi-species start to condense in one or more region of the search space with relatively high mean evaluation; if the landscape is 'smooth' in these regions, then the runtime operator correlation will start to climb. If operator correlation is low in these regions (the landscape is rugged), offspring of fit individuals may have low evaluation, even though mean evaluation is high; these offspring will thus be selected against in competition with individuals in other sub-populations whose mean evaluation is not significantly lower but whose operator correlation is high. Areduction in the value of self-adaptive rates indicates the increasing exploitation of a particular region of the search space;

- Phase 3, equilibrium: going from phase 2 to phase 3, a single quasi-species dominates and appears to maintain some sort of cohesion over time in a particular region of the search space. Whether this occurs in a rugged or smooth region of the landscape depends in part on the strength of selection, and the phenomenon of *canalisation*, as I shall demonstrate in chapter 5. Evolution of quasi-species along neutral networks is possible, and small satellite populations may appear through mutation to regions of comparable or superior evaluation. These will then either briefly flower and die in competition with the dominant qausi-species, or come to be the dominant quasi-species themselves. If a population encounters a local optimum while traversing a neutral network, then it may disperse as the new optimum is exploited througha brief return to phase 2 behaviour. Exactly this sort of behaviour is observed under the Royal Road evaluation function. Self-adaptive rates typically acquire an equilibrium value and demonstrate steady state behaviour. If selection or 'convergent drift' is strong enough, and there is no mutational mechanism for introducing variation, the population may converge to a point, and stasis. Where mutation is supported, a theoretical result by (Vose, 1993b) suggests that in the long run, the population will converge in the limit around the optimum with the largest basin of attraction.

To summarise this view: phase 1 and 2 give the transient response of the evolutionary algorithm; phase 3 represents its steady state response. Phase 1 represents initial selection and terminates with identification of several high evaluation regions of the search space. Phase 2 represents selection between regions of the search space according to the relative stability of competing quasi-species and terminates with the identification of dominant quasi-species. Phase 3 represents the steady state response and the long term behaviour of the dominant quasi-species.

It is interesting to consider the relationship between the local ruggedness in the vicinity of the population, and the particular stage of the search (i.e. phase 1, 2 or 3) it happens to be in. In particular, if we can find a principled relation between equilibrium dynamics and local ruggedness, we may be able to choose operator sets such that a regions of known ruggedness will be discovered.

Kauffman, commenting on asexual evolution in NK landscapes, writes (my emphasis):

> *"The highest average maintained fitness need not occur on those landscapes with the highest peaks but will occur on those landscapes whose ruggedness best suits the character of the adaptive process.* ***In the present case*** [of the NK landscapes], ***for any fixed size of genotype N, population size, and mutation rate, some particular value of K will yield a landscape whose ruggedness is optimal for maintaining high average fitness.*** *Thus some presently unknown curve in the NK plane corresponds to the optimal landscape to match a given stochastic flow exploring that landscape in order to achieve the highest average fitness. Grounds [exist] to suppose that this curve is related to the mutation rate at which populations just begin to melt from local regions, hence that both evolvability and sustained fitness can be jointly optimised..."* (Kauffman, 1993, p108).

Here, Kauffman assumes a homogeneous ruggedness over the landscape as a whole. By relaxing this constraint and assuming the more general case of a landscape with heterogeneous ruggedness, his argument suggests that the operator set will be more or less adapted to *geographically different regions of any particular landscape*. This provides a basis for the claim that evolution will drive the population towards regions of the landscape for which the current operator set is adapted i.e. that region of the landscape with a ruggedness characteristic of the operator set.

Conjecture 1: *for a given operator set (genetic and selective operators, given finite population size) evolution will drive the population towards a region of the valuation landscape with a well defined, or characteristic, local ruggedness, as given by the runtime operator correlation coefficient, opcc-r (equation 3.16).*

Conjecture 1 thus describes the characteristic ruggedness of the landscape as the local ruggedness of the region of landscape identified by the population in the steady state. It may be interpreted in a strong and a weak sense. Under the strong interpretation, the characteristic ruggedness of an operator set will be *independent* of the evaluation function. Under the weak interpretation, the characteristic ruggedness will be dependent on both operator set and the particular evaluation function. Recalling the 3 phase model, in which I suggested that selection strength is likely to influence the local ruggedness of the region of the landscape to which a population will converge, at least four experimental conditions must be considered to test the conjecture (strong and weak selection vs. strong and weak interpretation). Furthermore, under conditions of weak selection, it is likely that opcc-r will be maximised; where selection is strong, a population may be held around a dominant master sequence in a region of high local ruggedness (that is, low opcc-r).

In chapter 5, I present a series of experiments that present a preliminary test of this conjecture from the point of view of the variation of evaluations in an equilibrium population obtained under various operator sets.

3.4.1 Evolutionary Dynamics of an Initially Converged Population

If the initial population is nearly converged (that is, a quasi-species) then there are several possible trajectories the search may follow. Either a) the quasi-species is stable, dynamic equilibrium is maintained, and evolution progresses along neutral networks until an optimum is discovered (phase 3); b) the population diverges and period of early phase 2 (or even phase 1) search is entered.

3.4.2 Evolution in a Fluctuating Selective Environment

In a sinusoidally fluctuating environment, for example, the valuation surface will smoothly undulate, rather than maintaining a fixed profile (cf. figure 2-7). The population cloud must now move through the search space in order to stay within a region of high valuation. A thought experiment will help clarify this point, as well as serving the secondary purpose of illustrating population stability in a static landscape. Imagine a population cloud sat around the peak of a valuation hill in a fixed landscape, the mutation-selection balance keeping the

quasi-species in a state of dynamic equilibrium (figure 3-8a). If the mutation rate is high, or the selection rate low, the population will be spread out over the hill (although in an asexual population, the cloud forms a ring below the summit at a height determined by the mutation-selection balance (Woodcock & Higgs, 1996)). A consequence of this is that in a static environment, if a valuation peak is very sharp (i.e. localised) a significant proportion of the population (or the recombination neighbourhood) of the population may inhabit the 'lowlands' where valuations are low. The quasi-species will then not be able to regenerate itself and the population cloud will disperse. This effect is illustrated with the novel Tower of Babel evaluation function in chapter 5.

If the valuation surface now changes slightly - the peak of the hill moves - the balance of more and less adapted individuals within the population will change, and with it the mean population valuation (figure 3-8b). Selection will now act so as to shift the balance of the population in the search space towards the location of the new hill. Depending on how far the valuation hill moves within each generation, the population will be more or less able to track the optimum. What is required is that the current population contains members, perhaps on the periphery of the cloud, who will be able to identify the location of the valuation optimum - or at least be within its basin of attraction - *after* the environment changes, as in figure 3-8b. If the valuation landscape changes too dramatically, however, sight of the optimum may be lost to the population, figure 3-8.

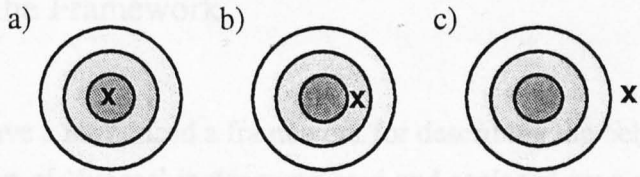

Figure 3-8: A population cloud in some search space. Hatching density reflects population density, **x** *marks the location of a valuation peak: a) the population is centered on the optimum; b) a slight change in the location of the optimum - there are still individuals within the population who may identify the peak and tracking is possible; c) the optimum moves outside the scope of the population, and all individuals receive a low valuation - the population is likely to diverge until the optimum is rediscovered.*

What this thought experiment suggests is that in a fluctuating environment, a more diverged quasi-species will be required in order for the optimum to be tracked (a higher mutation rate will be supported, for instance). This stands out immediately on considering the *distance* the

optimum moves over the search space. (If the distance moved 'per change' is only small, however, tracking may still fail if the rate of change is high). Why should this be the case? The population may be imagined to have a certain amount of *inertia* associated with it. When the optimum moves, it takes selection a certain amount of time to reorganise the balance of the population. For stronger selection, this may happen in a relatively few generations. If selection is weak, however, it may take rather more time to move the 'centre of gravity' of the population to the new optimum. If small but rapid changes in the location of the optimum occur, figure 3-9, then a population cloud of a given size may initially appear to track the optimum through always being able to identify its location, but the mean valuation will fall and the optimum may even be lost.

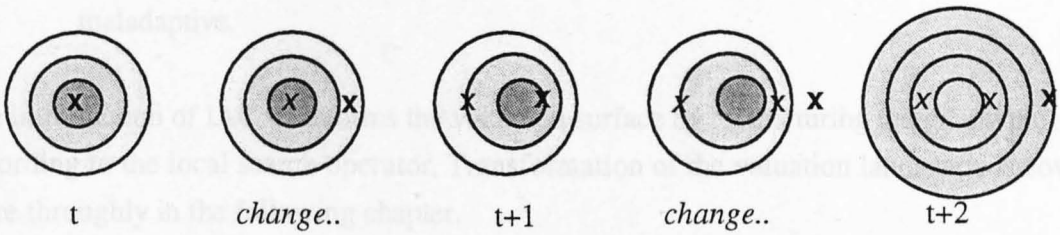

Figure 3-9: A rapidly moving optimum may evade a population moving slowly under weak selection. Hatching density reflects population density, **x** *marks the location of a valuation peak.*

3.5 Reviewing the Framework

In this chapter, I have a introduced a framework for describing the behaviour of evolutionary systems. The notion of 'fitness' is deconstructed and replaced by a set of related measures defined over specific search spaces - the *evaluation* surface as depicting evaluations over the phenotypic search space; the *valuation* surface as those values visualised over the genotypic search space; and, in the general case, the surface of selective value as a surface over a suitably defined population. The transformation of these surfaces through self-adaptive rates and within generation local search is considered in more detail in the following chapter.

The underlying structure of landscapes has been described in some detail (3.2.2), since this has a considerable effect on the 'ruggedness' of the landscape. In particular, operator neighbourhoods are shown to be a useful way of describing the operator induced microstructure of the search space. Using a simple counting argument based on

neighbourhood sizes, a comparison between the relative effectiveness of crossover and mutation at various stages of an evolutionary search is obtained.

The role of development has been shown to act as a mapping between the evaluation and valuation surfaces. A variety of *different* valuations are possible for a population assessed under the same evaluation function through several different mechanisms (figure 3-3):

- by restructuring the genotypic search space structure through the use of different, or self-adaptive, genetic operators;
- through the use of different developmental maps;
- through the application of within generation local search, whether adaptive or maladaptive.

The introduction of IAC transforms the valuation surface by restructuring the genotypic space according to the local search operator. Transformation of the valuation landscape is covered more throughly in the following chapter.

Finally, a three phase model of the evolutionary process has been described. It is suggested that an operator and population dependent measure, the runtime operator correlation coefficient, is appropriate for measuring the 'local ruggedness' of a landscape in the vicinity of a population. Relating the three phase model to this measure results in a conjecture as to the long term behaviour of an evolving population, specifically that operator parameter sets are adapted to (i.e. will tend towards) regions of the *valuation* landscape with a characteristic local ruggedness.

Chapter 4 - The Dynamic Transformation of Performance Landscapes

> "The construction of ... simple models is a basic part of the strategy of theory building in population biology. The reason that elementary general models are useful, despite their simplicity and unrealism, is that even the simplest evolutionary processes are hard to understand. Thus, simple models serve as an essential supplement to intuition, which is often misleading...It is not trivial to keep all (the) interacting parts of the problem straight. Simple models can serve as a check on less formal methods of deductive reasoning, as a basis for constructing more realistic models, and as an unambiguous standard of comparison for purposes of discussion."
>
> (Boyd & Richerson, 1988).

4.1 Introduction

In the previous chapter, I suggested how the notion of landscape ruggedness was necessarily related to both the structure of the genotypic search space, and the particular developmental map applied, as well as the actual evaluation function. In this chapter, I shall consider in rather more detail the way in which the different landscapes (valuation, evaluation, selective value) represent transformations of each other. Throughout, I shall demonstrate the transformations with a series of thought experiments that refer to a simple evaluation function applied to a sample population. The actual behaviour of an evolving population over a series of landscapes transformed by local search is investigated empirically in the following chapter.

The rest of this chapter proceeds as follows. Sections 4.2 and 4.4 offer a simple worked example of the transformations arising from applying particular selection functions and development functions to a small, contrived population being evaluated under two well known, but trivial, evaluation functions. As such, they recap on 3.2.1.2 and 3.3.1 respectively. The intervening section 4.3 describes the dynamic transformation of the valuation landscape arising from the dynamic restructuring of the search space by a self-adaptive mutation rate. In 4.5, I characterise the transformational effect of within generation local search operators on the valuation landscape and the resulting surface of selective values for certain selection functions. This represents a rather more exact interpretation of the way in which learning 'smoothes' a performance landscape than is traditionally (qualitatively) offered. The earlier used example population and evaluation functions shall be reconsidered

in the light of simple local search operators. Section 4.6 considers further surface transformations resulting from exacting an evaluation cost following the application of a given operator. Finally, a possible interpretation of operator behaviour is offered in terms of an image processing analogy.

4.2 The Surface of Selective Value

The surface of selective value represents the selective values assigned to the individuals in a given population and is thus a populatyion dependent measure. For proportionate based selection (PBS), the surface over genotypic individuals represents a simple scaling of the evaluation surface. For frequency dependent selection schemes, such as rank based (RBS) or truncation selection (TS), the surface of selective values over a subset of the genotypic search space that represents the current population represents a rather more significant transformation of the evaluation landscape.

Example 4.1. Imagine a set up utilising truncation selection with elitism and a threshold of $x\%$ (that is, the $100 - x\%$ worst performing individuals in a population with N members are discarded (i.e. their selective value is 0)), selecting individuals from the remaining $x\%$ with equal probability, $1/(x\% * N)$, with the proviso that the best adapted (elite) individual is always selected and placed into the selection pool. The MaxInt evaluation function over the current population may now be transformed into a function that defines a selective surface. The value of the surface over an individual represents the likelihood of it being selected for the next generation. This transformation makes concrete the distinction between evaluation and selective values. See, for example, figure 4-1, which considers the evaluations in a population of 10 binary coded individuals, {1, 3, 5, 6, 6, 7, 8, 9, 11, 13}, for the 4 bit MaxInt problem, and the selective surface over the population under proportional selection (figure 4-1a), elitist truncation selection with threshold 50% and linear rank selection (figure 4-1b).

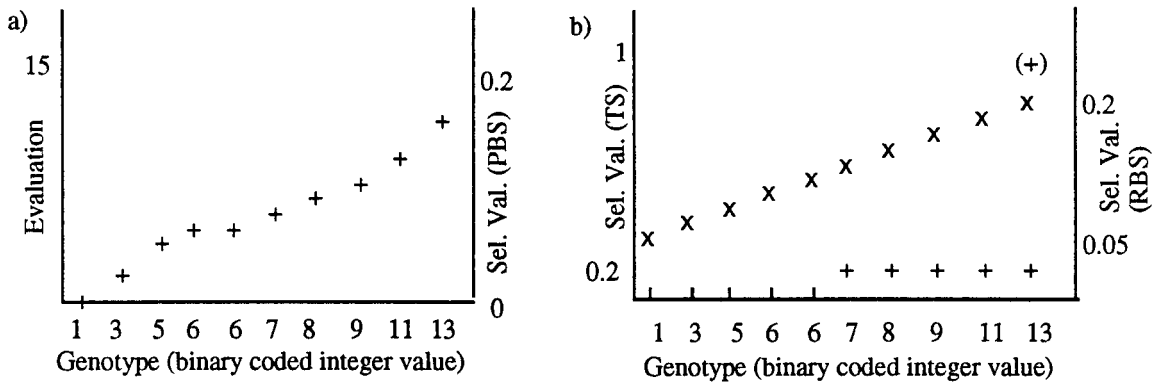

Figure 4-1: a) 4 bit MaxInt evaluation (and PBS selective value) against population members for Example 4.1; b) the distribution of selective values for the same population under elitist truncation selection, with cutoff 50% (+), and linear rank based selection (x) with rank minimum 0.5.

4.3 Transforming Landscapes with Self-Adaptive Mutation Rates

In the previous chapter (3.3.1), I discussed how the valuation surface could be tuned by an appropriate choice of development function, or genetic operators, for a given evaluation function. Modifier genes are often factored into the landscape view of evolution as providing neutral networks within the whole genotypic search space (Altenberg & Feldman, 1987; Stephens et al., 1997). However, under the framework presented in chapter 3, it seems that this view is not appropriate since different modifiers induce different search space structures and hence different valuation surfaces. A potentially more useful coevolutionary model is made possible by considering the landscapes over two coupled search spaces, one corresponding to the modifier search space, the other to the search space that covers the selected portion of the genome[35]. During the rest of this section, I shall apply both approaches. In addition, I demonstrate how the application of the framework allows us to phrase useful research questions, which are posed for future work to address further.

[35] Note that by employing a coevolutionary model, it becomes difficult to distinguish whether the modifier or the evaluated portion of an individual is more or less responsible for guiding the population to a particular region of the overall search space.

Self-adaptive mutation rates coded for by neutral modifier genes allow each individual in a lineage to effectively explore a uniquely structured search space, and hence a unique valuation surface, for that lineage. Where the evaluation function is fixed, a self-adaptive mutation rate should allow for phase 1 exploration of the search space and late phase 3 convergence around some optimum. Exploitation of landscapes correlated to the self-adaptive operator during phase 2 represents evolution over a *non-stationary* valuation landscape as the modifier gene coevolves with the evaluated gene(s) which act as the primary determinants of selection. Thus we are presented with a first puzzle - what is the nature of this coevolutionary process?

Closely related to this is a complementary second question: what performance measure provides the basis for selective evolution through the modifier search space?

- (Altenberg & Feldman, 1987) and 2.5.3 described a view of genetic modifiers in which selection of perfectly transmitted[36] modifier genes is driven by different *marginal fitness* values. Recall from 2.6.2 that the marginal fitness corresponds to the marginal *selective value* of the modifier. For PBS, this is given as the mean instantaneous evaluation of all individuals bearing a particular modifier allele. Note that under this view, the actual value of the modifier does not influence the calculation of its its marginal value. Over time, the marginal values change according to the spread of evaluations in the current generation that resulted from an application of the unchanging modifier rates in the previous generation.

- The more recent analysis by (Stephens et al., 1997) and 2.6.2 employs the single landscape description by considering the evolution of modifiers to be driven by the *effective fitness* of individuals carrying those modifiers. In contrast to the above approach, the effective fitness incorporates a component related to the value of the modifier (specifically as it relates to the likelihood of producing particular individuals from the current population through an application of the currently set operator). The effective fitness of an individual in the current generation thus represents its expected

[36] i.e. there is no source of variation available to the modifier, so it is not self-adaptive - what is described is the modifier value that is stable to invasion by other modifiers.

concentration in the next[37]. Rather than considering evolution over a set of valuation landscapes, the appropriate evolution surface is there argued to be the effective fitness (or mean evaluation of the reverse mutation neighbourhood) over the genotypic search space.

4.4 Transforming Valuations by Development

As discussed in section 3.3.1, for an evaluation function defined over the phenotypic search space, the valuation surface may be transformed by applying different developmental maps from the genotypic search space into the phenotypic space, figure 4-2a. Again, the particular selection function used generates the surface of selective values visualised over the genotypically represented population, figure 4-2b.

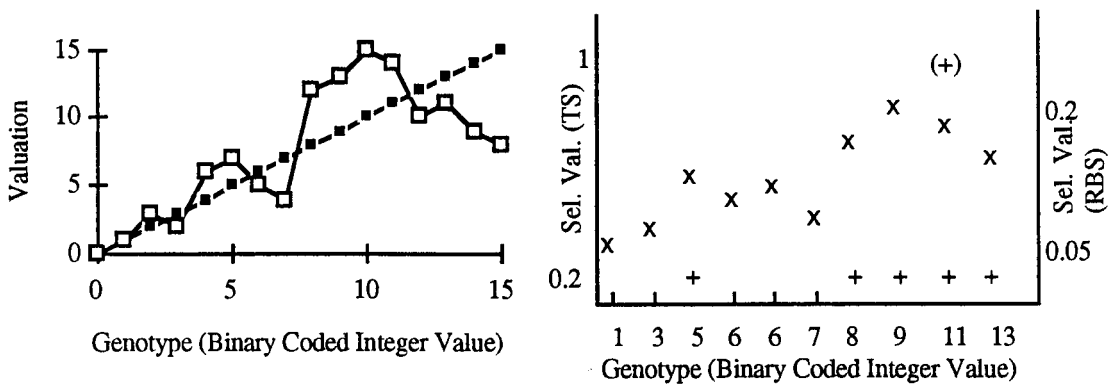

Figure 4-2: a) The 4 bit MaxInt problem using binary coded (solid) and Gray coded (empty) individuals; b) RBS (x) and TS (+) selective values for the Gray coded MaxInt using the population of Example 4.1.

Where a non-deterministic development (or local search) operator is applied, the valuation surface is non-deterministic (i.e. noisy), although the evaluation surface may be fixed. Where the selection function is noisy, (as, for example, in the case of the RBS of similarly evaluated

[37] Note how this definition more closely matches the biological sense of fitness as *measure* of representation within a population (albeit in this sense an *expected measure*), as opposed to a *prescription* of selective probabilities as used in quantitative genetics and EC.

individuals), valuation and evaluation surfaces may be deterministic, even though the surface of selective values is not. Noisy evaluation functions, on the other hand, are likely to induce noisy valuation and selective value surfaces.

4.5 Transforming Landscapes Through Learning

The idea behind introducing a within generation, local adaptive search (learning) operator is to improve the evaluation of each individual. In its application, the local search operator may define moves through the genotypic or phenotypic spaces. For example, in the binary coded TargInt function[38], a learning function may be defined at the genotypic level (e.g. single bit Steepest Ascent (SA) learning on the genotype) or the phenotypic level (e.g. add or subtract one to the integer value of the individual). In either case, local search may be viewed either as augmenting the development function, in which case the effects are limited to a transformation of the valuation surface; or, where the local search is applied deterministically, as transforming the evaluation function[39]. In both cases, the valuation function will be transformed. Unless otherwise stated, I shall view learning as a modification of the developmental map.

In the limiting case of within lifetime search to a local optimum, the valuation of each individual is collapsed onto the evaluation of the local optimum in whose basin of attraction (induced by the learning operator) the individual lies. This approach corresponds to the memetic algorithms described theoretically by (Radcliffe & Surry, 1994), depicted in figure 4-3 (based on (Whitley et al., 1994)). A more limited form of local search describes the operator in terms of the *fixed size* neighbourhood it induces, and this shall be introduced by way of example (a Markovian analysis of these more general *hybrid GAs* (i.e. GAs that employ a more restricted form of within generation local search) is given by (Whitley, 1995)).

---

[38] TargInt is the general case version of MaxInt, where the idea is to find the prespecified target integer (MaxInt defines the target as the *largest* integer in a given range).

[39] e.g. (Whitley et al., 1994; Whitley, 1995) describe an equivalence evaluation function for a simple algorithm *without* learning that behaves identically to a partially deceptive evaluation function under SA learning.

In their widely reported paper on the interaction of learning and evolution reviewed in 2.7, Hinton & Nowlan (Hinton & Nowlan, 1987), hereafter H&N, introduced a simple learning scheme that incorporates several interesting properties. An individual was represented by a string of 20 loci, each of which could take on the value **0**, **1** or **?** with initial probabilities 0.25, 0.25 and 0.5 respectively. The task was to match a target sequence of all **1**s. For individuals containing one or more **?**s, random 'learning' was allowed. A learning trial comprised of setting all the **?**s in a particular individual to 1 or 0, each with probability 1/2, with a maximum of 1000 learning trials per individual. Learning stopped immediately if the target sequence was attained. Two parameters thus govern the extent of learning applied to an individual: the *learning distance* (that is, the scope of the learning neighbourhood given by the number of **?**s in an individual) and the *number of learning trials*. These parameters are discussed in more detail below (4.5.1).

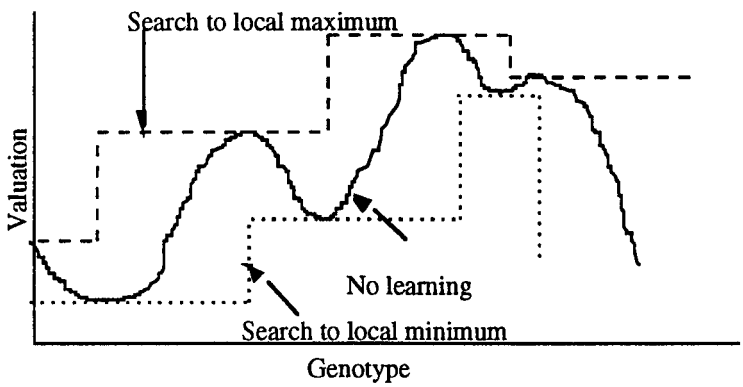

Figure 4-3: Transformation of the valuation landscape by learning and fault induction to local maxima/minima (adapted from a figure by (Whitley et al., 1994)*).*

In the H&N model, plasticity effectively replaces an individual by a virtual population of phenotypic individuals. The recorded valuation of the genotypic individual is then given by the evaluation of the best adapted (phenotypic) member of that virtual population (i.e. the one returning the highest evaluation). This distinction (*of within generation local search transforming the valuation surface* as opposed to the evaluation surface) is a key point of this section. At no time is the evaluation surface transformed - each phenotype has associated with it an evaluation; acquiring a new phenotype through learning allows a genotypically specified individual to acquire a new valuation. The acquired phenotype receives the same

evaluation as if it had been inherited directly[40]. Consequently, we may make a second conjecture that follows on from Conjecture 1 of 3.4.

The essence of the conjecture is that a population whose individuals employ local search will be adapted to a different region of the evaluation landscape than a population employing the same evolutionary operators (i.e. same recombination and selection operators) but no local search. This arises as a consequence of the valuation landscape being transformed by local search. The population moves towards a region of the valuation landscape to which it is adapted (as claimed by conjecture 1). Since the valuation landscape represents a transformation of the evaluation landscape according to developmental (i.e. local search) operators the local search population will equilibrate in a different region of the evaluation landscape than if no local search were applied.

Conjecture 2: *populations containing individuals who apply within generation local search will be driven to a region of the search space where the evaluation surface has a different (characteristic) local ruggedness compared with the local ruggedness about a steady state population with similar evolutionary operators but no local search.*

Transformations of the evaluation landscape in terms of ruggedness are as follows: for adaptive local search, the valuation surface represents a smoothed version of the evaluation surface. *Smoothing* occurs as a result of 'truncating' from the landscape local minima and emphasising the best adapted local maxima; as a consequence the mean valuation in that particular region of the transformed landscape is likely to be increased by comparison with the mean performance of the evaluation surface over the same region.In addition, we might expect that opcc-r, taken on the evaluation surface, will be *lower*than the value obtained for an evolved population of directly evaluated individuals, given otherwise equivalent operators. For maladaptive local search, the valuation surface represents a sharpened version of the evaluation surface. *Sharpening* occurs as a result of truncating local optima and emphasising the best adapted local minima; the effect on mean valuation (compared to the mean evaluation in the same region of the space) is unpredictable. In this case, opcc-r for the steady state population is likely to be higher than the value obtained for a population eveolved over the

---

[40] Where a learning cost is enforced, this is not strictly true - the acquired phenotype may receive a *lower* evaluation than one directly expressed. However, the parsimonious approach is to consider cost as a further transformation of the valuation surface, (see 4.6).

same landscape but withouth the 'benefit' of maladaptive local search. Once again, the strength of selection may qualitatively, as well as quantitatively, affect the extent to which this conjecture holds.

Under maladaptive local search, such as Steepest Descent, and a *maximisation* problem, selection maximises evaluations according to the *highest* evaluated individuals, each of whom report the lowest local evaluation. In adaptive LS, and a *minimisation* problem, (so again, Steepest Descent, say) selection favours the *lowest* evaluated individuals, each of which report the lowest local evaluations. One would not expect, therefore, a population employing a particular local search algorithm in an adaptive sense on a minimisation problem to behave in the same way as a population employing the same local search algorithm in a maladaptive sense on a maximisation problem over the same evaluation function, and vice versa.

Although the valuation surface may be transformed by local search, it is not necessarily the case that the surface of selective values over a population will be affected. By considering the population as a whole, and with learning available to every individual, two 'styles' of learning may be identified with respect to the relative valuations of individuals before and after learning: *rank respectful* learning and *rank transformational* learning. The transformational effect of each style of learning then depends on the particular type of selection (rank based or proportional) being used (4.5.2).

It should also be remembered that for learning operators defined in terms of local search at the genetic level, the transformational effect may be different where different basic developmental functions are defined. This point will be illustrated in 4.5.3 where the effect of SA learning on the MaxOnes and MaxInt evaluation functions will be compared.

4.5.1 Learning Parameters.

In this section, I shall consider in more detail the parameters relating to learning distance and the number of learning trials. Sufficient conditions for deciding whether local search is *probabilistic* or *deterministic* will also be presented[41].

4.5.1.1 Learning Distance: Defining the Learning Neighbourhood.

The particular class of learning operators that I am considering are, like the mutation operator, applied independently to particular individuals. As such, each learning operator defines a 'learning neighbourhood' local to each individual. The size of the learning neighbourhood may be thought of in terms of how 'far away' from the inherited phenotype an individual may learn. So for example, in the model presented by H&N, the number of queries in any particular genotype specifies a learning neighbourhood for each individual, such that an individual bearing q ?s has a learning neighbourhood that comprises 2^q members. If a ? had a 'default setting' in the event of the number of learning trials equaling zero, then q would represent the maximum learning distance (in Hamming units) away from the inherited, default phenotype that an individual could explore. Figure 4-4 demonstrates how the originally specified needle-in-a-haystack evaluation function is then transformed by adaptive plasticity into an *expected evaluation* function that varies smoothly over q.

4.5.1.2 Learning Intensity.

How extensively an individual may explore its learning neighbourhood depends upon the number of learning trials:

- If the learning neighbourhood is exhaustively sampled, (i.e. if the set of distinct virtual population members equals the size of the learning neighbourhood), and the best evaluation is always returned, then the operator is deterministic;

---

[41] In a probabilistic learning scheme, given any individual, the phenotype acquired through learning may be different from one application of the learning operator to the next. On the other hand, in *deterministic* learning procedures, for example SA learning, a given individual always acquires the *same* phenotype (Whitley, 1995).

- If the set of distinct virtual population members represents a randomly sampled subset of the learning neighbourhood, then the operator is probabilistic.

Note that even if the learning neighbourhood is searched exhaustively, the best evaluation returned may not be locally optimal in the strictest sense (i.e. with respect to a repeated application of the particular learning operator). Rather, the best evaluation returned through an exhaustive search of the learning neighbourhood of an individual is the *neighbourhood optimum* for the individual undergoing learning; where the number of trials is such that the neighbourhood is not sampled exhaustively, the best evaluation returned is a *sampled neighbourhood optimum* for the given individual. If a random (learning) sampling strategy is followed, increasing the number of trials may well improve the performance in the H&N example, figure 4-4b.

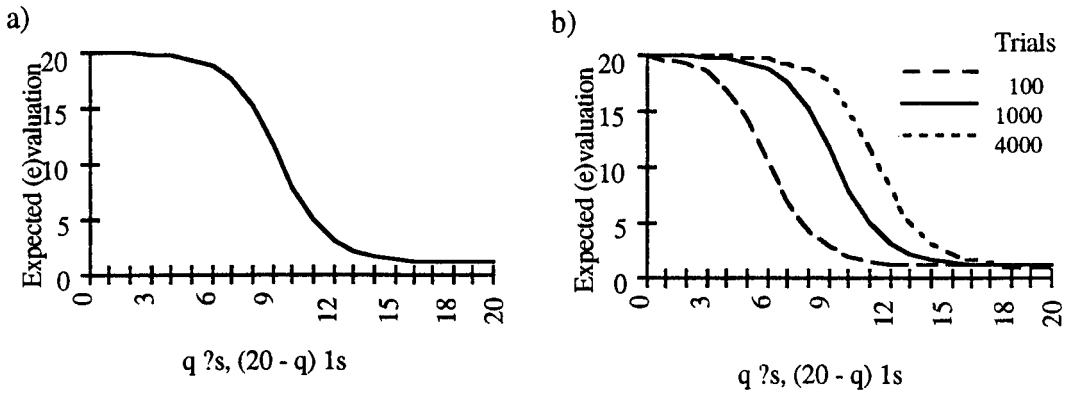

Figure 4-4: Transformation of an impulse evaluation function by learning for the evaluation function of H&N. a) the original experiment, in which the maximum number of trials is 1000, based on an analysis by (Harvey, 1993); b) for different limits on the maximum number of trials.

4.5.2 Learning and Its Effect on Rank Ordering.

The function transformations that are of particular interest here are those that alter the (relative) selective values of individuals in a population. I distinguish two forms of transformation: *v-transformation*, a transformation of the valuation function in terms of re-ordering the rank of individuals; and *s-transformation*, transformation of the selective values of an individual which maintains the rank order within the population. Table 4-1 summarises the susceptibility of PBS and RBS to these forms of transformation.

By defining the direct valuation, v_g, of an individual, g, and the *benefit* from learning, b_g, received following learning by that individual, the valuation, V_g, of g as passed to the selection function is given by:

$$V_g = v_g + b_g \tag{4.1}$$

The conditions for v- and s-transformation are given in 4.5.2.1 and 4.5.2.2 below.

Selection Type	Rank Respectful Learning	Rank Transformational Learning
PBS	*s-transforming*	*v-transforming;* *s-transforming*
RBS	*No effect*	*v-transforming*

Table 4-1: The possible transformational effects of rank respectful and rank transformational learning depending on the style of selection.

4.5.2.1 v-transformation.

If, for any individual, g, in the current population, the rank of an individual before learning is different to its rank after learning:

$$Rank(v_g) \neq Rank(V_g) \tag{4.2}$$

where $Rank(h_g)$ is the rank of a directly evaluated individual, g, with evaluation h_g, in the population *before* learning, and $Rank(H_g)$ is the relative ranking of individuals *following* learning.

4.5.2.2 s-transformation.

If, taken over all individuals, g, in the current population, the set of selective values before learning is different to the set of selective values following learning:

$$\{s(v_g)\} \neq \{s(V_g)\} \tag{4.3}$$

where $s(h_g)$ is the selective value of an individual, g, with valuation h_g, and $\{s(\ )\}$ is the set of selective values over the population.

Although trivial, the distinction between v- and s-transformation implies that the effectiveness of hybrid algorithms is at least partially dependent on the style of selection used. This in turn suggests that analytic treatments of even simple functions (e.g. unitation functions) must be tempered in the extent to which general conclusions are drawn from particular cases.

4.5.2.3 Rank Respectful Learning.

The simplest models of learning may offer the same benefit to each individual in the population, or scale the evaluation of each individual by a constant factor. RBS will be unaffected by either of these forms of learning; PBS methods are unchanged by uniform scaled improvements in valuation, but the resulting selection pressures *are* affected by uniform additive benefits applied to each individual.

For example, for the case of simple valuation proportionate selection, where rank order is respected (i.e. there is no v-transformation) the selection pressure *against* an individual is *reduced* (the selective value is *increased*) if:

$$\frac{v_g + b_g}{\bar{v} + \bar{b}} > \frac{v_g}{\bar{v}} \tag{4.4a}$$

and hence if:

$$\frac{b_g}{v_g} > \frac{\bar{b}}{\bar{v}} \tag{4.4b}$$

where $\bar{v}$ and $\bar{b}$ represent the population mean direct valuation and learning benefit respectively.

In the stricter case of uniform benefits to all individuals, where:

$$b_g = \bar{b} \tag{4.4c}$$

equation (4.4b) reduces to:

$$\bar{v} > v_g \tag{4.4d}$$

If the *lowest* evaluation in the plasticity neighbourhood is presented to the selection function then a 'fault induction', as opposed to learning, model is the result. In such a case, the selection pressure *against* an individual is *increased* (the selective value is *decreased*) if the inequality (4.4b) holds, where b_i represents the absolute valuation *penalty* incurred due to plasticity, and (4.4a) is rewritten as:

$$\frac{v_i - b_i}{\bar{v} - \bar{b}} < \frac{v_i}{\bar{v}} \tag{4.5}$$

For PBS, then, the notions of smoothing or sharpening the surface of selective values have a definite interpretation. If the range of valuations (i.e. best - worst) is the same before and after local search, then compared to the range of selective values for the no search case: a) the range of selective values for the case of *adaptive* local search will be *reduced*; and b) the range of selective values for the case of *maladaptive* local search will be *increased.*

4.5.2.4 Rank Transformational Learning.

Most interesting forms of learning will be rank transformational. As in the rank respectful case, under PBS the selection pressures applied to individuals (and hence the selection function itself) are transformed through learning. In both PBS and RBS schemes, where the relative ranking of individuals is transformed by the learning operator, the selection function is essentially being passed results from a transformed *valuation* function. Strictly speaking, under both rank transformational and rank respectful learning, the selection function itself may not be transformed at all. However, the valuation function will be transformed.

Where rank order *is* affected, then on first impressions it would appear that the transformation of the selective surface is likely to be complex, whatever the selection scheme. However, there is a powerful, simplifying assumption that can be made. As populations converge, adaptive local search applied to one individual is increasingly likely to discover similar adapted phenotypic states to other genotypically distinct individuals. (Mayley, 1997) refers to this as the 'Hiding Effect', since genotypic variation is hidden from selection through the many-to-one genotype-phenotype map (development function) that local search may implement[42]. Under proportional selection, I have already shown how one might

[42] Cost free local search is assumed. Genotypically distinct individuals that discover similar phenotypes and hence receive similar valuations may be distinguished by reducing the (e)valuation by an amount that is some

reasonably expect that adaptive local search will result in a smoother selective profile over the population. In RBS, the effect of Hiding is to 'flatten' the selective values over the similarly evaluated individuals (3.2.1.2) and again preserve genotypic diversity.

For the case of maladaptive local search, it is possible that *different* minima are identified, compared to the case of adaptive local search which is likely to discover *similar* adapted states. Why should this be so? Assuming that the population is gathered around a single optimum, then individuals may be approaching the optimum from different directions i.e. climbing different facing slopes. On this view, if each individual were pushed down the slope it was climbing, each would be likely to identify a different maladapted state, with possible consequences for rank order. Rather than hiding maladapted states as adaptive local search does, maladaptive search *reveals* maladapted states in the vicinity of an individual.

4.5.2.5 Learning in Quantitative Genetics

The formal, quantitative genetics view on learning is that it increases the variance of selection (weakens the strength of selection) (Boyd, 1985; Anderson, 1995). This may be seen as a result either of Hiding, or as a consequence of reducing the range of selective values. In chapter 7, I consider the effect of plasticity and IAC in quantitative genetics models rather more exactly.

4.5.3 By Way of Example...

In figure 4-5, I show how the valuation surfaces for the 4 bit MaxOnes and MaxInt evaluation functions are transformed through two single bit local search operators: adaptive local search through SA learning; and maladaptive local search through the complementary operator, fault induction through Steepest *Descent* 'learning'. Note that although the same bit level operation is being applied in each case, the development function that realises assessed individuals determines whether or not the learning operator is rank respectful or rank transformational. So for example, for the case of MaxOnes, figure 4-5a, the valuation surface is translated along the valuation axis and rank order is preserved (apart from 'outliers' where

---

increasing function of the amount of search required to find the phenotype; i.e. by costing plasticity. See, for example, (Mayley, 1996b).

boundary conditions prevent learning); whereas in the MaxInt case, figure 4-5b, the valuation surface is significantly transformed (and differently so for each operator) and the original rank order is lost.

A further point to note is that MaxInt represents a one-one map from genotype to phenotype/valuation/evaluation, whereas in MaxOnes, the map is many-one (e.g. individuals 1, 2, 4 and 8 all receive a direct valuation of 1, figure 4-5a). This may have consequences for IAC, (see 4.5.5 below).

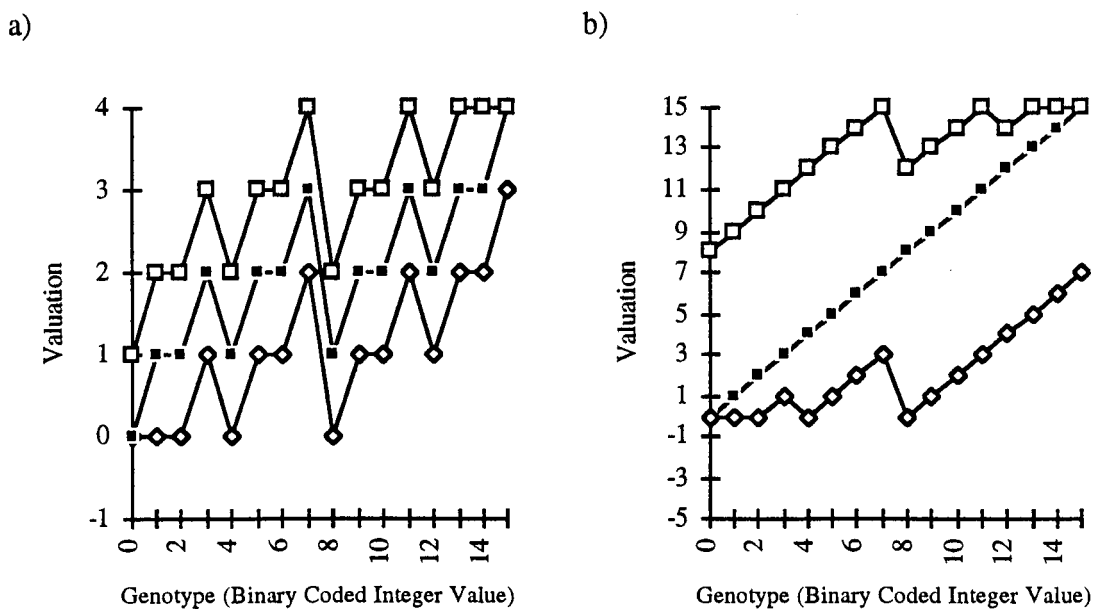

Figure 4-5: Direct valuation (solid square) and valuation following SA learning (empty square) and SD fault induction (empty diamond) versus integer phenotype for 4 bit MaxOnes (a) and MaxInt (b).

For the 4 bit binary coded MaxInt case, there are sixteen distinct individuals each with a unique evaluation. By introducing single bit, SA learning, the valuation profile of the population is transformed and there are only 8 distinct valuation scores, although there are still 16 distinct genotypes. What learning is doing in this case is hiding genetic diversity from the selection function with valuation equivalence. In the next section, I shall demonstrate how an additional transformation of the *population structure* occurs when IAC is supported.

Example 4.2. For the MaxOnes and MaxInt evaluation functions, application of the single bit SA and SD local search operators to the population of Example 4.1, (1, 3, 5, 6, 6, 7, 8, 9, 11, 13), transforms the valuation surfaces over the population to those of figure 4-6 and table 4-2.

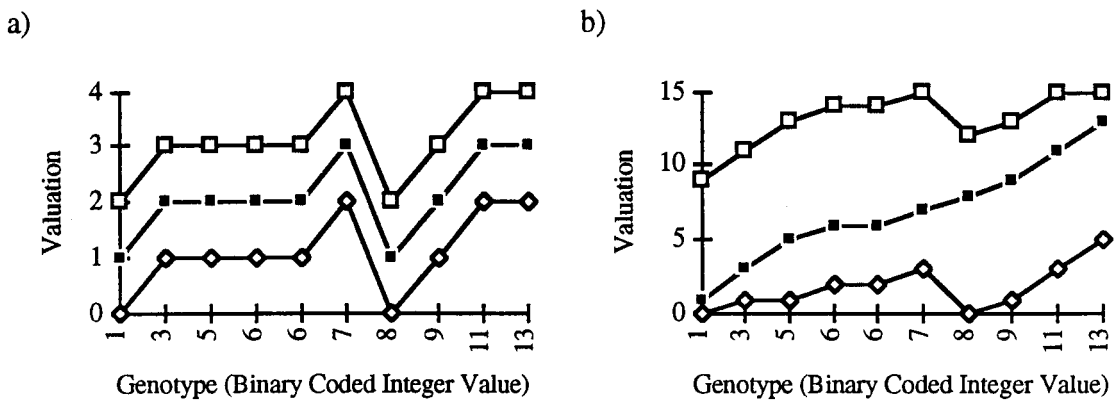

Figure 4-6: MaxOnes (a) and MaxInt (b) valuations for the population of example 4.2 with no local search (solid square), SA learning (empty square) and SD fault induction (empty diamond).

Naive Binary	Naive Int/Bitcount		'Educated' Binary/Int/Bitcount			'Lower bound' Binary/Int/Bitcount		
0001	1	1	1001	9	2	0000	0	0
0011	3	2	1011	11	3	0001	1	1
0101	5	2	1101	13	3	0001	1	1
0110	6	2	1110	14	3	0010	2	1
0110	6	2	1110	14	3	0010	2	1
0111	7	3	1111	15	4	0011	3	2
1000	8	1	1100	12	2	0000	0	0
1001	9	2	1101	13	3	0001	1	1
1011	11	3	1111	15	4	0011	3	2
1101	13	3	1111	15	4	0101	5	2

Table 4-2: A population of 4 bit binary coded individuals given by the population set {1, 3, 5, 6, 6, 7, 8, 9, 11, 13} under single bit SA learning and SD fault induction.

4.5.4 Cheating the Reaper - The Interaction of Local Search and Selection.

In the previous section, I demonstrated how the introduction of a local search operator may transform the valuation landscape over the genetic representation space. As I shall now show, the choice of selection function determines the extent of further surface transformation.

Example 4.3. Using the population of example 4.2 and table 4-2, 4 bit MaxOnes and binary coded MaxInt evaluation functions, figure 4-7 shows the selective surfaces over the population transformed by single bit SA and SD operators under proportional selection. Figure 4-8 shows the selective surfaces over the MaxInt population transformed through SA learning under the truncation selection of Example 4.1.

From figure 4-7a and b, under PBS, 'learning the leveler' narrows the range of selection pressures applied against population members in both the MaxOnes and MaxInt cases. This is demonstrated more dramatically by a plot of the variance of selective values in each case[43,44], figures 4-7c and d.

---

[43] The *mean* selective value for PBS, for a population with or without local search, is simply the reciprocal of the population size, 1/N.

[44] It may be that using the variance of selective values (i.e. normalised valuations for PBS) or directly expressed/acquired valuations (RBS) may provide a way of characterising the ruggedness of a landscape over a population. (Hart, 1994) describes optimisation in terms of minimising a function, f(x), such that the globally optimal solution x* is given by min f(x*). Accepting that "numerical procedures can only produce approximate answers", he treats a problem as solved to a degree of ε-accuracy if a solution x can be found such that $f(x) <= f(x^*) + \varepsilon$. It is possible that a similar approach can be used to define the characteristic ruggedness of an operator set - for example, as the probability distribution (or just the variance) of valuations over the population at dynamic equilibrium. Worst case behaviour may then be guaranteed by adding a constraint that individuals, x, in the equilibrium population should conform to the constraint that $f(x) \leq f(x^*) + \varepsilon$, where x* is the local optimum within the population and ε gives the maximum deviation from that optimum.

a)

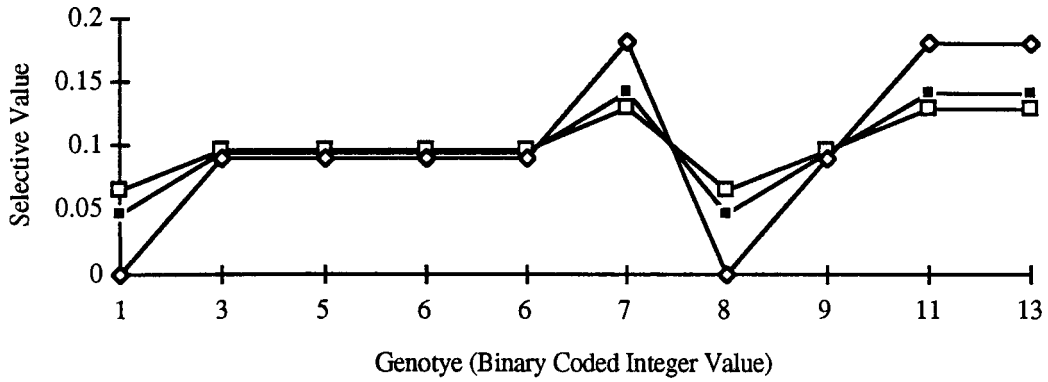

b)

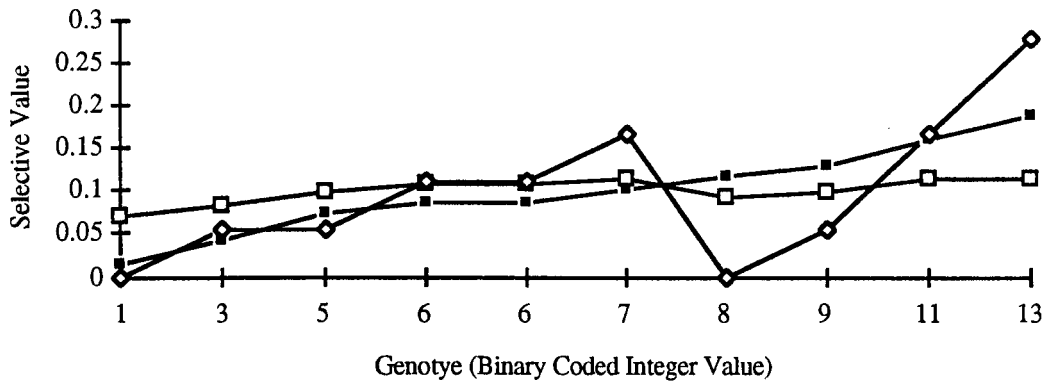

c)

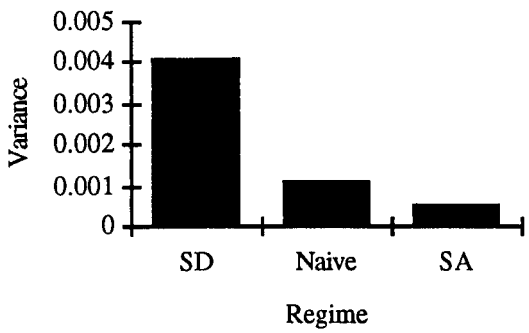

d)

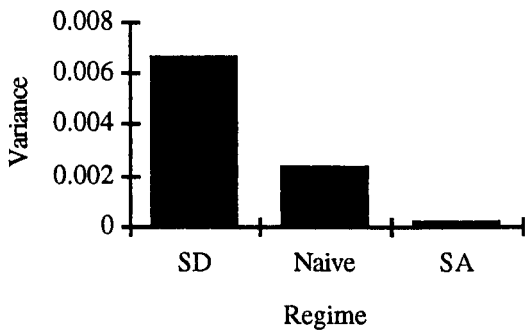

Figure 4-7: Proportional selective surfaces (and the variance of the selective va,ues over the search space) for the example population under SA learning and SD fault induction; a) (and c) MaxOnes; b (and d)) MaxInt. Direct valuation (solid square);valuation following SA learning (empty square) and SD fault induction (empty diamond).

For RBS, whilst the range of selective values remains the same, 'learning the rank transformational leveler' masks genetic diversity with phenotypic equivalence. For example, with MaxInt, the variance of expressed *phenotypes* is *reduced*, whilst the variance of directly expressed *genotypes* will tend to *increase*, over the no learning case. In the limiting case of hard TS, figure 4-8, genetically superior individuals are lost with respect to more effective learners. Where rank is respected, as in MaxOnes for the particular learning operator used here, there is no transformation of the selective surface.

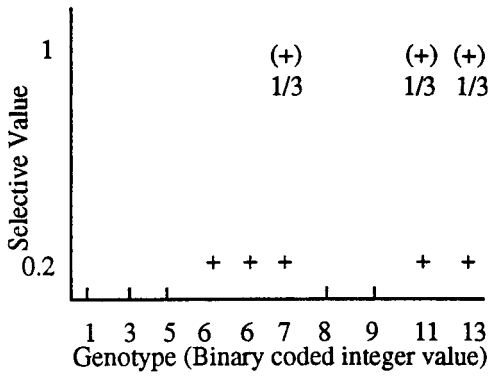

Figure 4-8: Truncation selective surface for the example population under SA learning and SD fault induction (the response is the same in this example) for the 4 bit MaxInt evaluation function.

4.5.5 Transforming Population Structure Through Inheritance.

I have already alluded to the transformation of population structure that results from the introduction of IAC. Where IAC is supported, the genotype corresponding to the phenotype acquired through learning is passed to the selected population, rather than the selected individual's inherited genotype. This has the consequence of transforming the population's genetic variance, as well as the evaluation received by each individual.

Recall how valuations on the 4 bit binary coded MaxInt function were transformed by single bit SA learning (figure 4-5b) so that only 8 distinct valuations were possible. Under IAC, the genotypes representing these 8 values are the only ones that may be passed to the breeding population. IAC, then, potentially restricts the amount of genetic variation in the breeding population.

4.5.6 Reaping the Benefit - The Interaction of IAC and Selection.

Having transformed the population structure through IAC, I now consider how the surface of selective value is affected.

Example 4.4. Just to tie up the ends, I complete the discussion of the simple 4 bit MaxInt problem by looking at the action of the previously specified truncation selection function on the original example population as transformed by the inheritance of SA acquired characteristics. See figure 4-9 below. Note that the genetic variation within the sub-population up for selection is limited to two distinct individuals - 14 and 15.

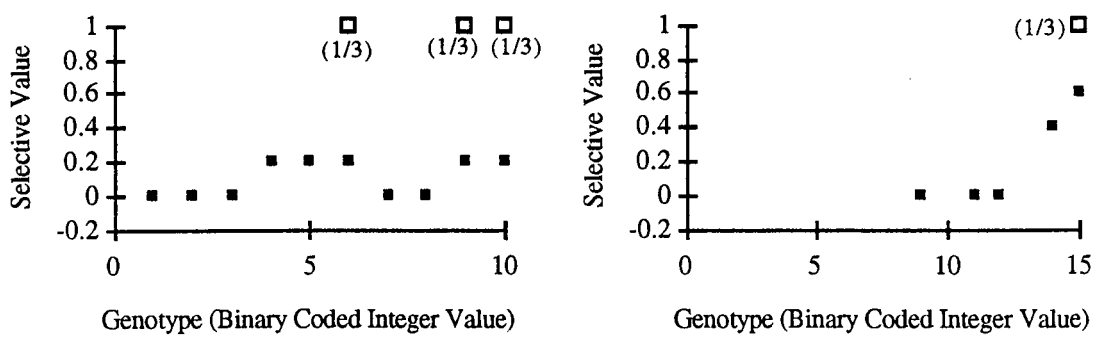

Figure 4-9: 4 bit MaxInt with SA learning and IAC: a) surface of selective value over the example population; b) surface of selective value for the effective population.

4.6 So What's the Damage? Costing Plasticity

I have shown how the introduction of learning may give rise to valuation equivalence of genetically distinct individuals. Where such equivalence resides, there is no way of using selection to distinguish between individuals that apply learning to acquire a particular trait and those that directly inherit it, and this may lead to the introduction of a significant amount of 'selective noise' particularly in RBS. It would thus be useful to be able to distinguish on grounds of valuation between individuals that inherit a trait and those that must acquire it. A biological example is the case of the Baldwin Effect for which the necessity of a learning cost has been demonstrated by (Mayley, 1996c) - the cost gives rise to the selection pressure necessary for the plastic trait to be driven out in favour of the appropriate fixed trait.

More generally, where the intention is to apply an operator under evolutionary (selective) control, a cost should be associated with that operator so as to allow the selection function to distinguish between individuals that make use of the operator and those that do not, all other things being equal. The operator will then only be applied where it is selectively advantageous and as such may be usefully placed under evolutionary control.

By introducing an operator cost, such as a cost for learning, the valuation landscape which has already been transformed through the application of the learning operator, is further transformed by the exacted learning cost, figure 4-10, (based on (Mayley, 1996c)). The costed valuation, $V_{g,c}$ is now given by:

$$V_{g,c} = v_g + b_g - c_g \tag{4.6}$$

where c_g represents the cost of learning, and $V_{g,c}$ replaces V_g in equations (4.2) and (4.3).

That is, the learning cost equals the difference in valuation between an individual that inherits a trait and a costed individual who has to learn it.

By considering the graph of expected evaluation for H&N (figure 4-10) it is easy to see how there is an implicit learning cost associated with learning in that experiment arising from the form of the valuation function.

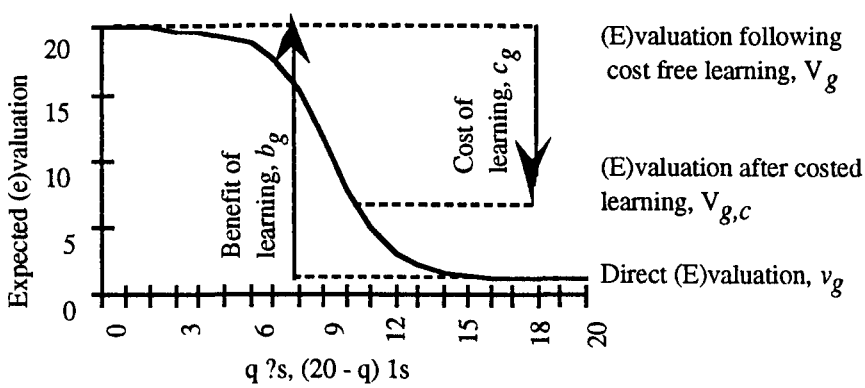

Figure 4-10: The cost of learning in H&N (based on (Mayley, 1996c)*).*

4.7 Filtering 'Spatial' Signals: "Landscape Processing"

In this section, I shall briefly introduce the notion of *landscape processing*, which applies the metaphor of image processing to valuation landscapes. In particular, I will suggest how plasticity and fault induction may be used to 'filter' a valuation landscape. The argument is informally presented, and relies on a simplistic two dimensional representation of the search space. The valuation landscape is visualised as a contour map, with thick contours signifying low valuation and fine contours high valuation, as in figure 4-11a.

It is now possible to consider filtering the 'valuation image' in a manner akin to image processing techniques. Such methods often make use of a 3 x 3 pixel 'mask' (the Moore neighbourhood in cellular automata terms) to generate the next state of the central pixel following filtering. The mask is passed over each image pixel in turn to filter the whole image. A low pass filter corresponds to the mask shown in figure 4-11b. The central pixel, after filtering, is set to the mean prefiltered value of the pixel itself and its 8 nearest neighbours. This is similar to an individual being evaluated according to the mean evaluation over the whole of a suitably defined plasticity neighbourhood. White noise ('speckles') is often removed from an image through a *median filter*, in which the central pixel is given the median value of those taken over the mask.

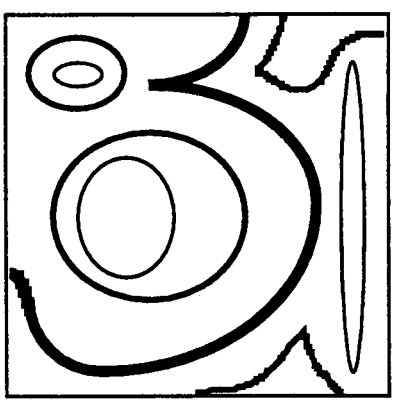

Figure 4-11a: An idealised 'landscape image'. Thick lines represent contours of low valuation, fine lines high valuation, over a search space structured by the genetic operators.

1	1	1
1	1	1
1	1	1

Figure 4-11b: An image mask as used in a low pass filter.

In hybrid GAs that employ within generation local search such as steepest ascent, learning corresponds to setting the focal 'pixel' to the *highest* value in the neighbourhood induced by local search, thus smearing any noise present, rather than removing it. Fault induction using steepest descent compares directly with using the *lowest* neighbourhood value (and 'blurring' of the image). Filtering through plasticity implements smoothing of the landscape by an individual and is thus *individual level filtering*.

Valuation surfaces may be further transformed by the use of particular selection functions as discussed previously. Since the selection function determines the range of allowable selective values over the population, selection function transformations of the evaluation surface are examples of *population level filtering*[45].

4.8 So Much for the Theory...

By using a simple worked example, I have discussed graphically the transformation of valuation landscapes by adaptive and maladaptive local search and the consequent effect on the selective surfaces induced by the particular selective conditions in use. Typically in PBS, adaptive local search (learning) smoothes the valuation and selective landscapes, whereas maladaptive local search (fault induction) sharpens them. A distinction has been made between two forms of selection and learning - proportional and rank based selection on the one hand, rank respectful and rank transformational learning on the other. These combine separately to produce two marked effects. First, there is the possibility that the range of selection pressures over the population is reduced (proportional selection, rank respectful learning); the limiting case of this is when genetic variance is masked through genetically different individuals acquiring similarly evaluated traits. Secondly, congenitally weak individuals may be hidden from the selection function at the expense of individuals with a higher direct valuation (rank transformational learning, either form of selection). The transformation induced by selection (particularly the truncation selection function considered herein) may also mask the range of direct objective values pertaining to population members, at least in so far as the differential selection of independent, 'fit' members goes. When the

---

[45] The image filtering metaphor may also be applicable to spatially structured populations. Specifically, if the population is arranged over a grid, the image mask may be taken as representing some structuring of the selective-reproductive stage of the EA.

inheritance of acquired characteristics is supported, not only is there a potential loss of phenotypic variation within the population, but also a corresponding loss in genotypic diversity.

In the following chapter, I shall illustrate the effects of surface transformation genetic operator choice and within generation local search in a series of GA experiments.

Chapter 5 - Characterising Evolutionarily Stable Populations

"I will now explain the process of - the name is blotted, I'm sorry to say. It will be illustrated by a number of - of....It seems to be either 'Experiments' or 'Specimens'-"
"Let it be Experiments," said the Emperor. "We've seen plenty of Specimens."
"Certainly, certainly!...We will have some Experiments...Our First Experiment requires a Machine. It has two knobs - only two - you can count them, if you like."

L Carroll, from "The Professor's Lecture" in Sylvie and Bruno Concluded.

5.1 Introduction

Typically, in a classically posed optimisation problem, the *evaluation* landscape over the phenotypic search space will be fixed. In this chapter, I set out to test the conjecture offered in 3.4 that a population evolves towards a region of the search space with a characteristic ruggedness. As discussed in chapter 4, various *valuation* surfaces for a given evaluation function may be achieved through the use of different development functions. Where within generation local search is applied to individuals, this may be thought of as creating a new development function. Consequently, the valuation landscape over a genotypic search space where local search is applied represents a transformed version of the valuation surface over that search space when there is *no* local search. The question that now arises is "how are the evolutionary trajectory of a population, and its equilibrium resting point (if indeed one exists), affected by the introduction of a within generation, local search operator?" In this chapter, I shall present two experiments that investigate this question empirically.

5.2 Experiments in Artificial Evolution

In 4.5.2.3, I suggested that for cases of local search where rank order is preserved, the resulting transformation of the surface of selective values for proportional based selection could be easily visualised: where local search is adaptive (*learning*) the landscape may be thought of as being *smoothed* (i.e. the range of selective values awarded is narrowed); where maladaptive local search is applied (*fault induction*) the transformed valuation and selective surfaces are rather more rugged than in the simple case (i.e. they are *sharpened*). For rank

based selection, preservation of rank order implies that there will be *no* transformation of selective values even if individual evaluations are affected. The question now arises as to the extent to which these essentially qualitative judgements can help us predict the equilibrium resting place of a population whose individuals may apply local search? In the previous chapters, I conjectured that the population will evolve towards a region of the valuation/selective surface to which it is adapted. Let us for the sake of argument denote the ruggedness of that region under some appropriate measure, as A. The convergence properties of the population, as regulated by the mutation-selection balance, are such that *the particular operator set is adapted to a region of the valuation/selective surface with a characteristic ruggedness* (e.g. A). Now, local search transforms the valuation/selective surface. When local search is applied, the population should thus move to the region of the transformed surface with ruggedness A, figure 5-1. I would thus expect that a population employing local search will settle in a region of the search space for which the corresponding *original* surface is a) more rugged than A for adaptive local search; b) smoother than A for maladaptive local search.

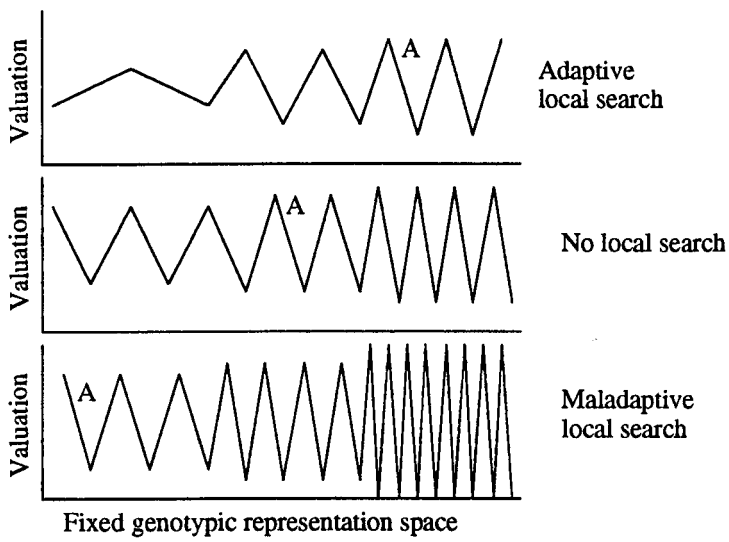

Figure 5-1: Transforming the valuation landscape through local search.

Under the view where populations flow towards an optimum within whose basin of attraction they lie, local search acts so as to transform the size of the basin of attraction. For adaptive local search, the basin of attraction of the local optimum with the best evaluation around its edge compared to neighbouring optima will have its basin of attraction extended. For fault induction, the basin of attraction of the optimum taking the lowest evaluation around its edge, again compared to other neighbouring optima, will have its basin widened.

In the following sections, I shall illustrate by experiment several of the points raised above. In the first case, using the well known NK landscapes of (Kauffman, 1993) and described in 5.3 below, I show how the application of adaptive and maladaptive local search operators influences the region of the search space in which an initially random population eventually settles (5.3.2). In a second experiment, using a novel evaluation function ("The Tower of Babel") I allow each individual to evolve the evaluation function on which it is to be evaluated (5.4). By varying mutation rate and strength of selection in the case of a simple GA with rank based selection, I demonstrate how global optima are *not* exploited by small populations (5.4.1). In addition, the evolution of genetic canalisation (tolerance to mutation), and how it is influenced by local search, is demonstrated (5.4.2).

5.3 The NK-landscapes

An increasingly popular family of test functions within the evolutionary computation community are Kauffman's NK landscapes (Kauffman, 1993)[46]. The NK landscapes are defined as follows: for a genome of length N loci, each locus contributes additively to the overall evaluation of the individual. Each locus is epistatically connected to K others (K_{max} = N-1), chosen either as neighbouring loci on the string, or as random loci. Since the effect of epistasis on locus evaluation is assumed to be unknown, 2^{K+1} random evaluations (0..1) are associated with each locus - the actual evaluation of a locus within a given individual then depends on the neighbourhood state (i.e. the state of the K epistatically connected neighbours) and the total evaluation of the individual is then given by:

$$F = \frac{1}{N}\sum_{i=0}^{N-1} F_i \tag{5.1}$$

[46] There is a considerable body of literature growing up around the NK landscapes (for a recent review, see (Altenberg, 1997b)). For example, the analyses of (Stadler, 1995; Stadler & Happel, 1995) and a discussion in terms of NP-completeness by (Weinberger, 1996). Recent reports in the GA literature include (Smith & Fogarty, 1996; Harvey, 1992; Thompson, 1996; Mayley, 1996a).

where F_i is the evaluation of locus i set to a random real value in the range 0...1 during the creation of a particular landscape[47].

5.3.1 Estimations of Ruggedness.

Since the NK landscapes are statistically defined, there is a pertinent methodological question that needs to be addressed: when reporting evaluations, should 'raw' values or ones normalised against the maximum value for that particular landscape be used? Given that large search spaces may be easily generated, it is not reasonable to expect that globally optimal evaluations will be available so normalising to maximum landscape values is not practical. Normalising *post hoc* against the greatest evaluation recorded during a particular run is possible, but this restricts the information available for comparing runs over different landscapes of the same order. In general, I will report on mean raw (unscaled) evaluations.

One of the aims of this chapter is to demonstrate that evolution may be guided towards regions of the search space with different degrees of evaluation ruggedness, according to the operator sets employed. Whilst it was suggested in 3.2.3 that the runtime operator correlation coefficient offered a useful measure of local ruggedness, cheaper alternatives are possible, such as simply the range or variamce or evaluations over a volume of the search space. In accord with the view of GAs as doing function optimisation, reported results often concentrate on instantaneous and time averaged best of generation and population mean evaluation or expression. Occasionally, worst of generation results are provided to give

---

[47] Notable extensions to the NK model include: a) (Kauffman, 1993), who implements a temporally fluctuating landscape by constructing the evaluation table over 2^{K+W+1} bits rather than 2^{K+1} bits, allowing 2^W possible environmental states which may be cycled through, for example; b) (Mayley, 1996b) uses a table containing L x N integers to define N distinct 'L-neighbours' of an individual. At the start of a run, for each of the L-neighbours, L values (between 0 and N inclusive) are chosen at random. L-neighbours of an individual are then realised by flipping within the individual the bit positions specified by the L values defining the particular neighbour. In this way, a one-to-N genotype to phenotype map may be provided; c) an extension from the haploid to the diploid case is given by (Bergman, Goldstein, Holsinger & Feldman, 1995), who liken the 2 allele diploid case (which gives 3 possible pairings) to the 3 allele haploid case; in the resulting evaluation table, each locus is associated with 3^{K+1} possible evaluations.

information on lower bound performance, but this is not generally the case. The use of statistical measures of variance or standard deviation is uncommon, although where these results are offered they tend to be measures of variance/s.d. over a number of test runs. This does not reveal anything about the spread of values over a *population*, since the variance of evaluations across a population is rarely, if ever, given[48]. However, measures of variance do not necessarily communicate useful information about the *range* of values, which may be important where worst case behaviour, for example, is required. Some techniques have been reported in the literature for considering the distribution of evaluations about an optimum, although they haven't been widely taken up: (Thompson, 1996) shows the mean evaluation of single mutant neighbours of a locally optimal individual discovered through evolution is greater than one may expect about a randomly discovered local optimum; (Parmee, 1996) comments on the standard deviation of evaluations of 'perturbations' about a solution, implying that evolution drives a population towards regions of the search space where this value is relatively low (i.e. the valuation surface is relatively smooth).

Returning to the NK landscapes, Kauffman suggests that increasing K for a given N increases the ruggedness of the evaluation landscape. By ruggedness, he means that for the single bit mutants of an individual there is a wide variance in evaluations. For a smooth (*correlated*) landscape, one expects small changes in evaluation over the mutation neighbourhood. The measures I shall employ to estimate the local ruggedness of the surface over a particular population extend this view to incorporate outliers: for each individual in the population, the best, worst and mean evaluation taken over the individual and its Hamming 1 neighbours will be recorded. This gives an impression of the local ruggedness of the evaluation surface as a spread of values in the vicinity of the population. The position of the mean score, compared to the best and worst evaluations, offers further clues to the structure of the landscape in that region. In addition, the best awarded evaluation over the entire run will be reported, as will the best of the final generation. The best ever serves as a worst case estimate of the global optimum, the best of final generation identifies the optimum discovered at equilibrium. For the local search algorithms, the mean *direct* evaluation (that is, the evaluation of an individual *without* the benefit of local search) will also be reported.

To give some impression of the distribution of evaluations for a random population over the NK landscapes, figure 5-2 shows a range of evaluation measures for 1000 randomly

---

[48] In the case of mutliple runs, there is correspondingly no mention of the mean and variance of the variance (*sic*) of evaluations taken across the populations of each run.

sampled individuals and their single Hamming unit neighbours. As claimed, for increasing K, the landscape becomes increasingly rugged (i.e. there is a widening range of evaluations over the Hamming one neighbourhood).

One particularly important feature of the NK landscapes is that for increasing K the distribution of optima throughout the search space is such that the highest peaks tend to congregate close to each other. This suggests that there are regions of varying degrees of ruggedness over the search space as a whole (Kauffman, 1993).

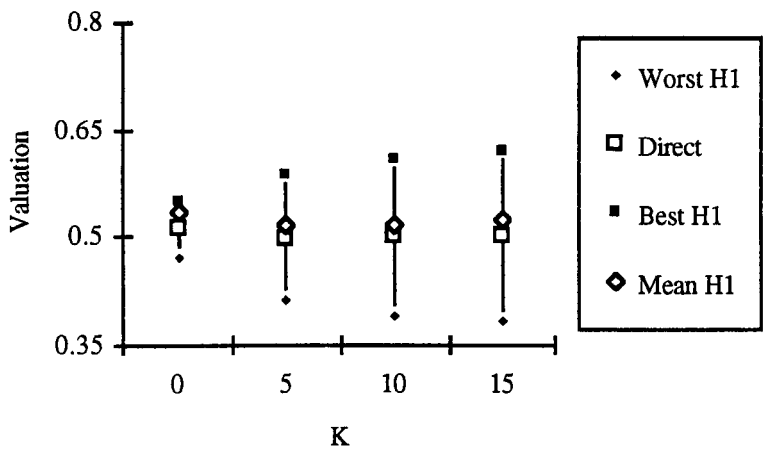

Figure 5-2: Mean evaluation measures for 1000 randomly sampled individuals and their single Hamming unit neighbours taken over K for N = 20:

5.3.2 Comparing the effectiveness of within generation local search strategies over the NK-landscapes

In this experiment, I shall compare the evolution of an initially random population of individuals over the NK landscapes using three different styles of local search, with and without the direct inheritance of acquired characteristics (IAC), with a simple GA (SGA) as a control. This differs from (Thompson, 1996) in that the aim is to see whether the different styles of within generation, local search influence the region of the search space to which a population converges. Another difference comes from the strength of selection: in accord with Wagner's theory of canalisation (Wagner & Booth, 1997), Thompson found that tolerance to mutation was highest where selection was weak. In this experiment, I shall use

quite strong selection in an attempt to rapidly converge the population. However, it is expected that phase 2 of the search will identify a region of the space where operator-evaluation correlation is high. Consequently, the ruggedness of the evaluation surface over the final generation, converged population should depend on the particular operators applied.

In each of the modified cases, exhaustive search of the evaluated values over the Hamming one neighbours is carried out. In the 'learning' case, GA-SA1, an example of single bit steepest ascent learning, the highest evaluation over the inclusive neighbourhood (i.e. the highest evaluation returned from the individual and its Hamming 1 neighbours) is returned as an individual's evaluation. In the second case, the 'fault induction' algorithm, GA-SD1, the *worst* evaluation over the neighbourhood is used; that is, the evaluated value is awarded according to a steepest descent algorithm. In the third case, GA-ASD, the assigned evaluation is the mean of the direct, highest and lowest recorded evaluations. In addition, two regimes supporting IAC are investigated, one using adaptive local search, the other using maladaptive local search (GA-ISA and GA-ISD respectively). In these cases, the genotype giving rise to the evaluation discovered by local search and used for selection is transmitted, rather than the genotype inherited from its parents.

The experiment is carried out over the NK class of landscapes, using Genesis 5.0[49] and Jones' NK generator[50]. GA parameters were set as follows: N = 20, with randomly selected neighbours; population size = 500; crossover and mutation rates 1.0 per individual and 0.0025 per bit respectively; linear rank selection with rank minimum 0.0 (i.e. *strong* selection). Reported results represent the mean of 10 runs (a different landscape using either random or next-door neighbours was used each run, with different landscapes across each condition), each lasting 100 generations except for the most rugged environment, K = 15, which was taken over 200 generations. Convergence is reported in terms of *bitwise convergence*, which represents the mean proportion of the dominant allele in the population at each locus, averaged over all loci. In all cases, the populations converged to a considerable degree at a single local optimum (table 1). For direct comparison of the effectiveness of the different regimes in locating optima, identical initial populations should be evolved over the same set of landscapes.

---

[49] Currently available from ftp://www.aic.nrl.navy.mil/pub/galist/src/genesis.tar.Z

[50] Available from http://www.santafe.edu/~terry/

The results of evolution over 4 landscapes of varying degrees of ruggedness are shown in figure 5-3. The best ever and best of final generation valuations as offered to the selection function are shown in table 5-1 - whilst these latter results are shown for completeness, emphasis is placed on the *range* of valuations over the population (figure 5-3) rather than the *absolute* valuations given in the table.

- For K = 0, all populations locate the optimum, and converge around it, with the exception of the population undergoing fault induction, GA-ISD. The apparently spurious result for this latter regime (i.e. the final best being lower the best ever result) is a consequence of the population converging on an individual at the bottom of some valley, as opposed to around the summit of the optimum. The apparently high valuations for GA-ISA, and low valuations for SGA (figure 5-3), compared to the others is presumably an artefact of the landscapes used during that particular run; the similar distribution of recorded values suggest that a similar region of the landscape has been located. Looking at the convergence properties, all regimes converge to the optimum except for regime GA-SA1. In this case, with deterministic, cost free learning, there is a degree of hiding of genetic variation through the local search operator (all Hamming 1 neighbours of the optimum sequence receive the evaluation of that sequence). The failure of the population in this regime to fully converge for any K is similarly explained;

- For K = 5, there appears to some evidence that the different regimes are settling in different regions of the search space. For the fault induction regime, GA-SD1, the range between the best and worst values over the Hamming 1 neighbourhood is considerably less than for the other regimes. GA-ASD also seems to have settled in a relatively smooth region of the landscape: in particular, not only is the range between best and worst over H1 reduced, but also the mean evaluation over H1 is close to the best. Although the adaptive search regime, GA-SA1 does not appear to have located a more rugged region of the landscape than the control condition, SGA, the fact that the population is maintaining a certain degree of genetic variation is revealed by the fact that the population is not genetically converging on the directly (innately) optimal genotype in the population (the direct evaluation is below that of the best H1 evaluation). For the two regimes capable of transmitting the effects of local search to their offspring, the adaptive regime GA-ISA completely converges the population, but regime GA-ISD fares only as well as random search;

- As the landscape becomes increasingly rugged, K = 10, regime GA-SD1 continues to locate smoother regions of the landscape than SGA, whereas GA-SA1 has now identified a rather more rugged region (the range between best and worst valuation over the Hamming 1 neighbourhood is now greater than for SGA). Also note that across all the regimes, the population is now identifying possibly globally optimal individuals during the search, but is not converging on them (table 5-1). Regime GA-ISA appears to be locating better optima than the other regimes, but since global optima are unknown the higher reported 'best' scores may just be artifacts of the random way in which landscapes are generated;

K	Evaluation	SGA	GA-SA1	GA-SD1	GA-ASD	GA-ISA	GA-ISD
	Best Ever	6.38	6.72	6.45	6.52	7.11	6.08
		(0.41)	(0.54)	(0.49)	(0.67)	(0.37)	(0.48)
0	Final Best	6.38	6.72	6.45	6.52	7.11	3.52
		(0.41)	(0.54)	(0.49)	(0.67)	(0.37)	(0.48)
	Convergence	0.999	0.988	0.999	0.999	1.00	1.00
	Best Ever	7.71	7.79	6.20	7.07	7.92	5.57
		(0.26)	(0.12)	(0.15)	(0.23)	(0.21)	(0.22)
5	Final Best	7.70	7.78	6.14	7.03	7.92	5.38
		(0.26)	(0.13)	(0.13)	(0.24)	(0.21)	(0.16)
	Convergence	0.999	0.987	0.999	0.999	1.00	0.565
	Best Ever	7.46	7.75	5.41	6.71	7.87	5.09
		(0.29)	(0.14)	(0.20)	(0.17)	(0.17)	(0.11)
10	Final Best	7.32	7.66	5.34	6.65	7.75	5.00
		(0.19)	(0.18)	(0.14)	(0.21)	(0.27)	(0.14)
	Convergence	0.999	0.982	0.999	0.999	1.00	0.542
	Best Ever	7.18	7.62	4.88	6.32	7.76	4.75
		(0.22)	(0.18)	(0.13)	(0.18)	(0.19)	(0.08)
15	Final Best	7.00	7.47	4.81	6.22	7.55	4.72
		(0.32)	(0.21)	(0.14)	(0.17)	(0.13)	(0.09)
	Convergence	0.999	0.985	0.999	0.999	1.0	0.536

Table 5-1: Best ever, best of final generation, and bitwise convergence measures for converged populations over landscapes of varying degrees of ruggedness. Standard deviations of evaluations, given to 2 d.p., are bracketed. Convergence properties of populations over runs within a condition were uniform.

- Finally, for the most rugged landscape, K = 15, GA-SD1 again locates a smooth region of the landscape, as does GA-ASD (although to a lesser degree); GA-SA1 moves to a rugged region of the space.

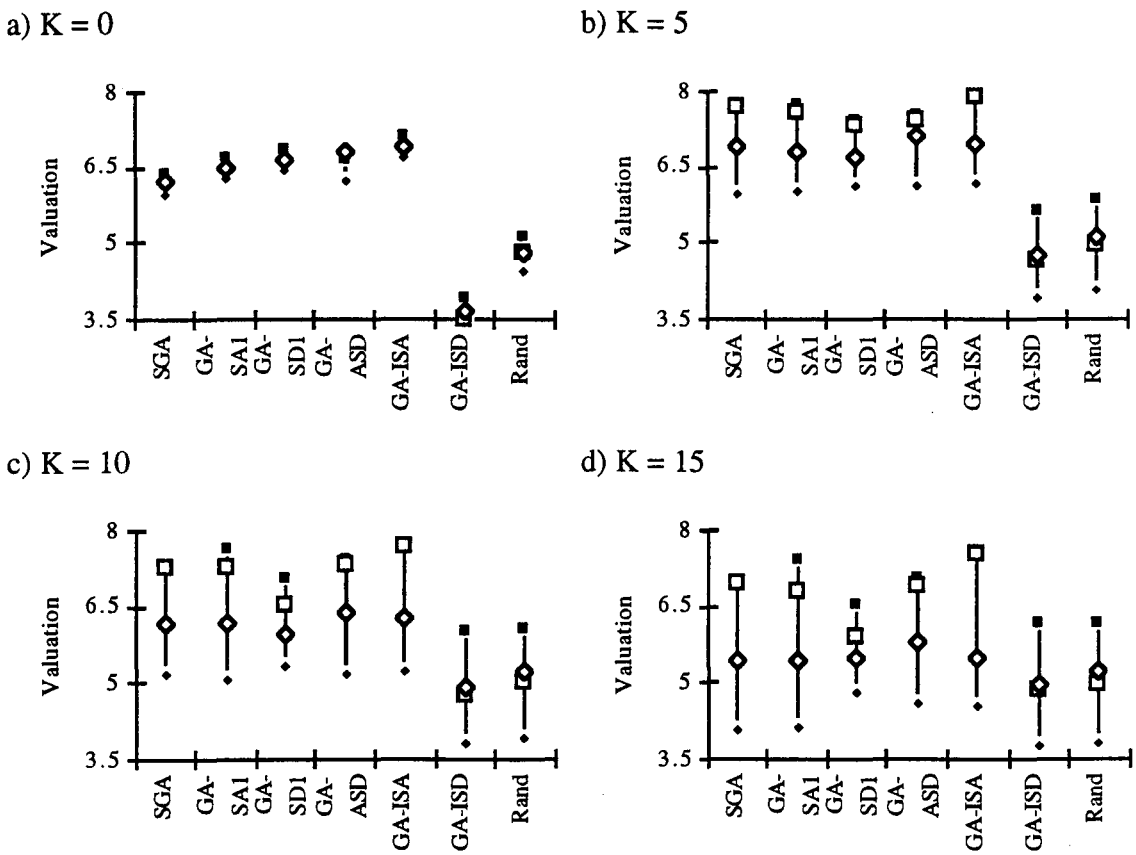

Figure 5-3: Performance measures for converged populations in landscapes of varying degrees of ruggedness. Solid symbols - best/worst valuation over H1 of each individual; empty square direct valuation; empty diamond - mean valuation over H1.

As I suggested in 3.4 above, the population in phase 2 of the search is likely to identify and exploit regions of the search space to which its operator set is adapted. In figure 5-4, I plot the bitwise convergence and the population mean evaluation during evolution of regime SGA in a rugged landscape (K = 15). I suggest that the three phases of search are identifiable - phase 1, until roughly generation 70, in which the population begins to converge but the population mean is unaffected; phase 2, in which the population continues to converge and the population mean begins to improve (exploitation of the an appropriately identified region of the search space); phase 3, starting at generation 130 or so, in which the population

converges to a stable limit and evolves through a region of similar mean valuation until another optimum is identified at around generation 170 and the population rapidly converges around a single point. Note that in the later generations, as complete convergence is approached and the population behaviour begins to approximate an adaptive walk, evolution is still possible (the rapid stepwise increase in mean evaluation around generation 170 as the population finally converges).

The steps in the mean evaluation curve appear to suggest the existence of 'metastable' population states, such as those that occur using the Royal Road evaluation function and as described in (van Nimwegen et al., 1996). Using that function, the metastable states occur when the population identifies an optimum located within a neutral network. Evolution is still possible along these neutral networks (the population still explores the genotypic search space) even though there is apparently nothing going on if the mean evaluation curve is to be believed. In the literature, neutrality is described in terms of connected genotypes (on a Hamming graph) with equivalent evaluations (*neutral networks*) or in 'coarser' terms as the equivalence classes that contain individuals bearing identical evaluations (Barnett, 1997). That is, neutrality is discussed relative to an evaluation landscape. In the general case of the NK landscapes, where there are no neutral networks, it would be interesting to know whether evolution along a pathway of equivalent population mean evaluations reflects traversal by the population through regions of identical local ruggedness. In the context of this thesis, the issue of neutrality may thus be raised in the context of landscapes other than the evaluation landscape, although this will not be pursued further herein:

- by considering the *stability* of populations at equilibrium, it may be possible to identify neutral networks on the population mean evaluation landscape[51] or in the sense of connecting regions of the landscape with an identical local ruggedness;

- neutrality may exist in terms of selective values of parents and offspring; for example, under RBS, each individual in the current population may beget an individual whose rank position in the descendant population is the same as the parent's rank in the current population, although each have received a different evaluation.

---

[51] That is, neutral networks defined on the population mean evaluation landscape in terms of 'neighbouring' *populations* with the same *mean* evaluation.

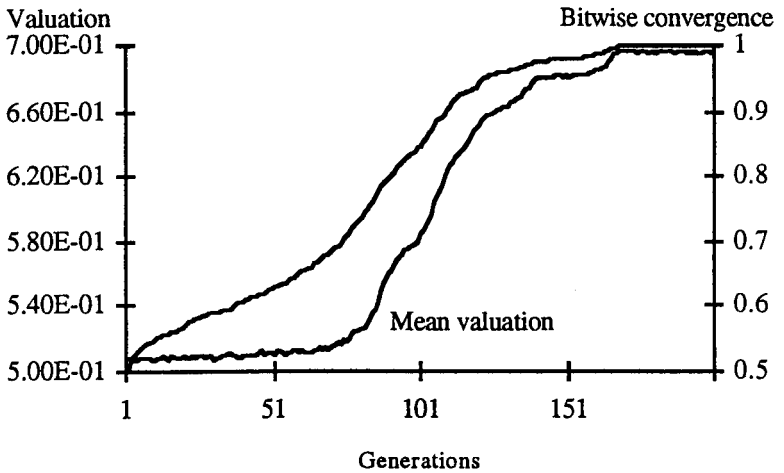

Figure 5-4: Bitwise convergence and population mean evaluation for SGA, K = 15.

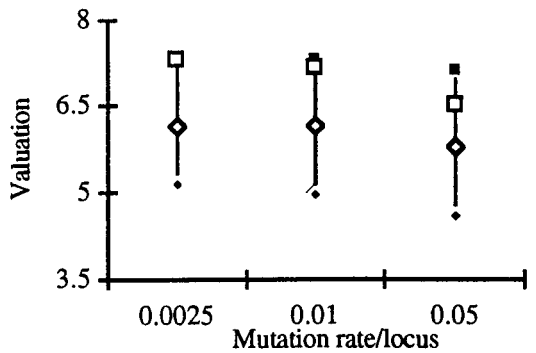

Figure 5-5: Equilibrium evaluation scores, K10, generation 100 for m-rate 0.0025, generation 250 for others.

	0.0025	0.01	0.05
Best Ever	7.46	7.44	7.37
Final Best	7.32	7.27	7.32
Convergence	0.999	0.995	0.945
Eqm at gen:	80	85	211

Table 5-2: Equilibrium population statistics, K10, generation 100 for m-rate 0.0025, generation 250 for others

Finally, the effect of increasing the mutation rate on the population search in shown in figure 5-5 and table 5-2. As mutation rate increases, the population best, mean and worst all fall slightly, along with the degree of convergence at equilibrium. In addition, the time to convergence also increases with the mutation rate. There are two ways of interpreting the fall in H1 evaluation measures, figure 5-5. Either, a region of different ruggedness is being identified; or a region of similar ruggedness is identified but the mutation-selection balance is pulling the population mean down from the top of the final generation optimum. (Note that both explanations may valid, although being complementary only one will be valid on any particular occasion). Comparing convergence limits, the population converges to a lower degree for increasing mutation rate. Since the best evaluations in the final generation are close across the three mutation rates chosen, it would thus seem likely that similar regions are

identified but the mutation-selection balance is holding the population just below an optimum. Again, it would be possible to validate this claim by tracking the region of space explored by each regime using identical evaluation surfaces.

5.4 The Evolution of Genetic Canalisation - 'The Tower of Babel.'

Canalisation refers to the ability of an individual to tolerate environmental or genetic (mutational) change - that is, it describes the sensitivity of an individual to change. (Wagner & Booth, 1997) suggests that genetic canalisation is most evident where the degree of genetic variation is high and selection is weak. In this section, I shall introduce a novel evaluation function - The Tower of Babel - which allows us to investigate the ability of a GA to discover canalised traits through the adoption of appropriate developmental maps.

By defining a suitably parameterised family of evaluation surfaces, it is possible to construct self-adaptive, evolvable evaluation functions (Lund & Parisi, 1994). Such functions may also be interpreted as modeling an evolvable development function (i.e. an evolvable genotype-phenotype map). For example, in the NK landscapes, which for increasing, non-zero K have the required property of increasing ruggedness, one might imagine the following: that the usual genotype-evaluation map ($\mathbf{G} \rightarrow R$) has been replaced with a genotype-phenotype map ($\mathbf{G} \rightarrow \mathbf{P}$, where $\mathbf{P}$ is defined over the real numbers) that corresponds to the original genotype-evaluation mapping; and that the evaluation function is now given by the identity map ($\mathbf{P} \rightarrow R$). Evolvable evaluation/development functions also bear a close resemblance to functions that are defined over *variable length genotypes*[52]. Typically, an increase in length will open up as yet unexplored regions of the search space and allow for potentially open ended evolution.

The idea is, then, that individuals are allowed to evolve the landscape on which they are to be evaluated. For example, it is possible to parameterise evolution over a family of NK landscapes in the following way. At the start of the run, N landscapes would be generated, one for each possible value of K, and these remain fixed throughout the course of a run.

---

[52] For example, (Harvey, 1992) describes an experiment involving a length changing genetic operator and a multiobjective evaluation function, in which one component awards an NK evaluation, the other encourages an increase in length.

Individual modifiers are then allowed to evolve that code for the landscape - i.e. the value of K - each individual is to be evaluated on. A population would be expected to evolve towards a region of the search space where the developmental map (i.e. K) - and hence the *valuation surface* - is adapted for the operator set applied. Evolution of the evaluation function *qua* developmental map evaluated according to an identity function then provides a way of investigating the evolution of (genetic) canalisation.

However, the uneven distribution of optima over the NK landscapes introduces an element of noise that may complicate our understanding of the observed system's dynamics, at least during the initial stages of this investigation. For this reason, I shall define a rather simpler parameterised family of unitation functions[53]. In the resulting experiment, I demonstrate that different operator settings of a simple GA are likely to move an evolving population to the region of the search space to which the *operator set* is adapted, by independently evolving the evaluation function each individual is rated on.

The particular family of evaluation surfaces used in this experiment draw inspiration from the MaxOnes evaluation function, which provides the basis for much theoretical work because it is so easily understood. The general evaluation function is given by:

$$\texttt{if}\ (\mathrm{H} \geq \mathrm{cutoff})\ \texttt{then}\ \mathrm{evaluation} = 0 \tag{5.2a}$$

$$\texttt{else}\ \mathrm{evaluation} = (L^2/cutoff)*(1-(\mathrm{H}/cutoff)) \tag{5.2b}$$

where L is the length of the target sequence; *cutoff* is the individually evolved modifier gene, length $I = \lceil \log_2 L \rceil$ bits (giving overall string length $L + I$) which binary code an integer bounded $1..L$, and H is the Hamming distance of the individual being evaluated from 1. If L is further constrained to be a power of 2, the calculation of I is further simplified. A particular evaluation function from within the family may then be strictly identified as an $(L + I)$ bit Tower of Babel function. For convenience, and given the definition of I as a function of L, this may also be termed an L bit problem.

For $L = cutoff$, equation (5.2a) reduces to:

$$\mathrm{evaluation} = L(1-H/L) = L - H \tag{5.2c}$$

which is the MaxOnes evaluation function.

---

[53] Recall that for unitation functions, the valuation and evaluation surfaces are identical.

The population mean evaluation, $\bar{e}$, is then given by:

$$\bar{e}N = \sum_{L=0}^{N-1} e_{c_i}(g_i) \tag{5.3}$$

where $e_{c_i}(g_i)$ gives the appropriate evaluation of individual g_i with *cutoff* c_i.

The motivation behind the Tower of Babel is to provide a function with obviously distinct regions of ruggedness about a single optimum (figure 5-6). These regions may be chosen between by evolution according to the stability of the population around the optimum, and within each region, given the current operator set. The long term goal is to predict stable, equilibrium distributions over the modifier gene according to the current operator choice.

Figure 5-6 shows some sample surfaces for the (20+5) bit problem. The area under each curve is the same (specifically, the area under each curve equals L/2)[54,55]. For two individuals with evolved *cutoff* values of c_1 and c_2, their evaluation is equivalent when they are both at a distance H_{eq} from the optimum, such that:

$$H_{eq} = \frac{c_1 c_2}{c_1 + c_2} \tag{5.4}$$

As the quasi-species approaches the optimum, decisions must be made (selected) about which curve the population is to be evaluated on. For the constrained case of a highly converged quasi-species, the decision about which of the two curves to be evaluated on is likely to be made when the master sequence is in the vicinity of H_{eq}. It is not yet clear whether this decision is made on the basis of population mean evaluation, effective evaluation or the marginal evaluation of the modifiers themselves.

---

[54] A similar, and simpler, family of curves, that may be more appropriate for experiment and analysis using proportional selection, is the following:

if H ≥ *cutoff*, eval = 0; else eval = (L + 1 -*cutoff*)(1-H/*cutoff*)

For two individuals with a different *cutoff*, c1 and c2, their eval(c1) = eval(c2) at H_{eq} = c1.c2/(1 + L).

[55] The normalised MaxOnes evaluation function, $e(g_i) = 1 - H/L_i$, where L_i is the length of individual g_i, may be suitable for use with a variable length genotype. Presumably, an increase in length will be favoured as a result of individuals close to the optimum suffering a *lower* drop in evaluation for a single mutation.

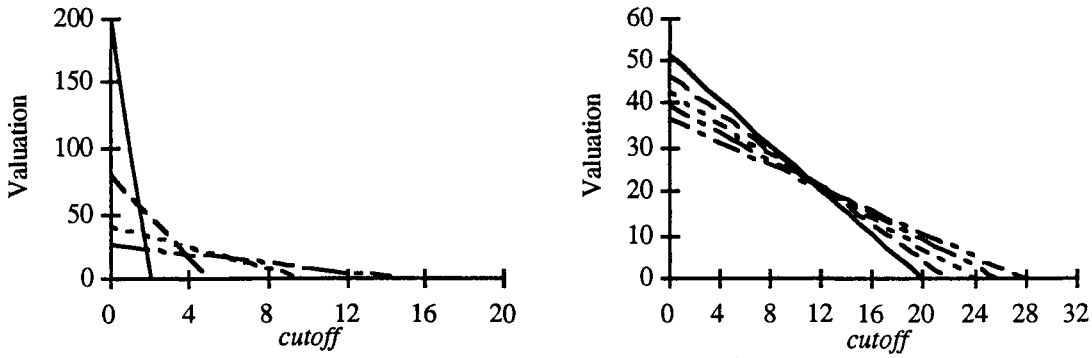

Figure 5-6 a) the 20 bit problem for 'cut-off' 2, 5, 10, 15; b) the 32 bit problem for cutoff 20, 22, 24, 26 and 28.

In the Tower of Babel, the interpretation placed on the modifier gene is very specific; in the first place, it is coded differently to the targeting part of the genome (binary coding - although Gray coding could have been used - in contrast to the literal coding of the targeting bits). Secondly, it is interpreted as a modifier gene in that the role it plays in the evaluation function may be regarded as tuning the ruggedness of the evaluation function over the L targeting bits through a developmental switch. The function was designed so that it encourages this particular interpretation, specifically a problem with 2^l regions of distinct ruggedness over the search space. The modifier is seen as identifying each region uniquely, with mutations to the modifier effectively introducing a jump between regions of different ruggedness (as measured over the L targeting bits) in the vicinity of the optimum. Although there do exist considerable regions of neutrality within the search subspace covering the targeting bits, introduction of the developmental modifier in itsef does not induce associated higher dimensional neutral networks to the search space[56]. This is in contrast to modifiers such as mutation rate modifiers which do include significant neutral networks within the expanded evaluation landscape.

5.4.1 The Tower of Babel and a Simple GA

Each experiment was carried out over 10 runs using Genesis 5.0, with parameters set as follows (unless otherwise stated): an initially random population of 100 members, targeting

[56] Neutrality exists mainly in the plains where evaluations are 0. There may also be a small number of neutral mutations within the modifier, specifically for individuals in the Hamming band about the optimum which represents an integer solution to c1.c2/(c1+c2), where c1 and c2 are distinct integer values of *cutoff*.

length (L) 32 bits, 5 evaluation parameter bits; two point crossover with crossover rate 0.6 per individual, and a fixed m-rate of 0.01/bit (0.05 bits/rate gene) for the function parameter bits. Linear rank selection is used so that there are no scaling problems. Experiments are run over a range of selection rank minima and targeting sequence m-rates.

Now, by assuming that one result of selection is to maximise the mean evaluation of the population, and maintain it over time (i.e. locate a stable solution) we can see how 'pointy' a single peak in an otherwise flat evaluation landscape the population will sit on. A high degree of canalisation will be reflected by a high value of *cutoff* (smooth landscape), a low degree by a low value for *cutoff*. The expectation is that a well spread quasi-species (resulting from relatively weak selection or a high mutation rate) is required to facilitate canalisation (Wagner & Booth, 1997).

In figure 5-7, I plot the evolution of the population mean *cutoff* for two conditions: a) over several (fixed) m-rates for the same rank minimum; b) over several rank minima for a single, fixed m-rate. The first thing to notice is that in each case the mean *cutoff* initially climbed to about 27. This is to be expected - in figure 5-6b, for an initially random population in which the expected distance an individual is away from **1** is 16 bits, the best evaluation is received by individuals evaluated on surfaces with a greater cutoff. It is only as the population starts to close the gap to **1** that evaluation surfaces of *lower cutoff* are favoured.

What about the long term behaviour of the evolved *cutoff*? For the runs over varying m-rate, and fixed, weak selection, figure 5-7a, an increase in the m-rate increases the equilibrium value of *cutoff* - where the m-rate is higher, the population prefers smoother regions of the evaluation space. That is, a developmental map favoring genetic canalisation is the result of increasing the mutation rate. Mean values for the modifier gene at equilibrium are given in tables 5.3a and 5.3b. In all cases, actual integral values for the gene in each run were within +/- 1 of the integer value of the reported mean.

For the runs over the same m-rate but differing strengths of selection, figure 5-7b, increasing the strength of selection forces a lower *cutoff* - for stronger selection, the population is able to converge to a rather higher degree and so a higher optimum in an albeit more rugged region of the space is favoured. In fact, to all intents and purposes, the population *does* identify the global optimum where selection is anything but weak. Only under weak selection does canalisation occur to any significant degree. One might expect a similar behaviour from

an increase in the crossover rate (i.e. increased convergence for increasing crossover rate means the population can survive in a more rugged region of the search space).

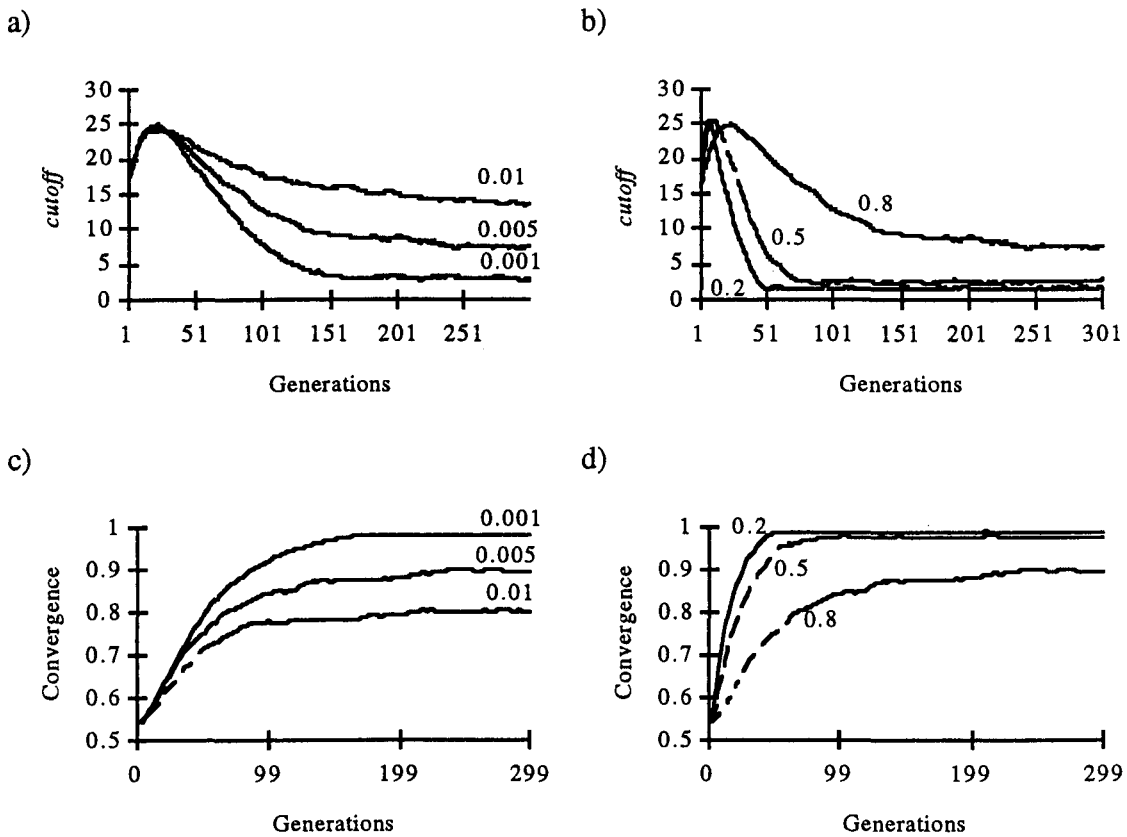

Figure 5-7: a) over targeting loci m-rate {0.001, 0.005, 0.01}, rank minimum 0.8, cutoff decreases for decreasing rate, b) over selection strength (rank minimum {0.2, 0.5, 0.8}, m-rate fixed at 0.005), cutoff decreases for increasing strength (lower rank minimum); c) & d) convergence plots for conditions a) and b) respectively.

m-rate	Convergence	Final *cutoff*
0.001	0.985	3.0
0.005	0.893	7.7
0.01	0.801	13.5

Table 5-3a: Bitwise convergence and mean cutoff at equilibrium against mutation rate under weak selection (rank min. 0.8).

sel. rank min.	Convergence	Final *cutoff*
0.2	0.988	1.7
0.5	0.973	3.0
0.8	0.893	7.7

Table 5-3b: Bitwise convergence and mean cutoff at equilibrium against selection rank minimum, m-rate 0.005/bit.

Turning now to the degree with which the population in each condition converges, table 5-3, the amount of variation in the equilibrium population increases for increasing m-rate and reducing selection strength. For intermediate and strong selection, the population is highly converged, *cutoff* is low and there is little canalisation. For weak selection or high mutation, two conditions that favour a high degree of genetic variation, *cutoff*, which is proportional to the degree of genetic canalisation, remains at quite a high level - the preferred evaluation surface is smooth.

5.4.2 Canalisation and Local Search.

By introducing a within generation, adaptive local search operator (i.e. learning), I have argued that the valuation landscape is smoothed. One would thus expect that a population whose members are capable of learning would evolve a different genotype-phenotype map compared to the case of a population of directly valuated individuals. For the fault induction model (i.e. within generation, *maladaptive* local search) I expect the population to evolve towards smoother regions of the direct valuation landscape than might otherwise be expected. That is, fault induction makes the landscape appear to be more rugged than it actually is.

Genetic canalisation relates to the ability of an individual to withstand mutational change. Where adaptive plasticity is available (i.e. where an adaptive local search operator is applied), then plasticity may buffer an adapted *phenotype* from mutational change, and canalisation (as evidenced by a high value of *cutoff*) is likely to be low. Additionally, a locus that may be hidden from selection by plasticity is not likely to fixate on any particular allele (except under drift) and genetic variation at that locus will be high. For the case of maladaptive local search, a high degree of genetic canalisation, as measured by a high equilibrium *cutoff*, is likely to be required.

In the following experiment, I will show how within generation local search can influence the evolution of an appropriate evaluation function. Two modes of local search are demonstrated - single bit steepest ascent (learning), GA-SA1, and single bit steepest descent (fault induction), GA-SD1. GA parameters not given in the caption are set as in section 5.4.1.

Considering figure 5-8, for the adaptive regime, the *cutoff* rate evolves to a lower value than in the simple case: the population prefers a rather more rugged evaluation landscape than one

would otherwise expect (i.e. it is more rugged than the one SGA goes for). As in the simpler regime, actual modifier values at equilibrium in any given run were within +/- 1 of the integer value of the reported mean. In other words, the learning operator *smoothes* the landscape, as expected. Under the single bit fault induction model, and weak selection, the *cutoff* rate has a higher equilibrium value than that of the SGA. That is, the population prefers a smoother evaluation surface because the fault induction operator effectively 'sharpens' the landscape.

Comparing bitwise convergence of the population at equilibrium for the three regimes, table 5-4, all three regimes appear to converge to a similar degree. This suggests that all three regimes consider themselves to be located around the optimum on selective surfaces of similar ruggedness. By assuming that the value of *cutoff* has converged to a steady value, then this indeed appears to be the case, since under rank selection, selective values across a population are fixed - i.e. the ruggedness of the selective landscape is fixed. During the early generations, the value of *cutoff* is allowed to evolve due to the different marginal evaluations of individuals carrying *different* values of *cutoff*. Using regime SGA on various *fixed* landscapes, the equilibrium convergence value is the same whatever the fixed *cutoff* value (not shown).

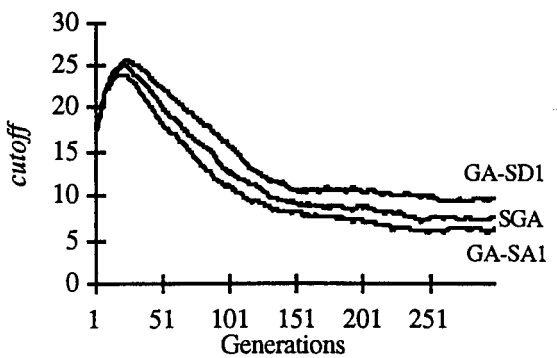

Figure 5-8: selection rank minimum 0.8, m-rates.0.005/bit over targeting loci, 0.01/bit over cutoff loci;

Regime	Convergence	Final *cutoff*
GA-SA1	0.891	6.4
SGA	0.893	7.7
GA-SD1	0.886	9.8

Table 5-4: Bitwise convergence and mean cutoff at equilibrium, time averaged over 10 generations, for SGA, GA-SA1 and GA-SD1. Rank minimum 0.8, m-rates 0.005/bit in the targeting sequence, 0.01/bit for cutoff *loci.*

5.4.3 The Dynamic Stability of Evolving Populations: A Case Study

In the previous sections, I have reported on the *mean* value of the *cutoff* gene across the population and also a simple measure of convergence. Such simple reports are typical of those frequently offered by the GA community, and do reveal some interesting - if inexplicable - properties. However, the richness of the underlying population dynamics is denied by such an approach (for example by taking mean measures over several runs, which is important for describing robust optimisation procedures but fails to characterise single run dynamics). In this section, I describe those dynamics for a single, asexual run with m-rates both set at 0.005/bit, crossover rate 0.0 per individual and selection rank minimum 0.5.

Typical measures - bitwise convergence, mean *cutoff* and mean evaluation are given in figure 5-9. The long term behaviour of figure 5-9a is stable, fluctuating apparently at random between 600-650 in the case of figure 5-9b. The long term equilibrium behaviour may be used by offering a time averaged mean evaluation taken over generations 100 to 150, for example.

Of potentially more interest are the number of individuals carrying each particular allele of the *cutoff* gene, and the distribution of populations over mutant classes (i.e. Hamming bands) (figures 5-9c to 5-13). For figures 5-10 to 5-14, 5-16 and 5-17, 'time-lapsed' views of the motion of population distributions are given, the initial and final population distributions for the given sequence being depicted separately to assist in the interpretation of these views. Early and late generation plots of the population and modifier distributions are given in figures 5-9c and d respectively. The population initially assumes a distribution centred on individuals 11 Hamming units away from the optimum, and skewed towards it. At equilibrium, the whole population us gathered around the optimum. The equilibrium distribution of the *cutoff* modifier maximises the payoff for individuals close to the optimum. However, the early distribution shows that several different values of the modifier are maintained within this searching population.

Figures 5-10 to 5-12 depict the same information (the evolution of the population distribution over mutant class, and the number of individuals bearing each cutoff allele) in three different ways. Figures 5-9c and d show the respective first and last generation population distributions for a single run. Figure 5-10 adds all points in between, showing the temporal evolution of the population over the course of the whole run. This demonstrates

how: a) the population is initially normally distributed over mutant classes 10 to 24, and ultimately focused on mutant class 0; b) how the cutoff allele takes on certain values in preference to others (e.g. 15, 5, 2 and 1 all dominate the population at certain points during the run, whereas allele 4, 8, 10 and 12 rarely ever appear. This is illustrated more clearly in figure 5-11, where for each mutant class/cutoff value, a generational count of the number of individuals in the population within the respective class is given; and in figure 5-12, which plots the same information over a single time axis. Considering the dynamics over the whole run, the first thing to note is that the dynamics of the population distributed over mutant classes appears to behave sensibly - the initially random population has a spread distribution centred around mutant class *L*/2 (figure 5-10a). Over time, the population becomes more converged and moves smoothly towards the optimum (figure 5-11a). In the long term, the population settles as a quasi-species mainly around mutant classes 0, 1 and 2 (figures 5-12a and 5-13a). The spread of the population at equilibrium may tuned by altering the selection rank minimum or mutation rate parameters.

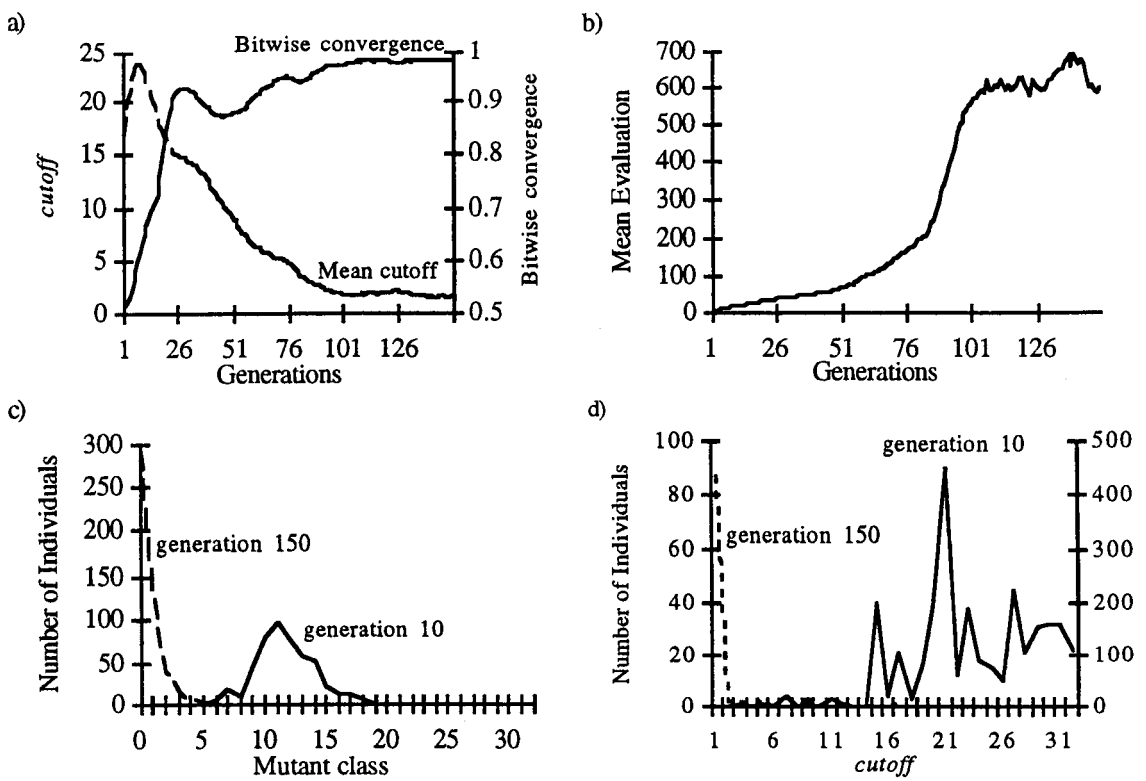

Figure 5-9: Single run population measures - a) bitwise convergence and mean cutoff *over generations; b) mean evaluation over generations; c and d) generation 10 and 150 distribution of the population over mutant classes, and values of the cutoff modifier respectively.*

On the other hand, the dynamics of the modifier gene do not appear to be so well behaved. Early on, the *cutoff* allele 15 dominates as the population exhibits phase 1 search, peaking around generation 25 (figure 5-12b). Other values then come to dominate, before being replaced by values offering higher peaks. However, not all values of *cutoff* come to dominate the population (figures 5-10b and 5-11b). Rather, it appears that certain values are ignored (for example, *cutoff* value 4 barely gets a look in at all during the run). Whether this jumpy dynamic is an artefact of the binary coding scheme used for the modifier gene, or whether it represents something far less trivial is left for further work. The value of the modifier that dominates long term behaviour for the particular parameters used here is the optimal value of 1, although there is a trace representation of other values (figures 5-12b and 5-13b) which represent the Hamming 1 neighbours of the dominant allele.

Figure 5-10: (overleaf) Time lapsed generational population distributions. The figures depict the population distribution a) over mutant classes, b) over cutoff alleles, for each and every generation in a single run. Early generation distributions are on the right hand side of each figure, final generations on the left.

Figure 5-11 (p146): The temporal evolution of the number of individuals a) in each mutant class, b) carrying each cutoff allele. Within each class/allele value, there is a generational time line running left-right.

Figure 5-12 (p 147): The temporal evolution of the number of individuals a) in each mutant class, b) carrying each cutoff allele. For each class/allele value, the number of individuals in the population within the respective class/bearing the given value, is plotted over time.

Figure 5-10a)

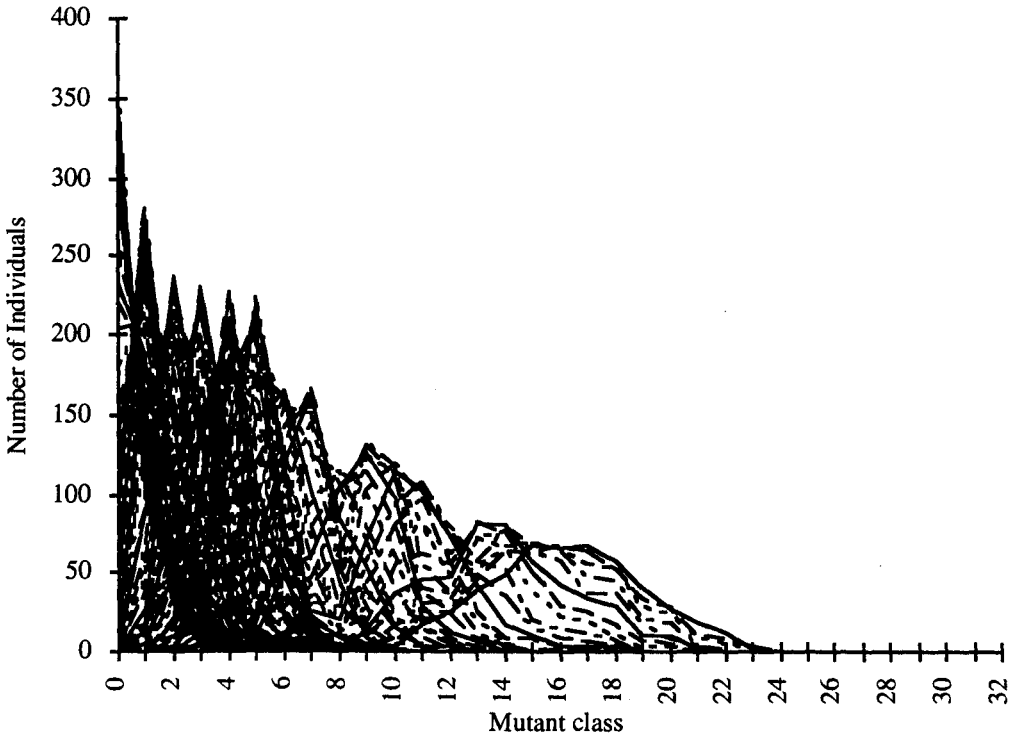

Figure 5-10b)

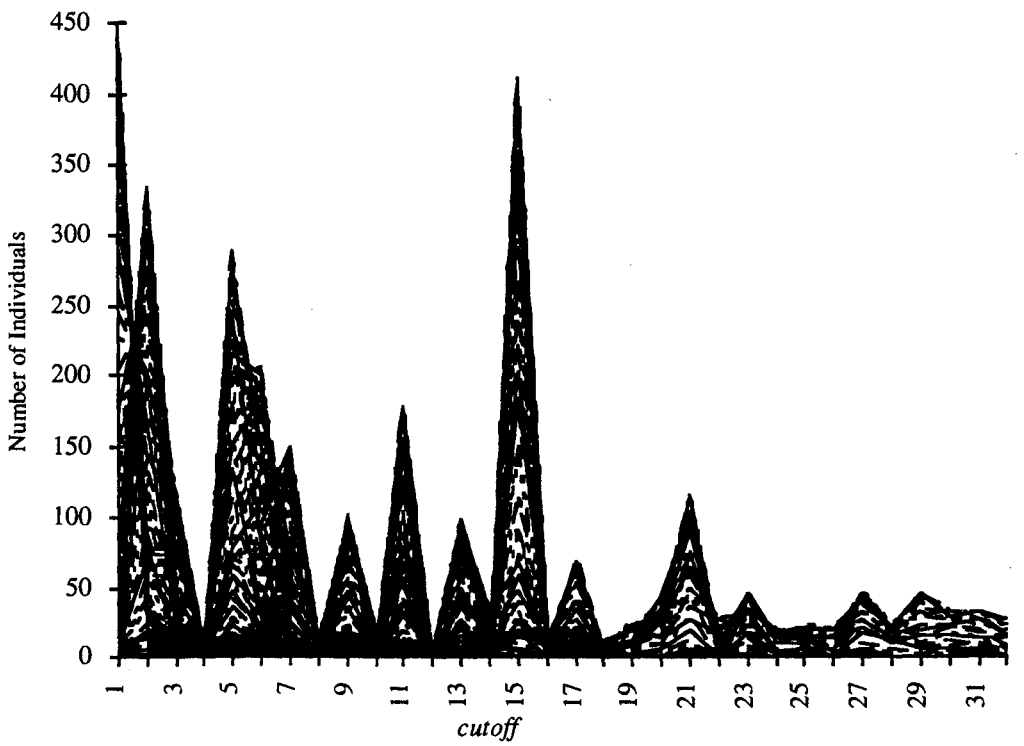

Figure 5-11:

a)

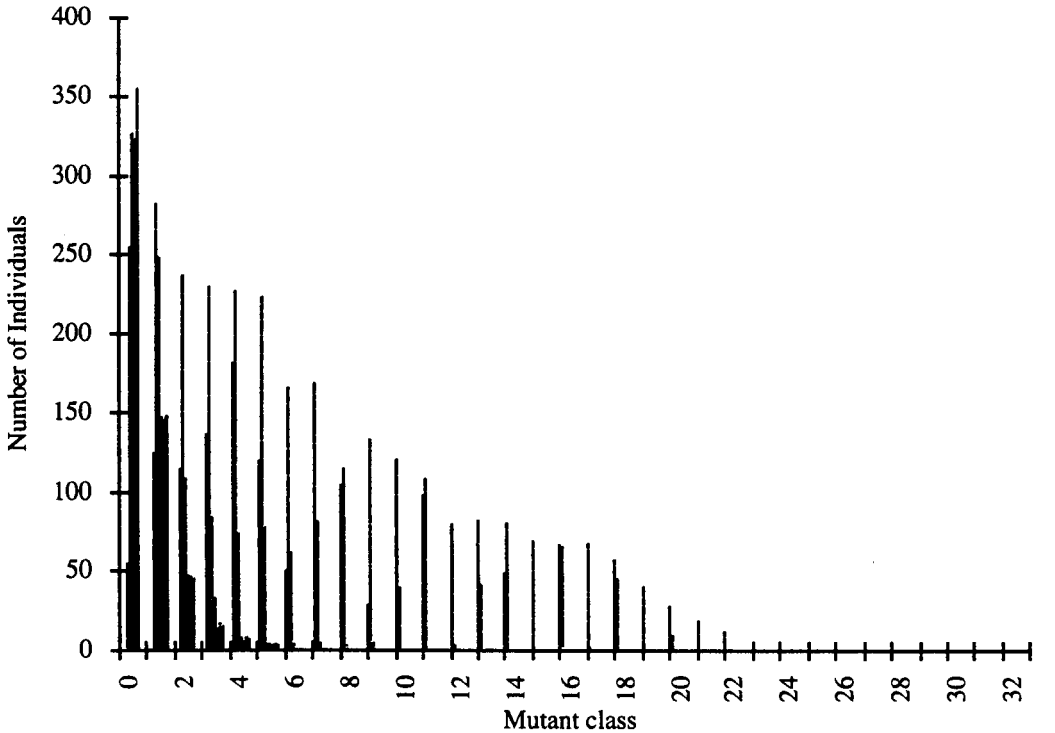

b)

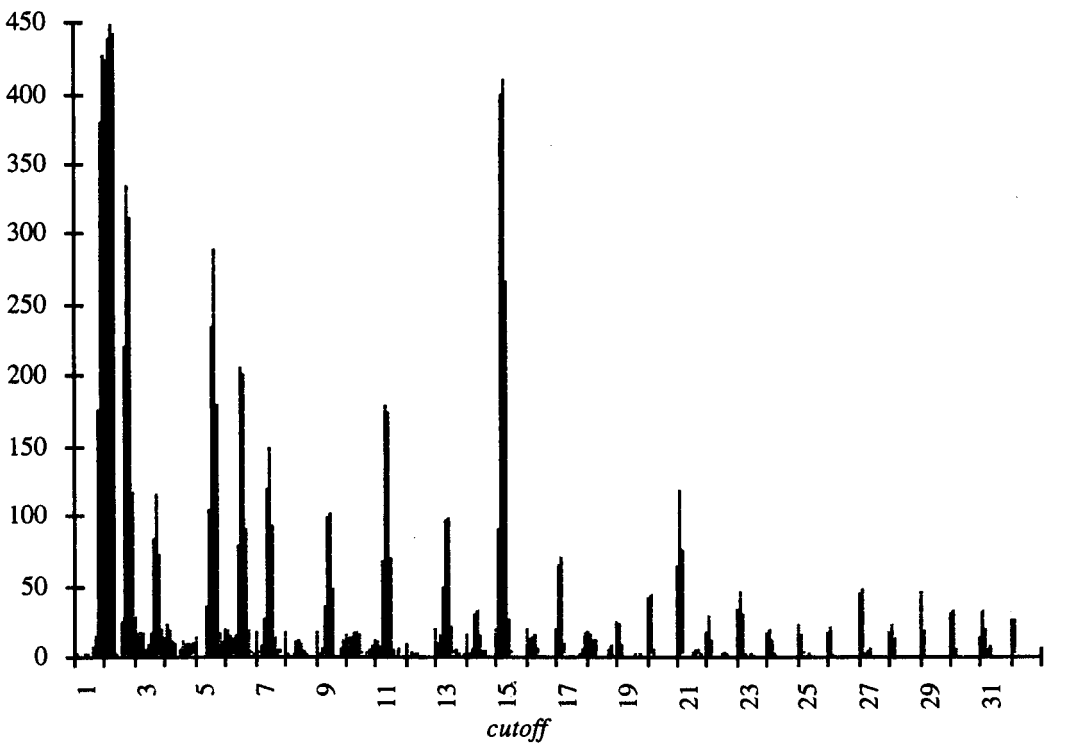

Figure 5-12:

a)

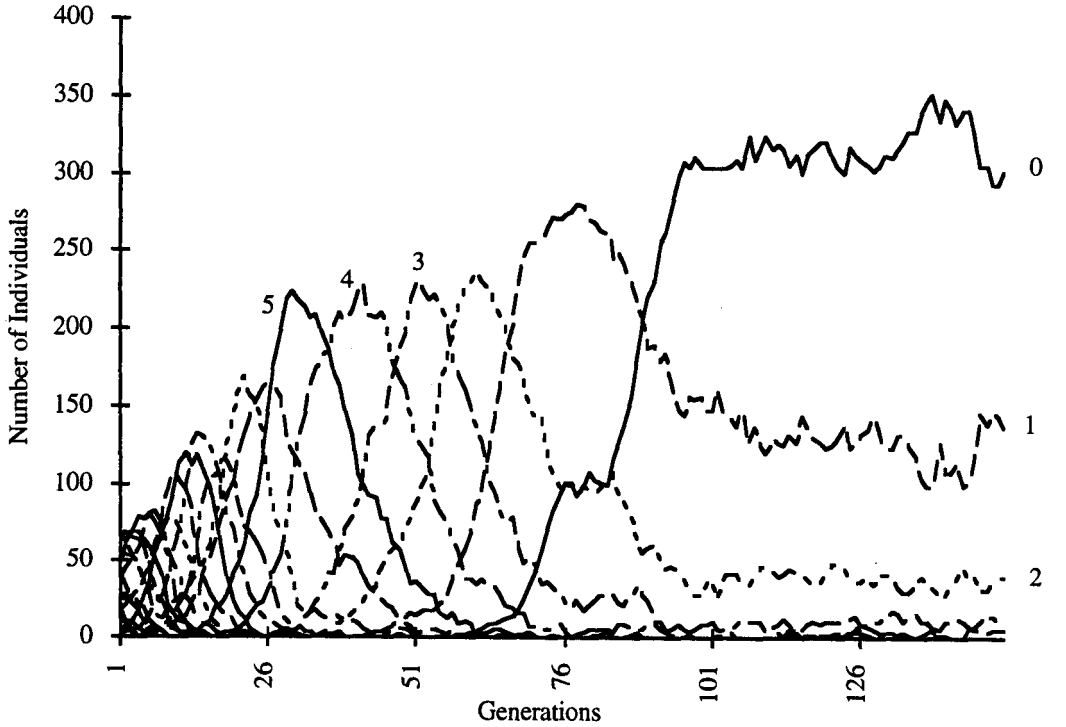

b)

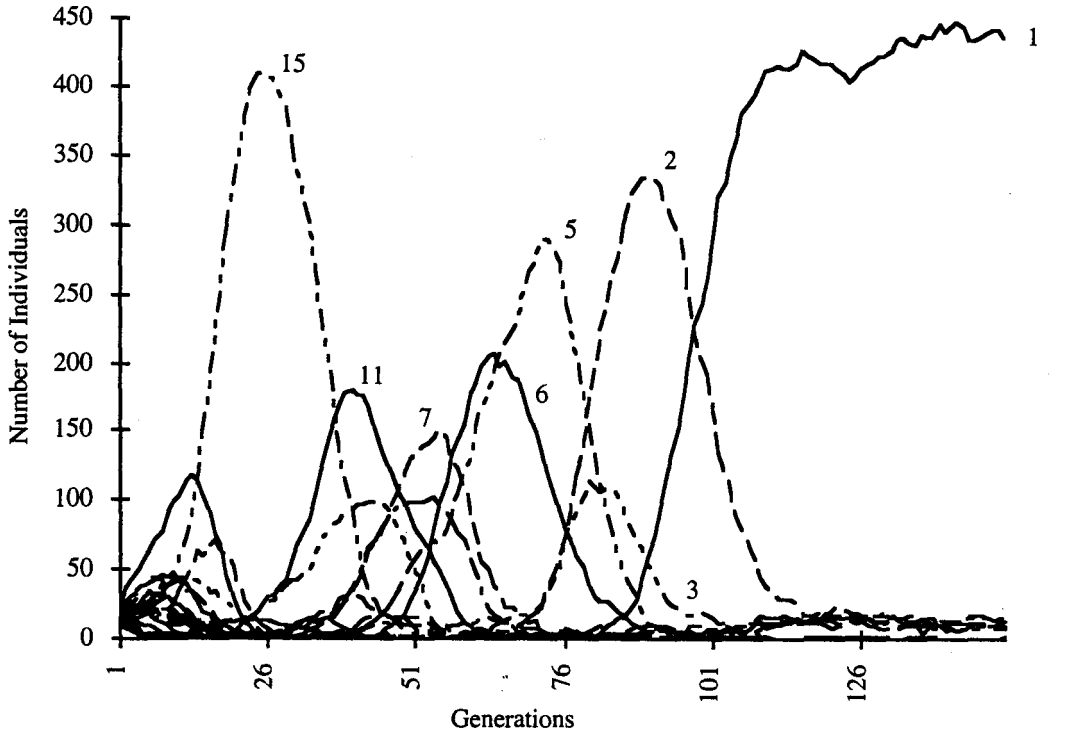

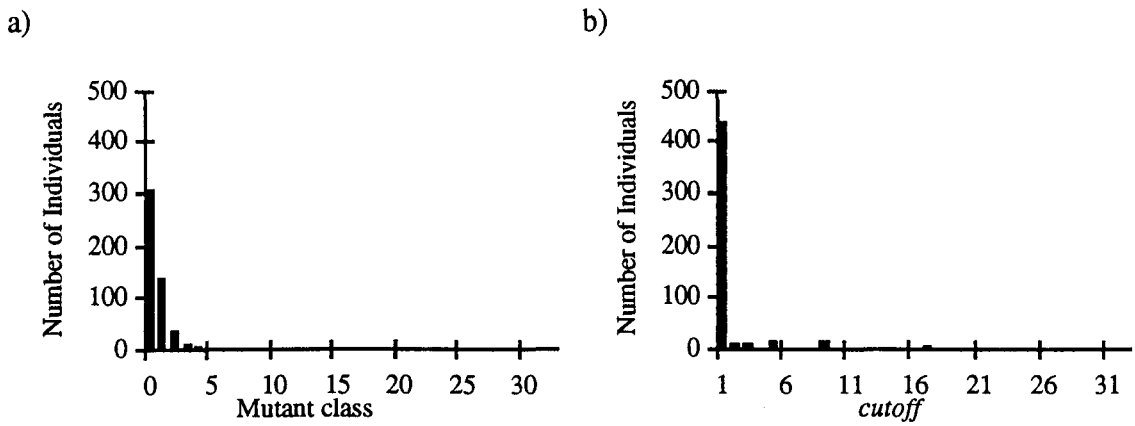

Figure 5-13: a) Generation 150 population distribution over mutant classes; b) generation 150 distribution of cutoff *alleles.*

Now, let us try and describe some of the features of the mean results by reference to actual population dynamics. Firstly, the dip in the bitwise convergence curve (figure 5-9) between generations 30 and 45. Referring to figure 5-14a, the population is moving from mutant class 6 to mutant class 5 during this period. In addition, the proportions of certain values of the modifier within the population appear to evolve in phase with each other (figures 5-15). So for example, the initially dominant modifier allele 15 flows into alleles 11, 13, and 14 at around generation 35 (figure 5-15c). These alleles in turn are mutated into alleles 9 and 7 about 15 generations later (figure 5-15d). Although these results are demonstrated for a particular run, they raise the question as to whether evolutionarily stable distributions of the modifier exist for given population distributions (and vice versa). In the search for general methods, it is suggested that analysis of the Tower of Babel evaluation function (which shall not be attempted here) may proceed not only along a quantitative genetics axis, but also in the game theoretic approach of evolutionarily stable strategies.

To finish this reconnaissance of the dynamics of the Tower of Babel, I shall consider the final stages of the search to equilibrium, generations 74-110. At the start of this period, the population appears to enter a metastable regime, the plateaux in the mutant class traces of figure 5-12a. During the same period, however, there is still a significant amount of activity in the evolution of the modifier gene, figure 5-12b. Taking a closer look, figure 5-16, note how the population is stable within mutant class 1, whilst the dominant allele of the modifier gene changes to give a steeper evaluation function, and as a consequence the population mean evaluation climbs rapidly (figure 5-9b). The drop in bitwise convergence during this period (figure 5-9a) is presumably as a result of the motion of the modifier.

a) b)

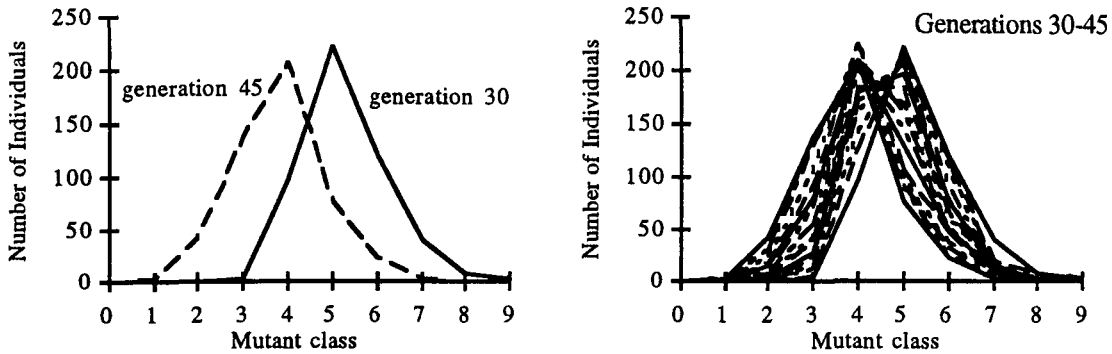

Figure 5-14: The population distributions over mutant classes a) at generations 30 and 45; b) over generations 30-45.

Figure 5-15 (below and overleaf): a and b) cutoff *distributions over generations 30-37 and 38-45 respectively. Different values of cutoff coming to the fore in phase with each other: c) values 11, 13 and 14; d) values 7 and 9.*

Figure 5-15a) Figure 5-15b)

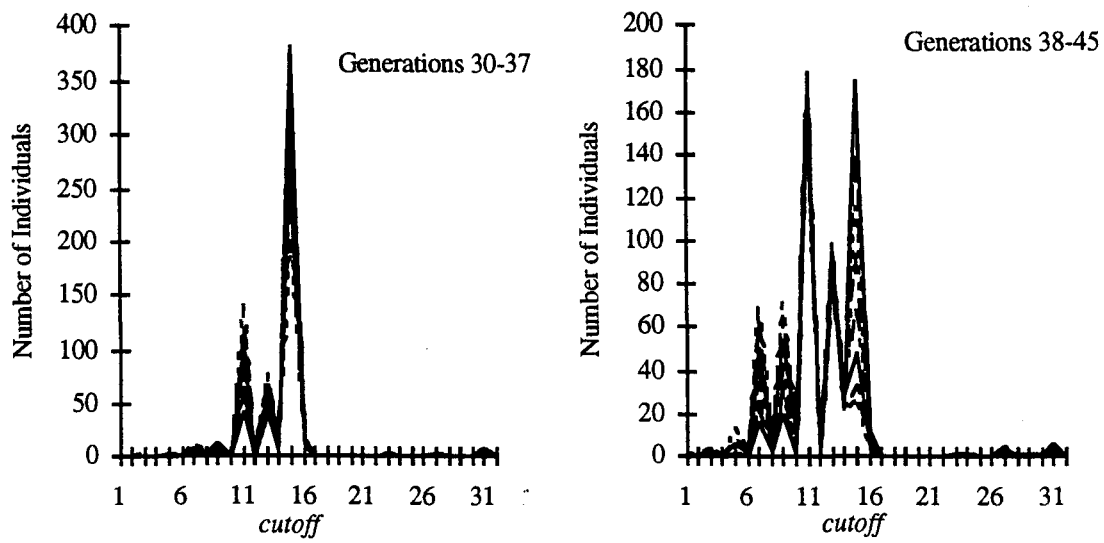

Figure 5-15c)

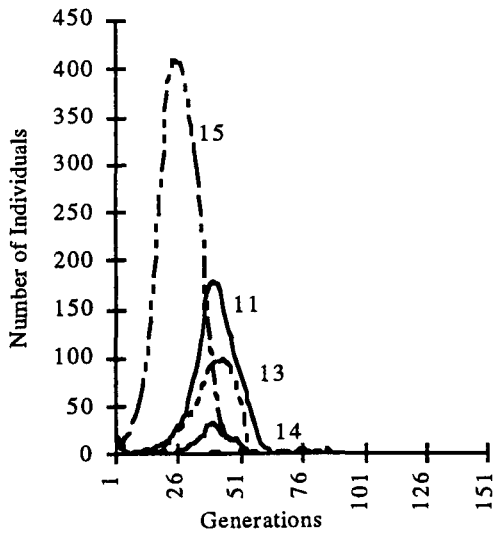

Figure 5-15d)

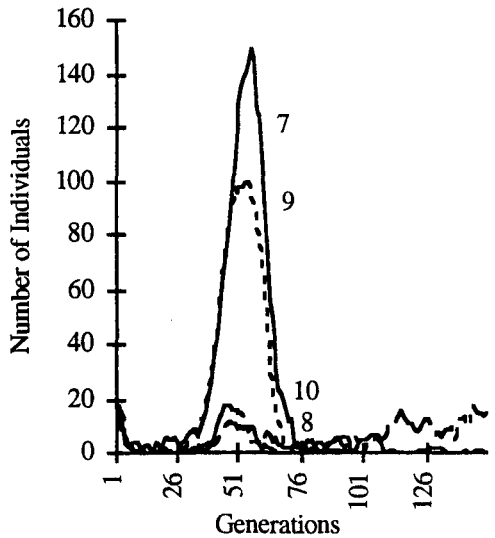

Figure 5-16a)

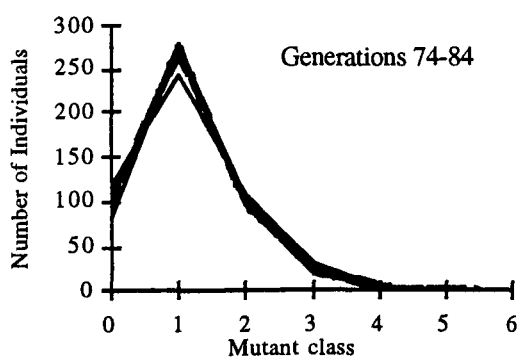

Figure 5-16b)

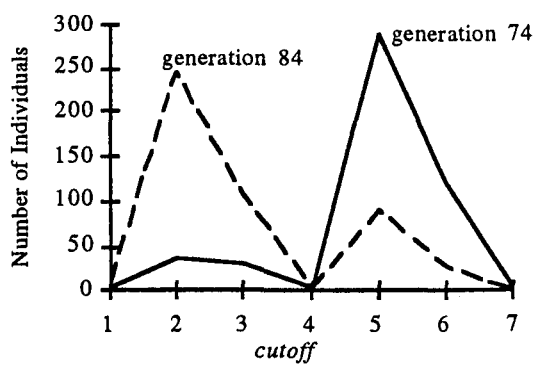

Figure 5-16c)

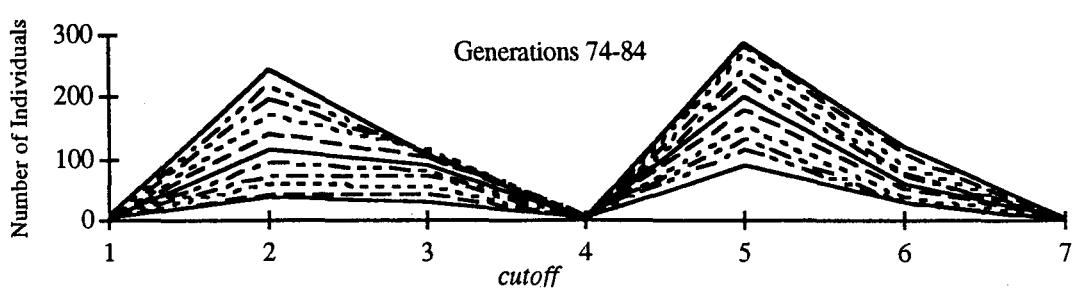

Figure 5-16: a) Population distribution over the lower mutant classes is steady for generations 74-84; b and c) the dominant modifier changing from allele 5 to allele 2 over the same period.

The final stage of the search, generations 85 until equilibrium is initially achieved around generation 110, is portrayed in figure 5-17. In this case, however, there is a shift in both the population distribution over mutant classes and the identity of the dominant modifier allele. Clearly, there is much still to be done in understanding the very rich dynamics exhibited by this relatively simple system.

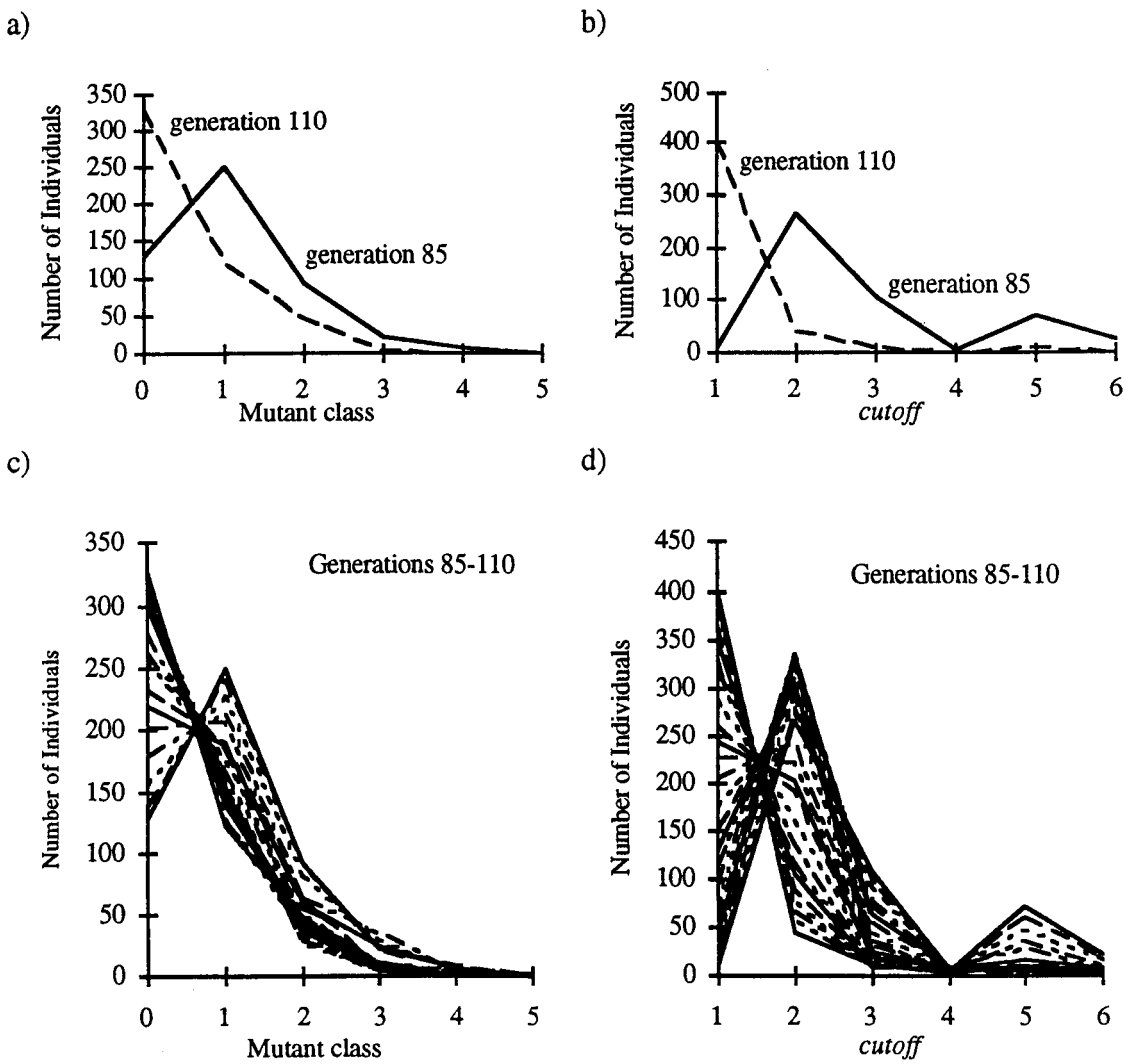

Figure 5-17: a and c) Population distribution over the lower mutant classes moves from mutant class 2 to class 1 over generations 85-110; b and d) the dominant modifier changing from allele 2 to allele 1 over the same period.

Finally, I shall consider the case of weaker selection (in particular, for rank minimum 0.8), where canalisation is likely to occur. Initial phase dynamics are similar to those reported above. Of rather more interest are the population distribution and distribution of modifier alleles at later stages of the run. Dynamic equilibrium of the population distribution appears to occur by around generation 250, and is of the form depicted in figure 5-18a. Note how the population is now centred around mutant classes 2 and 3 as the population is more free to drift away from the optimum. This effect is in accord with earlier work by (Woodcock & Higgs, 1996).

The modifier distribution does not settle so tidily, as shown in figures 5-18b-d. Consequently, the mean evaluation is not steady for later generations. What can be said, however, is that modifier allele 2 seems to dominate in the long term, with its single mutant neighbours ebbing and flowing as a result of the very weak selection on the modifier itself.

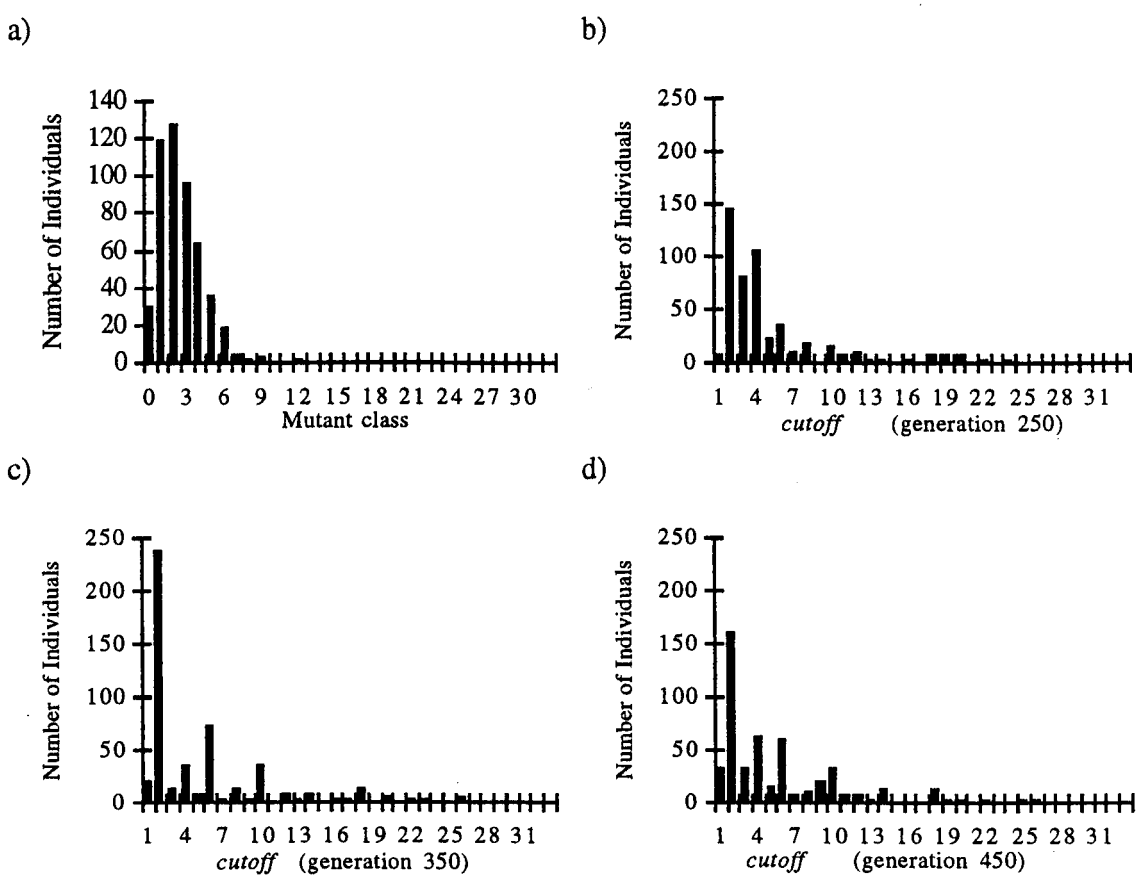

Figure 5-18: Selection rank minimum 0.8, m-rates both set at 0.005/bit. a) Generation 250 population distribution over mutant classes; b-d) generation 250, 350, 450 distribution of cutoff *alleles.*

5.5 Summary

In this chapter, I have demonstrated how genetic algorithms are ideally suited to identifying robust solutions to optimisation problems, although under certain conditions they may act as global optimisers. This stands in contrast to the public face of GAs, where they are often perceived solely as global optimisers. The view stated in chapter 3 of evolutionary search from an initially random population progressing through 3 distinct phases has been shown to provide a useful way of understanding the population dynamics of an evolving system.

Interesting evaluation functions are likely to present (e)valuation surfaces of varying degrees of local ruggedness over the search space. For a given operator set, I conjectured that the operators are adapted to regions of a landscape with a certain characteristic degree of ruggedness, as given by the runtime operator correlation coefficient (3.4). In this chapter, I showed how by using the variance of values in the vicinity of a population as an estimate of local ruggedness, different search regimes identified different regions of the landscape as characterised by the simpler, instantaneous measure. By introducing local search, valuation surfaces may be transformed and the region of the search space to which a population is attracted is changed. In particular, the use of fault induction appears to guide an evolving population towards smooth regions of the search space, whereas learning tends to lead the population towards rugged areas where higher optima are to be found.

By independently evolving the evaluation function on which individuals were evaluated, it was possible to show how under weak selection individuals evolved that were genetically canalised against the effects of mutation. The introduction of adaptive local search reduced the degree to which canalisation evolved (discounting the direct canalising effect of adaptive search), whereas fault induction increased the level of genetic canalisation supported. The evolution of genetic canalisation may be used to evolve implicit fault tolerance to faults which may be described in terms of mutation. Rather than concentrating on operators and parameter sets that will maximise the chances of locating a global optimum, if operators are defined such that the phenotypic effects observable to selection of mutation are similar to the effects on phenotypes of a (class of) faults, implicit fault tolerance will arise. Where problem solutions describe component values that are subject to tolerance limits, fault induction that evaluates a component at the worst case limit of its tolerance is likely to improve the stability of discovered designs.

Finally, I must answer the question "why are evolvable functions likely to be of ongoing interest?" One important reason is that evolvable evaluation functions may provide an implicit way of optimising complex problems. If a family of curves can be derived such that a) evaluation function parameters can be coded for as modifier genes; and b) there is some way of providing a selection pressure towards the parameter set of the evaluation function one wishes to optimise, then it may be possible to increase problem difficulty from a simple initial problem self-adaptively. A second motivation is that such evaluation schemes may offer a test bed for simulating the coevolutionary dynamics of evolutionarily stable strategies and evolutionarily stable populations.

In the next two chapters, I shall discuss the dynamics of evolution in a sinusoidally fluctuating selective environment. By interpreting quantitative genetics models of such a system in terms of signal processing, it becomes possible to probe the analyses in a novel way. The use of a GA simulation applying similar selective conditions adds a further dimension to this approach.

Chapter 6 - Evolutionary Signal Processing

"Our Second Experiment...is the production of [a] seldom-seen-but-greatly-to-be-admired phenomenon..."

L Carroll, from "The Professor's Lecture" in Sylvie and Bruno Concluded.

6.1 Introduction

In this chapter, I shall argue that the techniques of signal processing may be applied to the analysis and interpretation of quantitative genetic models of evolution that offer a similar mathematical form to systems identified in signal processing. By demonstrating that certain mathematical models from biology are identical to signal processing models, it becomes possible to use the concepts of signal processing to understand the consequences of particular model types (i.e. how they behave, what they predict) without having to recourse to large amounts of additional analysis or mathematical modelling.

Signals are quantities that convey information. *Signal processing* relates to a set of techniques that allow the manipulation of continuous or sampled signals. *Signal processors* are frequently encountered in the guise of *filters* which are broadly defined as systems whose output signal differs from the input signal in a well defined way. A more useful definition is of a *frequency selective system that attenuates certain frequency components whilst passing others unchanged.* This applies to both temporal and spatial (typically image processing) domains.

In accord with the view of EAs as function optimisers, there has been considerable interest in using EAs for both the optimisation of the parameters of canonical filters, and also in the discovery of the architectures of the filters themselves e.g. (Arslan & Horrocks, 1995). GAs have also been used for the online adaptation of a channel equaliser implemented in reconfigurable logic (Murakawa, Yoshizawa and Higuchi, 1997). By way of contrast, I will argue that it is possible to treat *evolution itself* as a filter of temporal signals.

Traditional quantitative genetics discusses evolving systems in terms of optimal phenotypes, population mean phenotypes and measures of selection strength and genetic variance. By

treating the optimal phenotype at any given time as the *input signal* to the evolutionary system, and the population mean phenotype as the system's *output signal*, it is possible to characterise the *transfer function* of the evolutionary system[57]. In 6.2 I shall draw on such transfer functions derived in quantitative genetics analyses (Lande, 1996; Charlesworth, 1993) and liken them to those of signal processing filters as derived for use by engineers (Baher, 1990). In 6.3, the filter characteristic is demonstrated using a GA experiment in which the population is set the task of tracking a sinusoidally fluctuating optimum; the low pass filtering property is clearly demonstrated in 6.4 by applying an input signal containing a significant amount of high frequency noise. In 6.5, I employ an evaluation function which describes a sinusoid whose frequency is itself dependent on time, and demonstrate the tracking properties of the *mutation rate* under such conditions. Finally, in 6.6, I speculate on how the ESP approach may be developed in the digital domain by describing several possible implementations of an evolutionary digital filter.

6.2 Filtering Temporal Signals

In this section, I shall describe the evolutionary filtering of temporally varying signals. If there is a way of putting the evolutionary transfer function into the form of a traditional filter transfer function, relating genetic variance and selection strength parameters to the design parameters of a particular filter, it will be possible to use the knowledge of filter design to tune the evolutionary parameters so as to obtain the required filter characteristics. Since the 'output' of the evolutionary filter is given by the mean of the population, this represents population level filtering.

In this initial study, I shall consider two quantitative genetics analyses of evolution in sinusoidally varying environments. It lies to further work to generalise the theoretical approach through the Fourier analysis or wavelet analysis of rather more complex (i.e. *general*) optimal signals. In 6.2.1, I shall provide the background to filter theory and

---

[57] (Kargupta & Goldberg, 1994) interpret the notion of a signal in a GA context rather differently, specifically as the selective value of an individual: "The confidence allocated to a hypothesis reflects the deterministic efforts toward a certain direction and is therefore called a *signal*." This stands in contrast to the definition of a signal as an optimal phenotypic target of selection, or the *expression* of an individual acted on by selection, which is used herein.

quantitative genetics analysis necessary for appreciating ESP in the continuous and discrete time domains. I shall then describe the ESP approach in the continuous domain with reference to the 'steady state' quantitative genetics model of (Lande & Shannon, 1996), and analogue filter theory, 6.2.2; and the discrete time domain by a comparison of the generational model of (Charlesworth, 1993) and digital filter theory, 6.2.3. To maintain consistency in notation, some simple substitution of symbols in some referenced equations has been followed. A direct comparison of the two biological models is presented in 6.2.4. Finally, in 6.2.5, I consider how an evolving population may itself be interpreted as a filter of adapted individuals.

6.2.1 Essential Theoretical Background

The theoretical basis of Evolutionary Signal Processing is provided by interpreting quantitative genetics analyses of evolution under temporally fluctuating selective conditions in terms of formally similar signal processing analyses. In this section, I shall provide a rapid introduction to both approaches. More details can be found in any domain specific introductory text (for example, (Baher, 1990; Stark, 1979), for treatments of signal processing; (Roughgarden, 1979; Bulmer, 1985) for an overview and a detailed consideration of quantitative genetics respectively).

6.2.1.1 Elements of Signal Processing.

As I mentioned in the introduction to this chapter, filters are systems that are capable of accepting (passing) or rejecting (stopping) certain frequency components of an input signal. As their names suggest, *low pass* filters pass low frequency input signal components, rejecting (i.e. significantly attenuating) higher ones; *high pass* filters pass high frequency components of the input signal and reject lower frequencies; and *band pass* filters reject input signal components of both high and low frequency frequencies. The *pass band* of the filter refers to the range of frequencies that the filter will accept and transmit without severe degradation in amplitude (i.e. without significantly *attenuating* the signal). The *transition band* represents the frequencies over which the filter response changes from a pass response to a stop response. The *gain* of the filter describes the ratio of the output signal amplitude to the input signal amplitude.

The behaviour of any given filter is typically described in terms of its *amplitude* and *phase* response. For an input sinusoid to the filter, the amplitude response describes the extent to which individual frequency components are attenuated by the filter. The amplitude response is traditionally presented in terms of the ratio of the power in the output signal to the power in the input signal, plotted over all frequencies, although a plot of the gain against frequency may also be used. Signal power is itself described in terms of the *gain squared.* Figure 6-1, for example, shows the amplitude response of a band pass filter. The phase response gives the phase delay between input and output signals, again for all frequencies.

The *cut-off* frequency, or *break point*, of an ideal filter refers to the frequency defining the edge of the pass band, although for real (non-ideal) filters it gives the transition point between stop and pass bands. This is taken to be the frequency at which the power of the output signal is *half* that of the power of the most efficiently passed signal. Power ratios are typically described in terms of *decibels*, defined as: $10\log_{10}(P_1/P_2)$. The cut off frequency is thus frequently referred to as the 3dB point, since $10\log_{10}(P_{out}/P_{in}) = 10\log_{10} 0.5 \approx -3\text{dB}$.

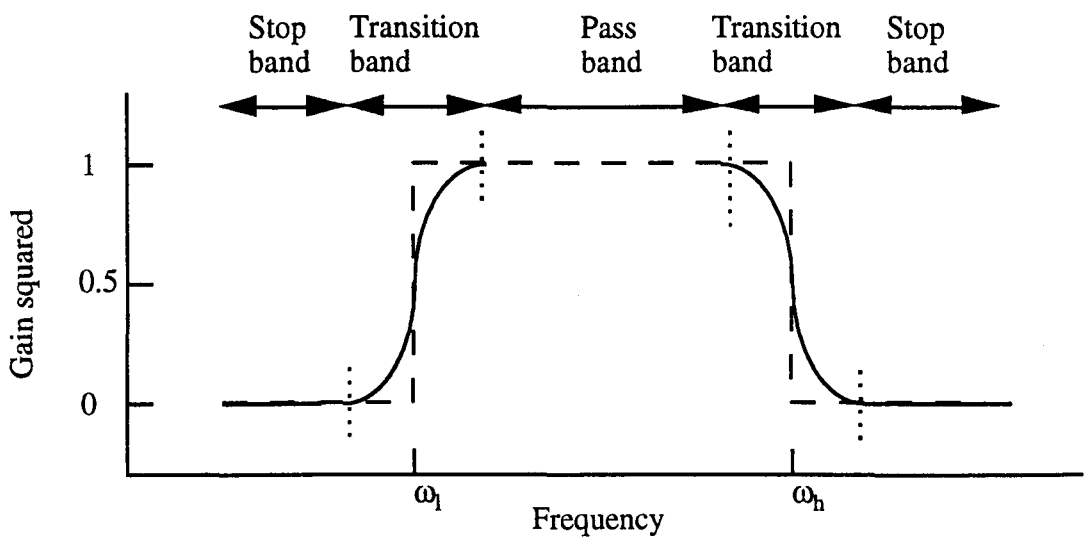

Figure 6-1: Elements of signal processing - the amplitude response of an ideal band pass filter (dashed lines) and a real band pass filter (solid lines).

6.2.1.2 Elements of Quantitative Genetics Analyses of Evolution Towards Nonstationary Optima.

Quantitative genetics methods are frequently used to analyse the evolution of a population's mean phenotypic expression in terms of the population distribution and the strength of

selection towards some optimum. The following derivation is taken largely from a standard reference to the area, (Bulmer, 1985).

The population distribution is given by:

$$f_p(y,t) = \exp\left(-\frac{(y-m(t))^2}{2V_g}\right) \tag{6.1a}$$

where y is an individual phenotype, $m(t)$ is the *population mean phenotype* at time t and V_g is the *additive genetic variance*[58] which is sometimes taken to be fixed over the course of the run[59] (Lande & Shannon, 1996).

Selective values, $w(y)$, arising from *optimising selection* (which seeks to minimise the distance between the population mean and some optimal phenotypic state) are given by:

$$w(y,t) = \exp\left(-\frac{(y-\theta(t))^2}{2V_s}\right) \tag{6.1b}$$

where θ is the *optimum* at time t, and V_s is the *variance of selection*. The *strength of selection* is inversely proportional to the variance of selection.

The population distribution *after* selection is proportional to the product of the population distribution and the selection function:

$$f_p(y,t+1) = \frac{w(y,t)f(y,t)}{\int w(y,t)f(y,t)dy} \tag{6.1c}$$

---

[58] The *additive genetic variance* is defined by (Roughgarden, 1979), p153, as "the variance among the phenotypes based on predicting the phenotypes from the additive effects [(on the expressed phenotypes) of individual alleles]." A quantity related to environmental effects is also frequently added to the additive genetic variance to give an overall population variance (although this environmental component may be included within the *selection variance*).

[59] Applying such a constrained analysis to the behaviour of a GA suggests that the model will only be applicable, if at all, when a steady state population behaviour has been achieved - i.e. when the population is in phase 3, dynamic equilibrium.

that is:

$$f_p(y,t+1) \propto \exp\left(-\left[y-\frac{V_g\theta(t)+V_s m(t)}{V_s+V_g}\right]^2 \Big/ \frac{2V_sV_g}{V_s+V_g}\right) \tag{6.1d}$$

The change in population mean phenotype that results is then calculated as:

$$m(t+1) = m(t)+\frac{V_g}{V_g+V_s}.(\theta(t)-m(t)) \tag{6.1e}$$

Following (Charlesworth, 1993), by letting:

$$k = V_g/(V_g + V_s) \tag{6.1f}$$

the change in population mean following selection, equation (6.1e), reduces to:

$$m(t+1) = (1-k)m(t)+k\theta(t) \tag{6.1g}$$

For overlapping generations, a corresponding expression governing the rate of evolution of the population mean for weak selection is given by (Lande & Shannon, 1996) as:

$$\frac{dm}{dt} = -\frac{V_g}{V_s}[m(t)-\theta(t)] \tag{6.1h}$$

This follows on naturally from (6.1g) since by following Charlesworth again and letting:

$$V = V_g/V_s \tag{6.1i}$$

where

$$V_s >> V_g \tag{6.1j}$$

we get for weak selection a large selective variance, hence small V and thus:

$$k \approx V \tag{6.1k}$$

from which a direct analogy between the discrete and continuous cases (of (6.1g) and (6.1h) respectively) may be drawn.

Where the genetic variance is allowed to change (e.g. the genetic operators fail to reintroduce variation at the rate at which selection removes it (cf. GA phase 2)), selection reduces population variance according to:

$$V_g(t+1) = \frac{V_s V_g(t)}{V_s + V_g(t)} \tag{6.1l}$$

However, (Anderson, 1995) offers an expression which introduces mutation (described by a mutation rate with a Gaussian distribution centred around a mean of 0 and with variance V_m). Convolution of the mutation distribution and the population distribution following selection gives a non-zero equilibrium population variance held steady through the mutation-selection balance:

$$V_{g,eqm} = \frac{V_m + \sqrt{V_m^2 + 4V_m V_s}}{2} \tag{6.1m}$$

Note that the population variance increases for increasing V_m (i.e. mutation rate) or decreasing selection strength (increasing V_s).

One final observation of this style of analyses comes from (Anderson, 1995), specifically that the introduction of *learning* may be modeled by increasing the variance of selection. This is covered in more detail in section 7.2.

6.2.2 Evolutionary Signal Processing I: A Continuous Model Realises an Analogue, Butterworth Filter

Filters that process a continuous input signal are known as *analogue filters*. In this section, I shall demonstrate how a quantitative genetics analysis of the behaviour of a continuously evolving population exhibits a similar mathematical form to a well known low pass analogue filter.

6.2.2.1 A Brief Introduction to Analogue Filter Theory.

(The following is based on (Baher, 1990), but any introductory textbook to signal processing is likely to follow a similar derivation).

In electronic filter design, filter characteristics are typically given by the transfer function, *h(t)*, which describes the relationship between an impulse or general sinusoidal input to the system, *f(t)*, and the resulting output, *g(t)*:

$$g(t) = \int_{-\infty}^{\infty} h(\tau) f(t-\tau) \mathrm{d}\tau \tag{6.2a}$$

This is often transformed from the time domain to the frequency domain through a Fourier transform, to give the relation:

$$G(\mathrm{j}\omega) = H(\mathrm{j}\omega) F(\mathrm{j}\omega) \tag{6.2b}$$

where $\mathrm{j} = \sqrt{(-1)}$, and ω corresponds to angular frequency, such that $\omega = 2\pi f$.

For a linear, time invariant system, the steady state response to a unit sinusoidal input is a sinusoidal output with a certain degree of *attenuation* and *phase lag*. So for example, for a generalised sinusoidal input:

$$f(t) = \exp(j\omega) \tag{6.3a}$$

the output is given by:

$$g(t) = H(j\omega)\exp(j\omega) \tag{6.3b}$$

with the transfer function given by:

$$H(\mathrm{j}\omega) = |H(\mathrm{j}\omega)| \exp(\mathrm{j}\psi(\omega)) \tag{6.3c}$$

The transfer function is itself typically expressed in terms of the *gain squared*:

$$|H(\mathrm{j}\omega)|^2 \tag{6.3d}$$

and the phase shift:

$$\psi(\omega) \tag{6.3e}$$

Several different analogue filter types are defined, typically according to the characteristics of the amplitude response (the phase response tends to be non-linear, although linear phase analogue filter designs are available). The two best known analogue filter types are the Butterworth and the Chebyshev filters. Butterworth filters present a flat response in the pass band, but the transition band is large; Chebyshev filters offer a more rapid fall off from pass to stop bands, but the gain over the pass band is not flat (it contains a ripple). Although canonical filter prototypes are usually presented with reference to a low pass characteristic, the other responses (high pass, band pass etc.) may be derived through the application of well-known transformations.

The particular analogue filter type I shall be concerned with is the maximally flat, low pass Butterworth filter, whose well known amplitude response is given by (Baher, 1990), equation (5.73) as:

$$|H(j\omega)|^2 = \frac{1}{1+\left(\frac{\omega}{\omega_c}\right)^{2n}} \tag{6.4a}$$

where n is the degree of the filter, ω_c the cutoff frequency. Increasing n reduces the width of the transition band and consequently improves the filter's cut off characteristic.

The phase response of a first order Butterworth filter with cutoff frequency 1 is given by (van Valkenberg, 1982) as:

$$\psi(\omega) = -\tan^{-1}\omega \tag{6.4b}$$

The response for other cut-off frequencies, ω_c, is obtained by applying the transformation:

$$\omega \rightarrow \frac{\omega}{\omega_c} \tag{6.4c}$$

6.2.2.2 Evolution as an Analogue Filter.

From an analysis by (Lande & Shannon, 1996) on an evolutionary model employing overlapping generations (as, for example, in a steady state GA), a continuous sinusoidally

varying evolutionary signal (optimal phenotype) with frequency b and amplitude A, is given by:

$$f(t) = A \sin bt \tag{6.5a}$$

This tracked for large t by an evolving population whose mean expressed signal, $g(t)$, represents a solution to equation (6.1h):

$$g(t) \approx \zeta(b) A \sin(bt - \beta) \tag{6.5c}$$

$$\text{where} \quad \zeta(b) = \frac{V_g}{V_s \sqrt{\frac{V_g^2}{V_s^2} + b^2}} \tag{6.5d}$$

for additive genetic variance, V_g; and variance, V_s, of *weak* selection towards $f(t)$.

Substituting with V (equation 6.1i), this gives:

$$\zeta(b) = \frac{V}{\sqrt{V^2 + b^2}} \tag{6.5e}$$

$$\text{and phase lag } \beta(b) = \cos^{-1} \zeta \tag{6.5f}$$

In terms of the signal processing model, this gives the gain transfer:

$$|h(t)| = |g(t)| / |f(t - \tau)| = \zeta \tag{6.6a}$$

Dividing the gain term, equation (6.5e), through by V, and then squaring, gives the following 'gain squared' term:

$$\zeta^2(b) = \frac{1}{1 + \frac{b^2}{V^2}} \tag{6.6b}$$

which is of a similar form to the independently used Butterworth amplitude response given in equation (6.4a). Translating between these models by substituting ω for b, the response derived by Lande & Shannon in equations (6.5) may now be characterised as describing:

a low pass, first order Butterworth filter ($n = 1$), with cutoff frequency V.

Note the gain never exceeds 1 (i.e. evolution is acting as a *passive* filter), and equals 1 for a constant input signal (i.e. $b = 0$).

A similar conclusion about the evolutionary transfer function is reached by reflecting on the phase response, equation (6.5f) (Rowe, 1997). From equations (6.4b and c), which gives the phase response of a first order Butterworth filter as:

$$\psi_B(\omega) = \tan^{-1}\frac{\omega}{\omega_c} \tag{6.6c}$$

and setting $\omega_c = V$,we can obtain:

$$V^2 + \omega^2 = V^2\left(1 + \tan^2 \psi_B(\omega)\right) = V^2/\cos^2 \psi_B(\omega) \tag{6.6d}$$

Incorporating this into Lande & Shannon's amplitude response, equation (6.5e), by substituting $b = \omega$, gives:

$$\zeta(\omega) = \cos \psi_B(\omega) \tag{6.6e}$$

Substituting this in turn into the expression Lande & Shannon give for the phase response, equation (6.5f), we get:

$$\beta(\omega) = \psi_B(\omega) \tag{6.6f}$$

That is, the phase response of the evolutionary filter corresponds to that of a first order Butterworth filter.

Interpreting this model, what is the effect of changes in genetic variance, V_g, and selection variance, V_s, on the filter cutoff, V? The cutoff represents the ratio of these two terms, so increasing either the selection strength (i.e. decreasing selection variance) or the genetic

variance (for example, through an increase in mutation rate) will increase the cut off frequency - the population will be better able to track faster moving optima. This is demonstrated using values given in (Charlesworth, 1993) for the sexual and asexual case, table 6-1.

V	Low pass cutoff/gens
0.005 (asexual)	1257
0.02 (sexual)	314

Table 6-1: Cut off frequencies (as period length) of the evolutionary low pass analogue filter for example variances given by Charlesworth.

6.2.3 Discrete, Non-Overlapping Generations.

Filters acting on a *discrete-time* or *sampled data* signal are classed as *digital filters*. Rather than employing a Fourier transform, analysis is presented using a z-transform which represents an analytical counterpart to the numerically applied *discrete Fourier transform* (Stark & Tuteur, 1979). However, for *causal* signals (which are signals that are only defined for time $t > 0$), the transformations resulting from each are identical (Baher, 1990). In this section, I demonstrate how a quantitative genetics analysis of an evolving system with discrete non-overlapping generations exhibits certain formal similarities to the response of a particular band pass digital filter.

6.2.3.1 A Brief Introduction to Digital Filter Theory.

As with analogue filters, digital filters are characterised in terms of their amplitude and phase response; so there exists, for example, digitally realised Butterworth or Chebyshev filters. One advantage over analogue realisations is that particular phase responses (either flat, or linear) may designed in. The transfer function of a digital filter may be realised either recursively or non-recursively. For a system with input sequence $\{f(n)\}$ and output sequence $\{g(n)\}$:

- a recursive filter generates the current output signal $g(n)$ as a function of $\{g(n-1), g(n-2),...; f(n), f(n-1),...\}$;
- a non-recursive filters derives $g(n)$ as a function of $\{f(n), f(n-1),...\}$.

Recursive filters are typically used to implement *infinite impulse response* (IIR) filters which have a maximally flat phase response. Non-recursive filters are typically used to realise *finite-duration impulse response* (FIR) filters, and are defined by Baher, equation (7.121) as:

$$g(n) = \sum_{r=0}^{N} a_r f(n-r) \tag{6.7a}$$

That is, the filter output represents a weighted sum of the current and previous N input samples.

The linear phase response of an antimetric (i.e. asymmetric) FIR filter, which is suitable for implementing approximate, ideal high and band pass filters, is given by Baher, equation (8.108), as:

$$\psi(\omega) = -\frac{\omega N T_s}{2} + \frac{\pi}{2} \tag{6.7b}$$

where ω is the environmental frequency, $\psi(\omega)$ is the phase response, N is the 'depth' of the filter (i.e. $g(n-0)...g(n-N)$ samples are used in finding $g(n)$) and $T_s = 2\pi/\omega_s$ is the sampling period for sampling frequency ω_s.

For an input signal to a digital filter to be completely recovered at its output, the *Sampling Theorem* states that the sampling frequency must be *at least twice* the highest frequency component contained within the signal:

$$\omega_s \geq 2\omega \tag{6.7c}$$

The frequency *exactly twice* that of the highest frequency signal component required to be passed is known as the *Nyquist frequency* and is designated ω_N. For a low pass digital filter, the sampling frequency must thus be at least twice the cutoff frequency. In the case of a filter presented with a band limited signal in the range ω_l to ω_h, the Nyquist frequency is given by:

$$\omega_N = 2(\omega_h - \omega_l) \tag{6.7d}$$

6.2.3.2 Evolution as a Digital Filter.

In his treatment of the evolution of recombination, (Charlesworth, 1993) presents a theory one may expect to be rather more in accord with traditional (i.e. *generational* as opposed to *steady state*) GA experiments, specifically the tracking of a sinusoidally varying environment with *discrete* generations. Equations (6.8) to (6.11a) below are taken from this source.

Appeals to digital filter theory in interpreting Charlesworth's work will be made in two ways: firstly, by comparing the form taken by the system input-output expressions; secondly, by comparing the phase responses derived from each model.

As in the continuous case, two measures of variation are required: the additive genetic variance, V_g, and a quantity $V_s = V_e + 1/S$, where V_e is a the environmental variance (arbitrarily set to 1) and S is the strength of selection. V_g is itself a function of V_s and the generational variance due to mutation (and hence the mutation rate; see for example, equation (6.1m)). Definitions of k (equation (6.1f)) and V (equation (6.1i)) also apply.

Comparing System Structure.

In a fluctuating environment, with constant equilibrium values for V_s and V_g, and optimal phenotype $f(n)$ in generation n, the mean phenotype in generation n, $g(n)$, is given by Charlesworth, in accord with equation (6.1g)[60], as:

$$g(n) = (1-k)^n g(0) + k\sum_{i=1}^{n}(1-k)^{i-1} f(n-i) \tag{6.8a}$$

$$\approx k\sum_{i=1}^{n}(1-k)^{i-1} f(n-i) \tag{6.8b}$$

approximating for large n. Note that the current mean is given here as only a function of the previous optima, and *not* the previous mean expressions, as was the case for equation (6.1g).

[60] Using a change of notation, equation (6.1g) states: $g(n) = (1-k)g(n-1) + k.f(n-1)$.

One reason for developing the ongoing experimental work reported in 6.3, below, is to check the range over which the many approximations used in deriving this expression are likely to be valid. Also note that if within generation learning is allowed, then *g(n) may* be a function of the current environmental state, *f(n)*.

Now, inspection of the form of equations 6.7a and 6.8b shows them to be similar, with the exception of the limits of the summation. If some relation can be drawn between these limits, then it would seem that evolution with discrete, non-overlapping generations implements an FIR filter, with filter coefficients:

$$a_0 = 0; \tag{6.9a}$$
$$a_r = k(1-k)^{r-1} : r > 0 \tag{6.9b}$$

In the digital filter expression, the limit N represents the depth of the filter; that is, the number of previous samples (*plus* the current one) that it may use in calculating the current output. In the evolutionary analysis, the limit n represents the current generation number and the summation thus allows all *previous* generation optima to influence the current population mean. The approximation of $(1 - k)^n$ to zero suggests that $(1 - k)^m$ for m slightly less than large n will result in a negligible contribution to the overall mean by $f(n - m)$. This relaxes the constraint on the upper bound of the summation of (6.8b) being strictly n, but the question as to what is an appropriate setting remains[61]. It is possible that the upper bound may be set in relation to the selection strength and genetic variance: under weak selection and low genetic variance, (i.e. low k) there are likely to be some remnants of temporally distant populations in the current population (e.g. previously optimal individuals). In this case, the upper bound is likely to be *high* - distant optima will have some influence on the current population mean. For high k (strong selection, larger genetic variance) there will be little 'population memory' of earlier adapted phenotypes and only the most recent optima will feature in deciding the current mean - i.e. the upper bound on the summation will be *low*.

If the analogy between (6.7a) and (6.8b) does hold, then where the mutation rate is allowed to evolve, V_g and hence k will be non-stationary (see also equation (6.1m) which describes

---

[61] A 'pragmatic' setting for the upper bound, $n = M$, may be set as follows (Rowe, 1997): for $|f| < 1$, M may be chosen such that $(1 - k)^M \leq \delta < (1 - k)^{M-1}$, where δ is a small, appropriately chosen value, and hence $M = \lceil \log \delta / \log(1 - k) \rceil$. For small k, this gives $M \approx \lceil \log(1/\delta)/\log k \rceil$.

the equilibrium population variance where the mutation rate is non-zero). That is to say, the filter coefficients will be self-adaptive and an *adaptive filter* will be realised. This will be discussed further in 6.2.5.

Comparing Transfer Function Characteristics.

In a sinusoidally varying environment with non-overlapping generations:

$$f(n) = A\cos(\omega n) \tag{6.10a}$$

$$\text{where} \quad \omega = \frac{2\pi}{T} \tag{6.10b}$$

For the sinusoidal input given by equation (6.10a), Charlesworth approximates equation (6.8b) under conditions of small V (specifically $V << \omega$, presumably arising from weak selection) by an integral with solution:

$$g(n) \approx \frac{2\pi TAV\sin(2\pi(n-1)/T)}{4\pi^2 + V^2T^2} \tag{6.11a}$$

Substituting with equation (6.10b), we can reduce this to:

$$g(n) \approx \frac{\omega AV\sin(\omega(n-1))}{\omega^2 + V^2} \tag{6.11b}$$

We are now in a position to interpret this expression in the light of digital filter theory. The gain squared term and phase shift, ψ, of the 'filter' are given by:

$$|H(j\omega)|^2 \approx \left(\frac{V}{\omega} + \frac{\omega}{V}\right)^{-2} \tag{6.12a}$$

$$\text{and} \qquad \psi(\omega) = -\omega + \frac{\pi}{2} \tag{6.12b}$$

The amplitude response of the system limits the amplitude of the population mean to a maximum of *half* the amplitude of the input signal, and also gives a characteristic typical of a

band pass filter[62]. The maximum value the gain (squared) term can take is simply 0.5 (0.25) and is located at the mid-point of the pass band (the *band center*):

$$\frac{d|H(j\omega)|^2}{d\omega} = \frac{d\left(2 + V^2\omega^{-2} + \omega^2 V^{-2}\right)}{d\omega} = 0 \tag{6.13a}$$

and so $\omega V^{-2} - \omega^{-3}V^2 = 0$ (6.13b)

giving $\omega = V$.

Further, taking the 3dB point to be 3dB down on the maximum possible gain squared of 0.25, the two cut-off points are located at the solution of:

$$|H(j\omega)|^2 = \frac{0.5^2}{2} = 2^{-3} \tag{6.14a}$$

Using the expression for the gain squared given by equation (6.12a), this gives:

$$\frac{V}{\omega} + \frac{\omega}{V} = 2\sqrt{2} \tag{6.14b}$$

and hence:

$$\omega^2 - 2\sqrt{2}V\omega + V^2 = 0 \tag{6.14c}$$

i.e. $\omega = V\left(\sqrt{2} \pm 1\right)$ (6.14d)

which gives a pass band width of 2V.

Given that the sampling frequency is set to $\omega_s = 2\pi$, we can apply the Sampling Theorem, equation (6.7c), to identify the *maximum possible* value of the high cutoff frequency:

$$\omega_{hmax} = \pi \tag{6.14e}$$

---

[62] Indeed, the expression for the gain squared resembles a standard low pass to band pass transformation.

This in turn limits the ratio of additive genetic and selective variances to:

$$V_{\max} = \frac{\pi}{\sqrt{2}+1} \approx 1.3 \tag{6.14f}$$

Using values suggested by Charlesworth (V_s =20, V_g = 0.4 for a sexual population, 0.1 for an asexual population) this gives the cutoff frequencies (as generational periods)given in table 6-2. What this appears to suggest is that evolution is only capable of tracking fluctuating signals that have a period of between 100 and 3000 generations. A failure to track more rapidly fluctuating environments effectively is easily understood, but the band pass characteristic - and the consequent inability to track slowly moving or even stationary optima - appears anomalous. Given the low pass response predicted by the continuous model of (Lande & Shannon, 1996), the two together suggest that Charlesworth's model is not applicable at frequencies less than the mid-point frequency. A further comparison of the models is presented in 6.2.4.

V	Low cutoff, ω_l/gens	High cutoff, ω_h/gens	Mid-point, ω_h/gens
0.005 (asexual)	3034	520	1257
0.02 (sexual)	758	130	314

Table 6-2: Cut off and mid-point frequencies (as period) for the bandpass evolutionary digital filter for example variances given by Charlesworth.

The phase shift (equation (6.12b), not shown) is a linear function of the environmental frequency with an additional constant shift of 90 degrees (not 180 degrees as Charlesworth states). This contrasts with the continuous result for which the phase change was also a function of the selection strength and genetic variance (i.e. V in this case). Recalling (6.7b), the phase response of the evolutionary system is of a similar form to that of an asymmetric FIR filter. Taking the comparison to a limit, and equating these expressions, the following evolutionary filter characteristics are identified:

for a sampling period $T_s = 1$ generation, this fixes the filter depth $N = 2$.

Now, this stands in stark contrast to the time dependent, increasing depth suggested by (6.8b), not least because it limits the contribution from earlier generations so drastically. If we accept this result, then its interpretation suggests that selection must be very strong, which may contradict the assumptions used in generating equation (6.11a).

A further problem with discrete models is that they admit the possibility of aliasing. If the input signal to the digital evolutionary filter contains frequency components above the Nyquist frequency, then low frequency echoes of the high frequency signal may be passed down into lower frequency bands. There is thus a 'false' signal component that really exists and may be tracked by the population. The extent to which natural systems sample their environment through reproduction is unlikely to be periodic and discrete (although there, may be exceptions...). For example, if selective reproduction is an accurate reflection of fitness, then the fitness of individuals in an evolving population is unlikely to be evaluated a) instantaneously, and b) at discrete, periodic intervals. However, in the generational GA model, aliasing may be a very real problem, This may be avoided either by using a steady state GA, or by 'dithering' the signal applied to each individual; i.e. by adding a small random noise component to the signal during each individual's evaluation.

To recap, there is partial evidence on two counts (although one supervenes the other) that Charlesworth's analysis of evolution in sinusoidally fluctuating environments with non-overlapping generations describes an asymmetric FIR filter - the structure of the transfer function, and the equivalent phase response. Interpreting this result, however, leads to certain anomalies - the low amplitude of the population mean in the centre of the pass band; the band pass response itself and the setting of the cut-off frequencies; and the derived filter settings from the phase response. Note that these results are only anomalous if we assume that evolution will act to optimise the cutoff frequencies and mid-point in some way. Whilst this is an important question within the context of ESP as a model of evolution, its resolution will not affect the undoubted usefulness of the ESP approach for furthering one's understanding of the models produced elsewhere.

6.2.4 Comparing Continuous and Discrete Models of Evolution.

A direct comparison of the characteristics of the transfer function derived from the models of (Charlesworth, 1993) and (Lande & Shannon, 1996) suggests that they are not commensurate in predicting either the frequencies at which the evolutionary filters described are most effective, or the low frequency behaviour. In addition, the magnitude of attenuation at high frequencies is significantly different. This is clearly demonstrated in figure 6.2. The common quantity $\gamma\sigma^2$ does, however, represent a distinguished frequency component within each model (mid-point frequency firin Charlesworth's model, cut-off frequency in Lande & Shannon's). It seems likely that this quantity will play an important role in any adaptationist interpretation of the respective models, for example as the effective target of selection.

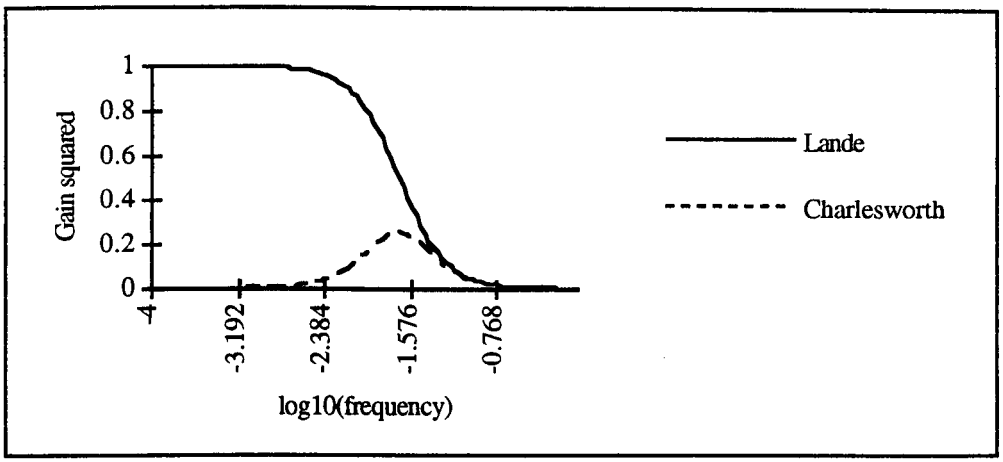

Figure 6.2: A comparison of the evolutionary filter responses predicted by (Lande & Shannon, 1996) and (Charlesworth, 1993) for an asexual regime.

In addition, the signal processing interpretation illustrates a significant common sense difficulty with Charlesworth's discrete model: through the band pass characteristic, evolution would appear *not* to track *low* frequency components. Charlesworth admits the model does not apply at low frequencies, but the analysis given above suggests that this is because the form of the derived response (a band pass characteristic) is not generally appropriate. Comparison with the analogue model suggests that the range of applicability of Charlesworth's model should be explicitly limited to explaining the tracking of non-stationary signals with significant frequency (specifically, greater than the mid-point). This restriction is in accord with the approximations used to derive the model (not applicable beyond (unspecified) low or high frequencies) but results from a straightforward interpretation of the model rather than a deep-seated familiarity with its derivation.

Finally, the filter depth apparent in the discrete model is set at the low value of two. That is, only the current and previous generation are taken into account when calculating the population mean. This would suggest that selection is strong and hence that there is little carry over from temporally distant ancestral populations. This interpretation is counter to Charlesworth's assumption that selection is weak.

6.2.5 Evolution as a Filter of Adapted Individuals.

Recall the amplitude response of an ideal band pass filter given in figure 6-1. Input frequencies within the pass band are passed unattenuated, frequencies falling outside the pass band are completely blocked. Now, the evolutionary filter is capable of transmitting information about which states are adapted to the current environment between one generation and the next by means of the inheritance (transmission) system. This information is transmitted on the basis of the relative adaptedness of the individuals within the population, and the particular mode of inheritance supported. In contrast to the received view of populations converging on a local optimum, we shall utilise the view of evolutionarily stable population clouds, or *quasi-species* (Eigen et al., 1989), that tend to gather *around* an optimum, (e.g. for the asexual case, the population may be held in check around an optimum by the mutation-selection balance (Woodcock & Higgs, 1996)). The evolutionary filter now acts as follows - *individuals* are passed between one generation and the next on the basis of their selective value. The product of an individual's current generation proportion within the population and its selective value, which describe its *expected* proportion in the next generation define the evolutionary filter response over the search space[63]. The 'width' of the quasi-species thus defines the extent of the pass band of the evolutionary filter. By allowing the mutation rate to change self-adaptively, but keeping the range of selective values constant (for example, by using rank based selection) the width of the quasi-species -the width of the filter pass band - may change. This is the method proposed for tuning the evolutionary filter that I shall describe below. Where the input signal fluctuates, and the population is able to track this signal whilst maintaining its level convergence, this corresponds to translating the pass band over the search space frequency.

The filtering properties of the population faced with the sinusoidally varying input signal are a consequence of the following: for higher input frequencies, the optimum moves through the search space rapidly. If the population is to track the signal, at least one individual in the population must lie within its basin of attraction. For high frequency input signals, the population must cover a large volume of the search space if it is to anticipate the position of

---

[63] The expected proportion over current proportion represents the gain of the evolutionary filter for the particular phenotypic 'frequency'. This resembles the interpretation of the selective value of an individual as a signal used in (Kargupta & Goldberg, 1994).

the optimum in the next generation. For lower frequency signals, even a highly converged quasi-species will be able to track the optimum. Following this line of reasoning, if a 'narrow' population cloud is faced with a rapidly changing optimum, it will fail to track the signal at all well and the population mean expression will represent a significantly attenuated version of the input signal.

6.3 Evolutionary Signal Processing in Genetic Algorithms

Over the last decade, GAs have been deployed with some success over a range of stationary optimisation problems. On well defined, temporally varying optimisation functions, however, there has been little work done. Where the environment is non-stationary, the use of convergence around an individual to identify when an optimum has been found is not appropriate, since individuals are assumed to be specialists rather than generalists, and so the population can never settle on a single individual that is optimal over a period of time. In the previous section, I described how evolution may act as a well-behaved signal processor. In this section, I will present an incremental GA experiment in which the evolutionary algorithm exhibits a similar set of behaviours.

The analyses offered above employed a sinusoidal input signal. Such an evaluation function has already been described in the GA literature by Cobb (Cobb, 1990); this function was also used by Dasgupta (Dasgupta, 1994) as an initial testbed for his structured genetic algorithm. In its original formulation, the evaluation function optimum changed *between* generations and was given by minimising the error between an individual, p, and a time varying optimum, h_t,:

$$h_t = 10 + \sin(\alpha \times Generation) \tag{6.15a}$$

$$e_t(p) = (p - h_t)^2 \tag{6.15b}$$

$$f_t(p) = -\log_{10} e_t(p) \tag{6.15c}$$

where p is a 32 bit Gray coded phenotypic individual over the range [0.0, 2.0], derived from 32 genotypic 'targeting' bits (as opposed to 'rate' bits which specify individual mutation rates, for example); h_t gives a 32 bit environmental state that varies sinusoidally

over time, changing between generations, with rate parameter, α, and period $2\pi/\alpha$; the individual error is given by e_t and this value is used for selection; f_t is the reported evaluation value (scaled for convenience, and apparently recasting the experiment as a *maximisation* problem).

Previously reported results using this evaluation function (Cobb, 1990; Dasgupta, 1994) have concentrated on the ability of the *fittest* population member at any one time (i.e. the one with the lowest error evaluation) to track the environmental state. Taking a cue from the theoretical/quantitative genetics approaches discussed above, I will focus on the behaviour of population means, although best of generation results are also reported in line with the previous work.

Cobb identified two strategies for coping with fluctuating environments - memory (e.g. through diploidy) and introducing variation at the *genotypic* level through *mutation*. An additional strategy, not considered by Cobb, but nevertheless likely to prove beneficial, is to introduce *phenotypic* variation through plasticity.

In this chapter and the next, the effect of three adaptive strategies investigated in *stationary* environments for evaluation landscapes of various degrees of ruggedness will be considered:

- an individually evolved mutation rate, that sets the bit rate at which an individual is mutated at the start of a generation (rate loci are not subject to the evolved rate);
- individual learning at an individually evolved 'learning' rate (this loosely corresponds to the *bias* of Turney (Turney, 1996));
- 'cultural inheritance' or the inheritance of characters acquired through learning.

6.3.1 Tracking Fluctuating Environments with Mutation.

In Cobb's original experiment, she compared the effectiveness of fixed mutation rates and an adaptive mutation rate strategy on the basis of offline (time averaged generational best) fitness. The adaptive mutation strategy switched from a low to a high globally applied mutation rate whenever the offline fitness dropped for a range of environmental velocities. For comparison, I shall use a simple GA with fixed and evolving mutation rates (m-rates

(units bit^{-1})) and will report offline and online (time averaged population best and mean) 'fitnesses' (i.e. evaluations), although I would suggest that these are not necessarily the most appropriate measures when 'settling time' (initial exploration) and equilibrium (steady state) modes are separated out. The self-adaptive mutation rate is a 7 bit binary coded value over [0.0, 1.0], and is used in the biased flipping of each targeting bit following crossover (i.e. the evolved mutation rate is not applied to the bits that code for the rate itself). In the experiments reported herein, mutation at the low background rate is applied equally to all bits in the genome (i.e. targeting and evolved rate bits alike). However, more careful tuning of the rate is possible by reducing the background rate to zero and employing a fixed m-rate that is applied *only* to the modifier gene. In this way, high levels of variation may be introduced into the modifier gene to guarantee a rapid response to a change in environmental conditions.

All experiments were carried out using a modified version of Genesis 5.0. Unless otherwise specified: graphed results represent the mean of 10 runs, with x-axis representing generation number. Given the time varying nature of the evaluation function, time averaged results are given in many cases, taken over a number of generations within a run. The time averaged results are then themselves averaged over runs.

The population size was 200; the background mutation rate was set at 0.001/bit in all cases; two point crossover was applied at a rate of 0.6 per individual to individuals selected using linear ranking selection with rank minimum 0.5. As in chapter 5, convergence is measured as the bitwise convergence *over the whole individual* (i.e. including the modifier rate bits).

6.3.2 To Compare Simple GAs with Fixed and Evolvable Mutation Rates.

In this first experiment, I shall concentrate on Cobb's original results and compare them to results obtained from an algorithm employing individual, evolved mutation rates. Regime GA-M corresponds to the evolved m-rate case; SGA corresponds to a simple genetic algorithm, with fixed m-rates, set at the 'optimal' mutation rates suggested by Cobb as follows -(α:m-rate): 0.001:0.01, 0.01:0.05, 0.05:0.1, 0.1:0.5, 0.5:0.5. These correspond to the environmental periods 6283, 628, 126, 63 and 13 respectively.

Figure 6-3a shows how the optimal fixed rate (SGA) offers better offline performance than the evolved rate. For the evolved m-rate regime, GA-M, the offline fitness is at best only

comparable to, and frequently worse than, that of a random population. The reason is that the population is converging around a point that lags the optimum thus limiting the available variation and as a result, offline fitness. Note how the steady state time average values offer a clearer picture of the equilibrium behaviour.

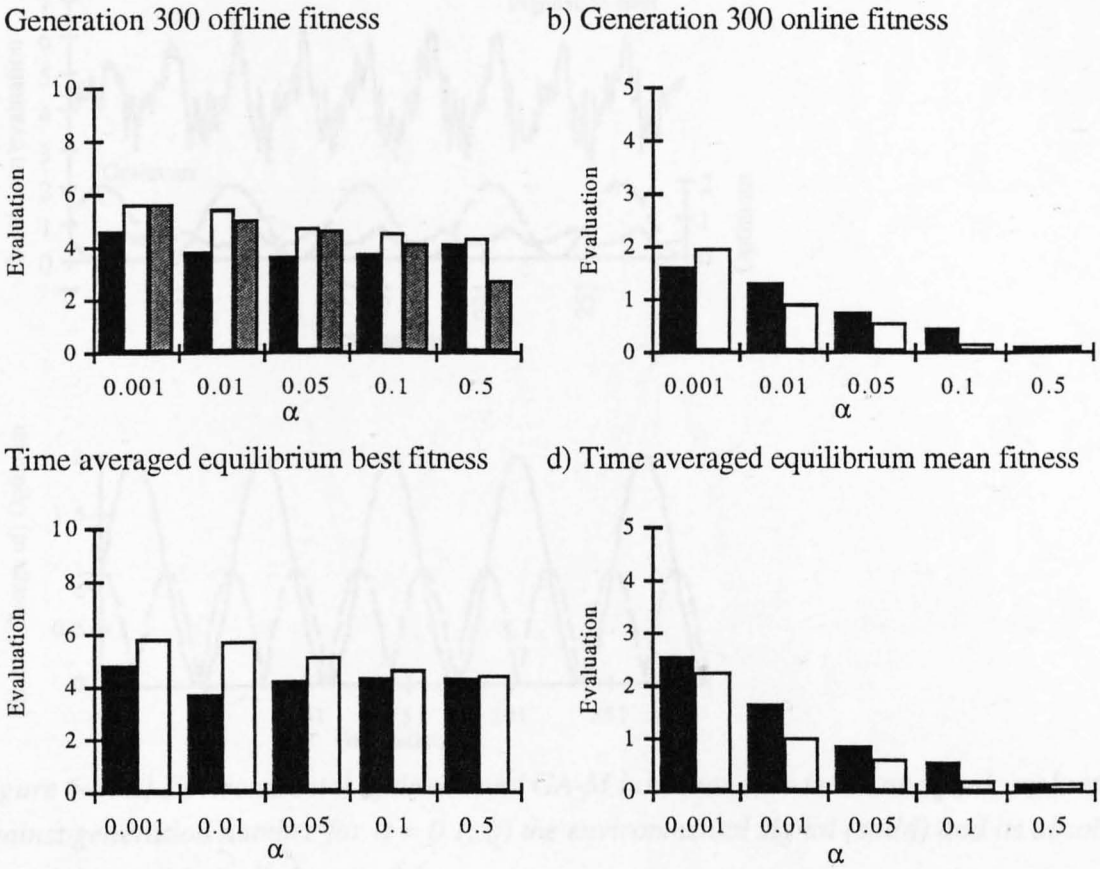

Figure 6-3: Fitness values against environmental rate, α, for GA-M (solid), SGA (empty), and Cobb's reported results (hashed) a) Mean offline fitness at 300 generations; b) mean online fitness; time averaged best (c) and mean (d) evaluations taken over the last 126 generations.

Looking at the instantaneous mean and peak evaluations for GA-M for a particular environmental rate, (specifically, $\alpha = 0.1$, figure 6-4a), the scores oscillate with a period half that of the environmental signal (i.e. twice the frequency) with the mean performance lagging the peak performance. Comparison of the phase of these fitness traces with that of the environmental optimum $h(t)$, suggests that the population best fitness peaks just after the rate of change of $h(t)$ is at a minimum (i.e. when the absolute value of the gradient of $h(t)$, that is $|\cos(\alpha t)|$, figure 6-4b, is at a minimum). Close inspection of the mean evaluation trace

suggests that rather than describing a sinusoid it describes an inverted, rectified sinusoid; that is, the population mean dynamic tracks the rate of change of the input signal, climbing for decreasing rates of change.

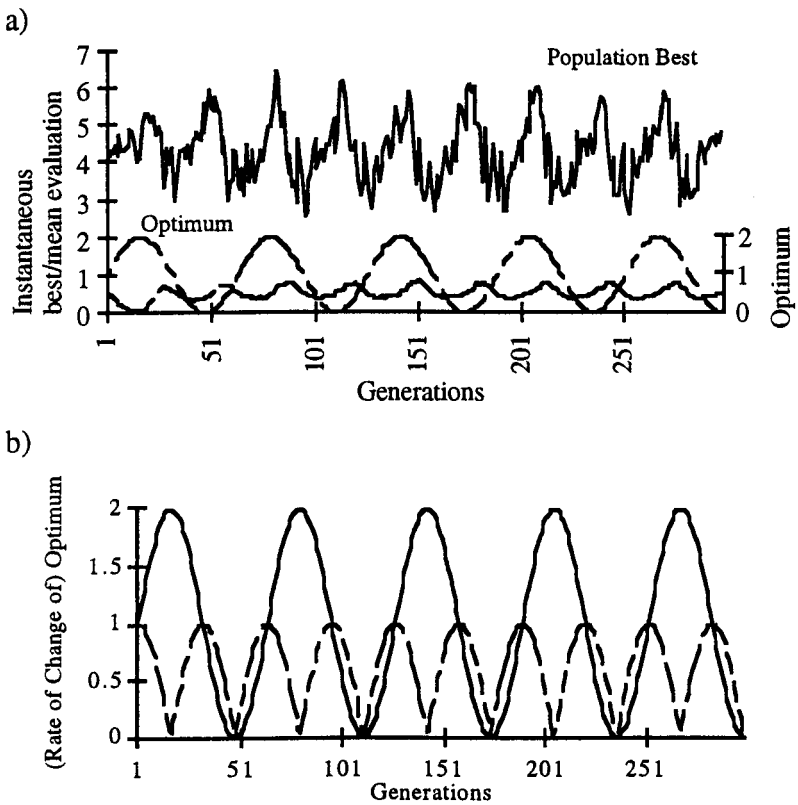

Figure 6-4: a) Environmental optimum and GA-M instantaneous mean and peak evaluation against generation number for $\alpha = 0.1$; b) the environmental signal (solid) and its absolute rate of change (dashed), for $\alpha = 0.1$;.

As far as the actual evaluations go, once the system has settled the time averaged means seem to be on a par with the scores recorded by the initial random population. The equilibrium convergence of each population, however, varies considerably over the environmental frequency within each condition, and across conditions for any particular frequency (figure 6-6). Equilibrium convergence levels, then, appear to be dependent on the mutation rate, as predicted by equation (6.1m).

Whilst the improvement in online fitness of GA-M over the fixed, 'optimal' rate may seem slight, just how significant the different approaches are in terms of the ability of the population *as a whole* to track the environment is revealed by the behaviour of the mean

population *phenotype*, figures 6-5, which depict single runs (the behaviour being typical); see also figure 6-3d. The SGA requires a high, fixed mutation rate to source enough variation

a) SGA, α= 0.1

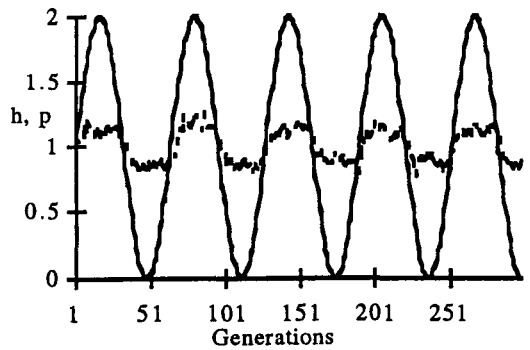

b) GA-M, α= 0.1

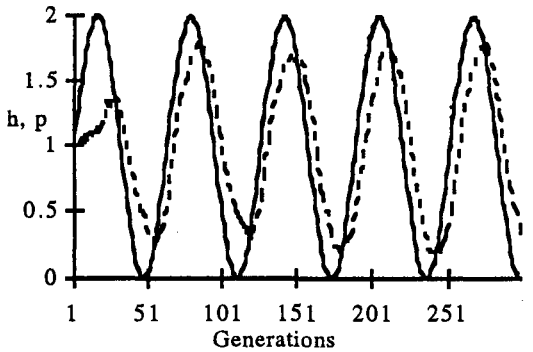

c) SGA. α= 0.05

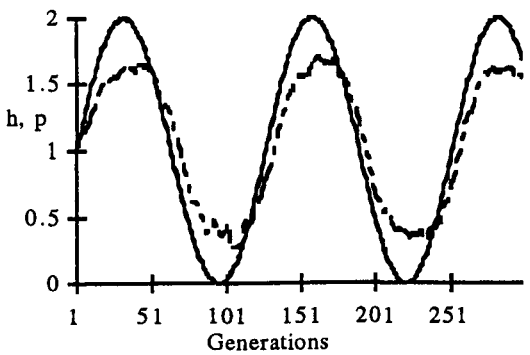

d) GA-M, α= 0.05

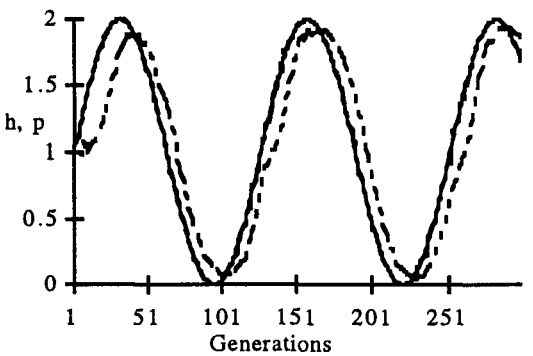

e) SGA with lower m-rate, α = 0.1

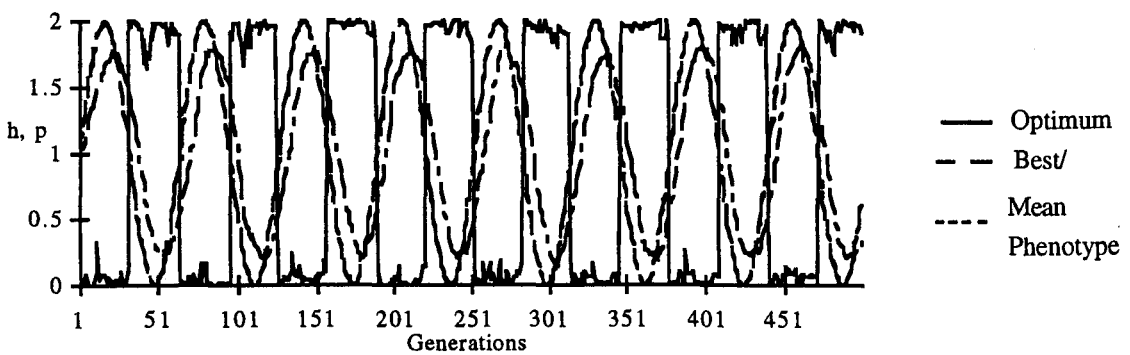

Figure 6-5: Tracking a moving environment over 300 generations- a) and b), α = 0.1; c) and d), α = 0.05. In all cases, the trajectory of the environmental value over time is very closely followed by the best adapted phenotype of each generation. The solid line represents the environmental target, almost ideally tracked by the best of each current population. The dashed lines represent the ***mean phenotypic expression*** *of the population during the current generation - a) fixed m-rate of 0.5; b) evolved m-rate + 0.001 background mutation;*

c) fixed m-rate of 0.1; d) evolved m-rate + 0.001 background mutation; e) Longer run behaviour, for fixed m-rate 0.045, $\alpha = 0.1$; worst individual expression also shown.

for the efficient *offline* tracking of a rapidly moving environment; however, this high m-rate has a detrimental effect on the ability of the population as a whole to track the environment, figure 6-5a. For the case of GA-M, whilst offline performances are down (though not noticeable on the trace shown), the population mean is able to track the environment far more efficiently, figure 6-5b. In a slower environment, where a lower rate of mutation is applied, mean tracking behaviour is much improved in both cases; and although the online and offline fitnesses are worse for GA-M compared to SGA, a) the final generation population convergence (measured as mean proportion of the majority bit at each locus over the population, increasing for decreasing m-rate) is higher (figure 6-6), and b) mean tracking is better, for all values of α. Also note in figure 6-5e how the *worst* individual in each population describes a noisy square wave completely out of phase with the optimum. Whether this behaviour can be characterised in terms of (an extension to) the defining filter equation is left as an open question.

Turning now to the degree to which the populations converge. Firstly, it appears that in general the degree to which a population is stably converged (i.e. the equilibrium convergence value) falls with increasing *environmental* rate. However, inspection of figure 6-6b suggests that in fact the degree of convergence is actually reflecting the mutation rate since for the two fastest moving environments, where the *same* fixed mutation rate is implemented, the convergence levels are similar. For the rapidly changing environment, the initial increase in convergence is largely a result of the convergence of the rate gene (convergence is measured over *all* genotypic bits).

It is interesting to compare the phase of generational mean fitness, convergence and mutation rate for a single environmental condition, figure 6-6c. In particular, the mean and convergence traces are in phase, whereas the mutation rate is 180 degrees out of phase (since the higher the value of -log(mutation rate), the lower the actual rate). With the mutation rate at a minimum, the mutation-selection balance drives the population closer towards the optimum; consequently, the bitwise convergence and mean fitness are close to their maximal values.

a) GA-M convergence

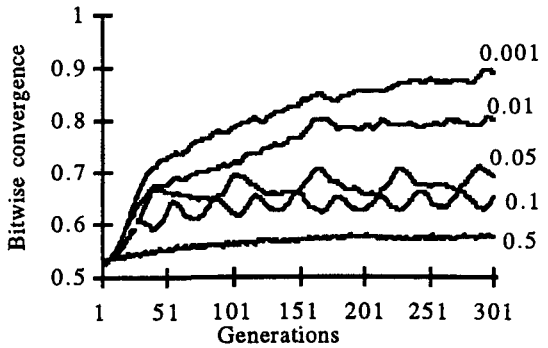

b) SGA convergence

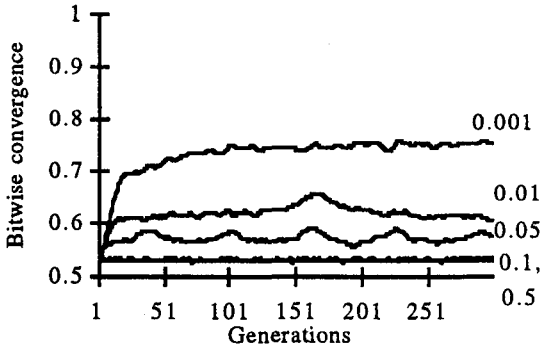

c) Phase comparison of population mean expression, convergence and evolved mutation rate

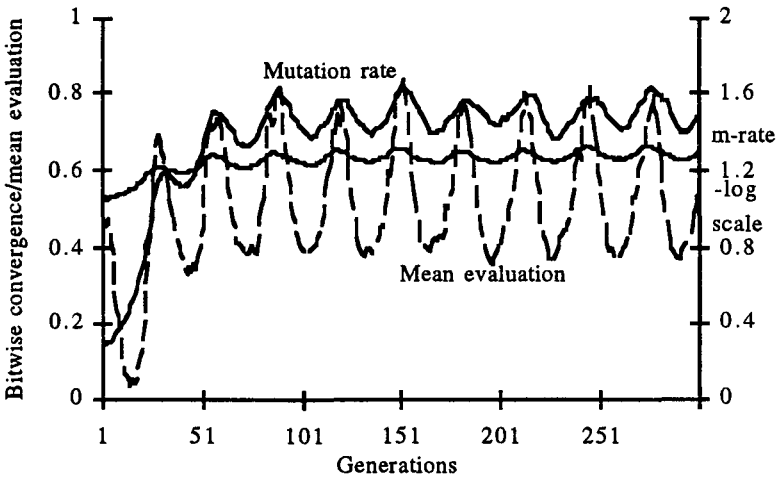

Figure 6-6: Decrease in convergence for increasing environmental rate 0.001, 0.01, 0.05, 0.1, 0.5: a) regime GA-M; b) regime SGA; c) phase comparison of convergence (middle trace), population generational mean fitness (lower trace) and (-log10) mutation rate (upper trace, right hand scale) for GA-M, a = *0.1.*

Figure 6-7 shows the mean instantaneous evolved mutation rate over α and table 6-3 compares the time averaged final generation mutation rate to the values selected by Cobb. Clearly, her intuition was correct in that higher mutation rates are required to track the environment successfully: increasing α increases the evolved mutation rate, although only as much mutation as is necessary is supported. The far lower evolved rates explain the better mean population tracking behaviour of GA-M. Note how for intermediate α the raw evolved mutation rate (figure 6-7) oscillates at twice the environmental rate as did the instantaneous mean and peak fitnesses, and for a related reason: the absolute rate of change of the sinusoidal input, $|\cos(\alpha t)|$, which is what is ultimately tracked by the mutation rate, attains its maximal value twice in every environmental period. For low environmental rates, the rate of

change is so low that the mutation rate meanders at random around a very low value; for rapidly changing environments, the selection pressure on the modifier is not strong enough to permit tracking and so the mutation rate settles at a steady equilibrium value.

α	0.001	0.01	(0.025)	0.05	0.1	(0.25)	0.5	(1.0)	(5.0)
SGA	0.01	0.05	N/A	0.05	0.5	N/A	0.5	NA	N/A
GA-M:									
1-300	0.023	0.031	0.034	0.044	0.061	0.136	0.238	0.287	0.231
150-300	0.002	0.004	0.008	0.016	0.032	0.102	0.194	0.200	0.128

Table 6-3: Evolved mutation rates over α with 4 additional rates (bracketed); mean rates given are taken over the whole run and the last 150 generations (this latter representing an equilibrium rate). For GA-M, there is an additional background rate of 0.001/bit.

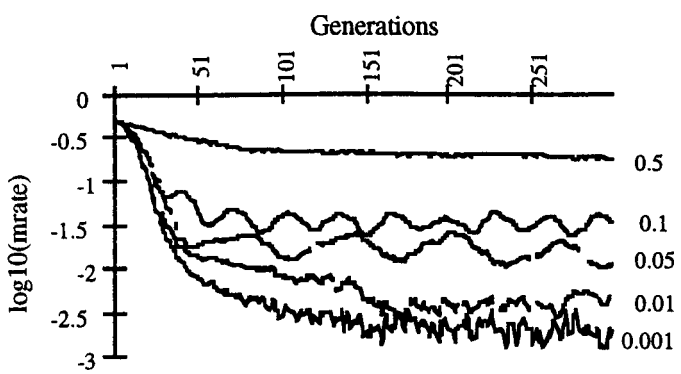

Figure 6-7: Mean instantaneous evolved mutation rates against generations for GA-M, α = 0.001, 0.01, 0.05, 0.1, 0.5.

Given that the variance of the steady state, evolved mutation rates is not too great, using the time averaged steady state rate as a fixed rate should allow us to tune the evolutionary filter for the particular input signal frequency. This intuition - that evolution sets the self-adapted mutation rate to some sensible value - may be strengthened by identifying which measures are apparently being optimised by means of the self-adaptive rate. This is achieved by seeing what happens to the fitness measures when different fixed m-rates of similar magnitude to the evolved rate for a particular α are applied. Table 6-4 shows online and offline fitness values for fixed m- rates in the vicinity of the evolved m-rate for α = 0.1, based on the value attained under self-adaptation (table 6-3 - specifically, the time averaged evolved rate was 0.061/bit taken over generations 1-300, 0.032/bit over generations 150-300; the first value represents the mean rate over the whole of the run; the latter the mean equilibrium rate).

Whilst offline fitness improves for increasing fixed m-rate, as an increasing amount of variation is available, the online performance peaks for m-rate 0.045, the mid-point of the two time averaged evolved values given in table 6-3.

m-rate/bit	0.015	0.030	0.045	0.060	0.075
Online fitness	0.49	0.55	0.56	0.55	0.53
Offline fitness	2.51	3.11	3.54	3.81	4.06

Table 6-4: Online and offline fitness versus fixed mutation rate in the vicinity of the evolved rate for $\alpha = 0.1$, selection rank minimum 0.5.

Further consideration of traces shown in figure 6-5 reveals both an attenuation of, and a phase lag between the mean value expressed by the population and the environmental value. I have already noted how increasing the mutation rate affects the stability of the population and its ability as a whole to track the environment - for higher mutation rates, the mean phenotypic value of a population fails to track the environmental target.

Changes in gain (or attenuation - that is, the peak-to-peak value of the tracking signal compared to that of the optimum) and lag (the phase difference in generations between tracking and tracked signals) result from altering the genetic variance, as regulated by the mutation rate. Table 6-5 and figure 6-8a shows how phase lag is affected by m-rate. Mean values are mean delays in generations between related target and population mean values passing through 1.0, discounting the first proper cross, and taken over 10 independent runs. Specifically, lag is decreased by a slight increase in the m-rate (±0.01 on a rate of 0.045), although only by a generation or so. For larger changes in the mutation rate, tracking properties of the population are altered and useful direct comparison becomes difficult. The amplitude of the mean expressed phenotype is similarly a decreasing function for increasing m-rate, although the effect on mean tracking ability is the converse to that of the lag. That is, as m-rate increases, whilst the lag between mean phenotype and the environmental target is reduced, the attenuation of the mean increases (i.e. there is a *worse* fit between the time delayed environmental signal and the mean value).

m-rate/bit	0.015	0.030	0.045	0.060	0.075
peak to peak	1.71	1.64	1.56	1.45	1.38
mean lag	7.70	6.39	5.74	5.15	4.80
lag sd	0.44	0.38	0.32	0.27	0.36

Table 6-5: Regime SGA - selection rank minimum fixed at 0.5; $\alpha = 0.1$; change mutation rate from 0.015 to 0.075 step size 0.015.

If the mutation rate is now fixed and the selection strength is varied, the effect on phase and amplitude is rather more pronounced (figure 6-8b and table 6-6). As the selection strength increases (reducing rank minimum) both the amplitude and the phase response improve. Recalling the response predicted by Lande, this is to be expected. However, the system response to the varying mutation rate appears anomalous (one would expect *improved* performance for both amplitude and phase for increasing rate). This failure in the predicted amplitude response possibly results from the increased population spread (decreased convergence, sparse distribution) arising out of the increased mutation rate, compared to the normal distribution assumed by Lande & Shannon.

rank min.	0.1	0.3	0.5	0.7	0.9
Online fitness	0.81	0.70	0.56	0.39	0.19
Offline fitness	3.39	3.48	3.54	3.71	4.09
mean lag	4.00	4.80	5.74	7.28	8.75
lag sd	0.27	0.28	0.32	0.21	0.99
peak to peak	1.75	1.68	1.56	1.29	0.59

Table 6-6: Regime SGA - mutation rate fixed at 0.045; $\alpha = 0.1$; selection parameterised by rank minimum, from 0.1 to 0.9 step size 0.2.

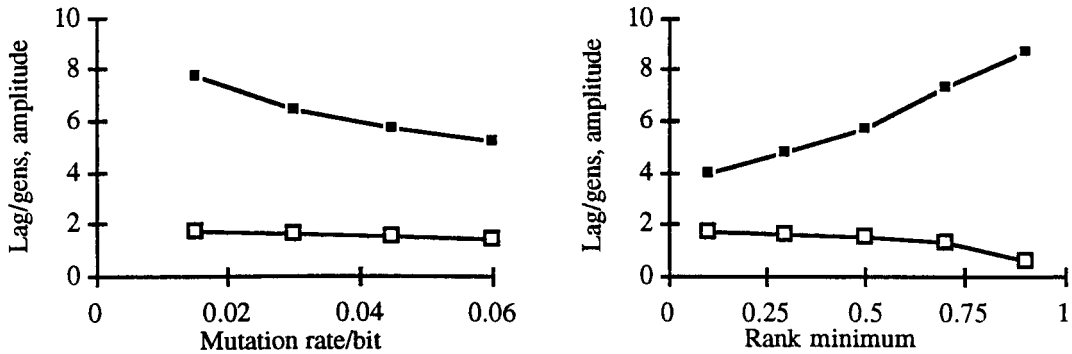

Figure 6-8: a) Equilibrium lag/gens (solid squares) and amplitude (empty squares) vs. m-rate/bit for linear rank selection minimum 0.5; b) Equilibrium lag/gens and amplitude vs. linear rank selection minimum for m-rate 0.045/bit.

The gain of the population mean against environmental frequency (figure 6-9a) has the form of a low pass filter. This is qualitatively in agreement with the predicted low pass filter response described in 6.2.2.2 above. The effect of *increasing* the mutation rate appears commensurate with *decreasing* the order of the filter, although this is not supported by Lande's theory and may just be an artifact of the two responses shown. The phase response obtained via equation 6.5f applied to the measured gain, behaves appropriately for low and high frequences, although there is some discrepancy for the intermediate frequencies. It lies to further work to establish the *quantitative* relationship between these two models more thoroughly (although I describe a possible approach to this problem in 6.5 below through the use of a self-adaptive mutation rate in an environment with a temporally dependent frequency).

6.4 Applications of ESP

In this section, I shall offer two examples of how evolving populations can process signals in rather more meaningful ways - extracting noise from an input signal, and a prefilter for distinguishing between two frequencies. Although the performance of the evolutionary filter is unlikely to compete with traditional signal processing techniques, these examples do illustrate that the ESP metaphor is robust.

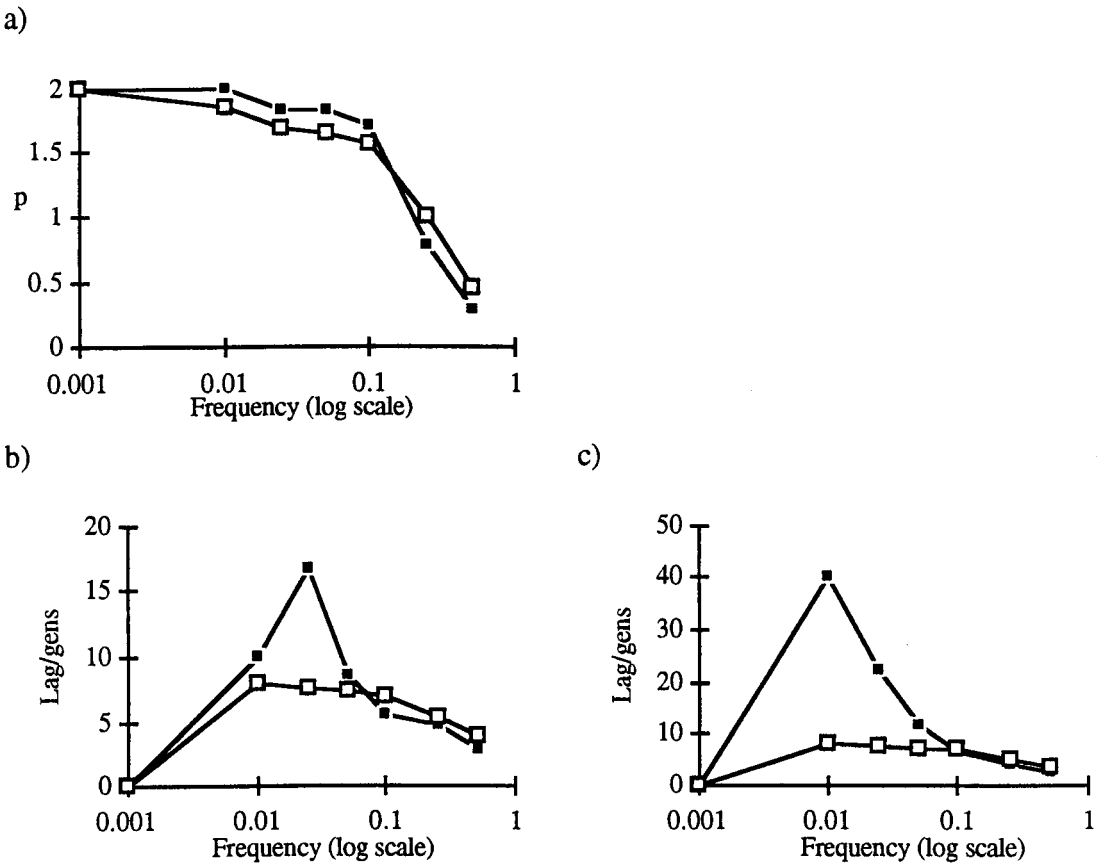

Figure 6-9: Population size 200, rank min. 0.5, fixed m-rates 0.02/bit (solid). 0.045/bit (empty) - a) amplitude vs. log frequency; b and c) Lande & Shannon's predicted (solid) and measured (empty) delay (generations) vs. log frequency, for m-rates 0.02/bit and 0.045/bit respectively.

6.4.1 Evolution in the Presence of Noise

In the previous section, I characterised the behaviour of an evolving system presented with a clean input signal. In this section, the filtering capabilities of an asexual evolving population (rank minimum 0.5) are demonstrated by providing an noisy input signal:

$$h_t' = a_0 + a_1 \sin(\alpha \times Gen) + a_2 \sin(\beta \times Gen) + a_3 \sin(\gamma \times Gen) \qquad (6.16)$$

such that $a_0 = 2.0 - (a_1 + a_2 + a_3)$; $a_2, a_3 < (a_1 = 1.0 - a_2 - a_3)$; and $\beta, \gamma >> \alpha$.

Note that the noise is *not* introduced on an individual basis (i.e. with each individual within a generation being evaluated according to a different optimum; it is well known that GAs can accommodate such noisy evaluation functions in terms of still being able to identify the optimum, although the time to convergence may be affected[64] (Miller & Goldberg, 1996)). Rather, *all* individuals within a generation are evaluated according to the same optimum. If the low pass filtering effect works, then the low frequency components of the input signal should be cleanly recovered.

The fixed m-rate is set close to the time averaged steady state evolved m-rate for the clean sinusoidal signal with frequency α. Results of a single typical run for 3 population sizes are given in figure 6-10, parameters set as in table 6-7.

Run	N	m-rate	Rank min.	a_1	a_2	a_3	α	β	γ
0	200	0.015	0.5	0.85	0.05	0.1	0.05	0.8	1.5
1	50	0.015	0.5	0.75	0.2	0.05	0.05	0.8	1.5
2	1000	0.015	0.5	0.5	0.2	0.3	0.05	0.8	1.5
3	200	0.015	0.5	0.75	0.20	0.05	0.05	0.1	1.5

Table 6-7: Parameters for an input signal given with added noise (equation (6.16)).

Note how the evolutionary filter does indeed appear to filter out the higher frequency noise, figures 6-10. Even for a low population size and relatively high noise levels, as in run 1, a significant clear up of the signal is accomplished, figure 6-10b. For larger population sizes, ever higher levels of noise can be removed (although acceptable limits are still to be identified). Where the noise has a frequency component of a similar order to the clean signal (e.g. as in run 3) the population tracks the skewed curve although the higher frequency noise components are still rejected, figure 6-10d. The popualtion itself has no direct knowledge of the clean signal (the evaluation function uses the *noisy* signal).

[64] Presumably as a result of a reduction in selection strength. Cf. landscape smoothing in chapter 4.

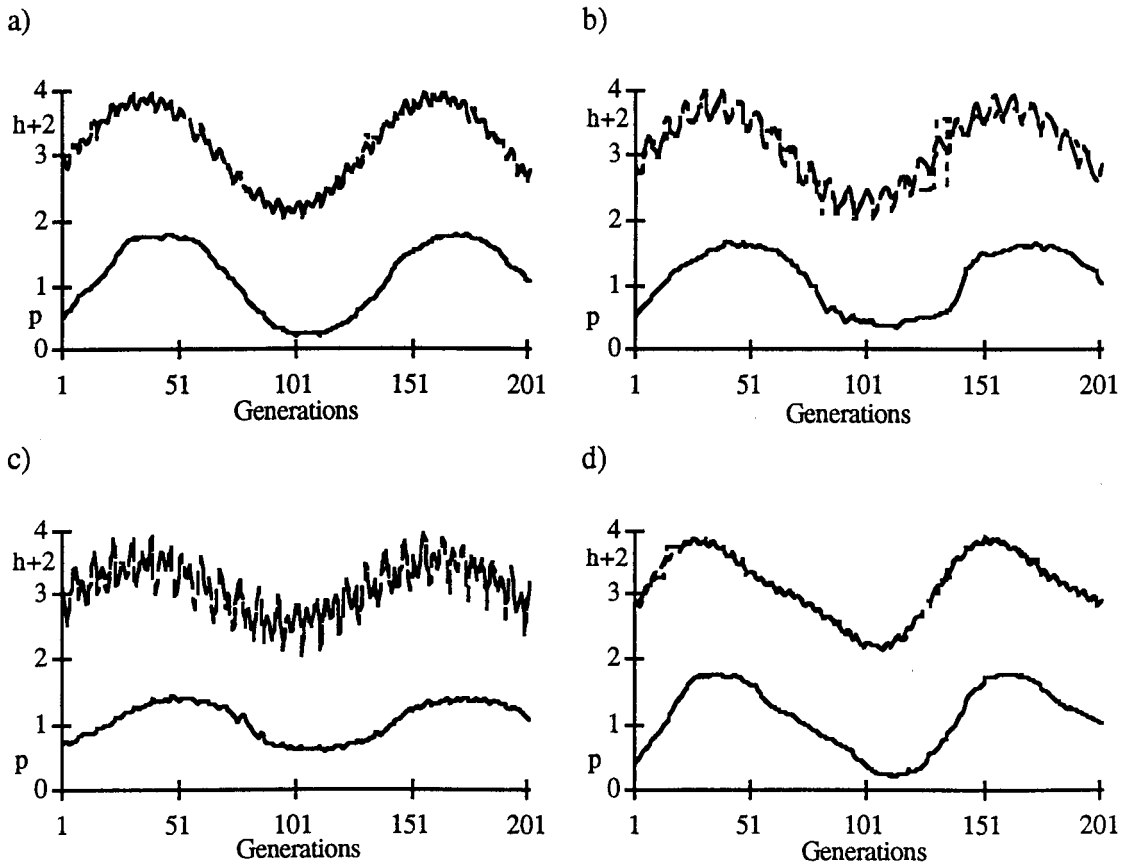

Figure 6-10: Target/population best curves (offset) and population mean expression for parameter sets given in table 6-3: a) run 0; b) run 1; c) run 2; d) run 3.

By comparing the individuals with the highest evaluation and the population mean expression, it is possible to extract a trace of the noise signal, figure 6-11. This is done by *delaying* the best trace by the amount of time the mean signal is delayed compared to the best; and scaling the mean trace to account for attenuation. Subtracting the scaled mean from the delayed best signal then give a reproduction of the *noise* signal. Note that we are only using *local* knowledge here (i.e. the best and mean expressions).

By subtracting the delayed *best of generation* curve from the population mean, *no* explicit reference is made to the optimum itself (although this value is used in evaluating each individual). What this means is that a similar method is applicable in situations using an *implicit* evaluation function, where information about the relative adaptedness of individuals is available, but knowledge of the evaluation function itself is not (as, for example, in some Alife ecological simulations).

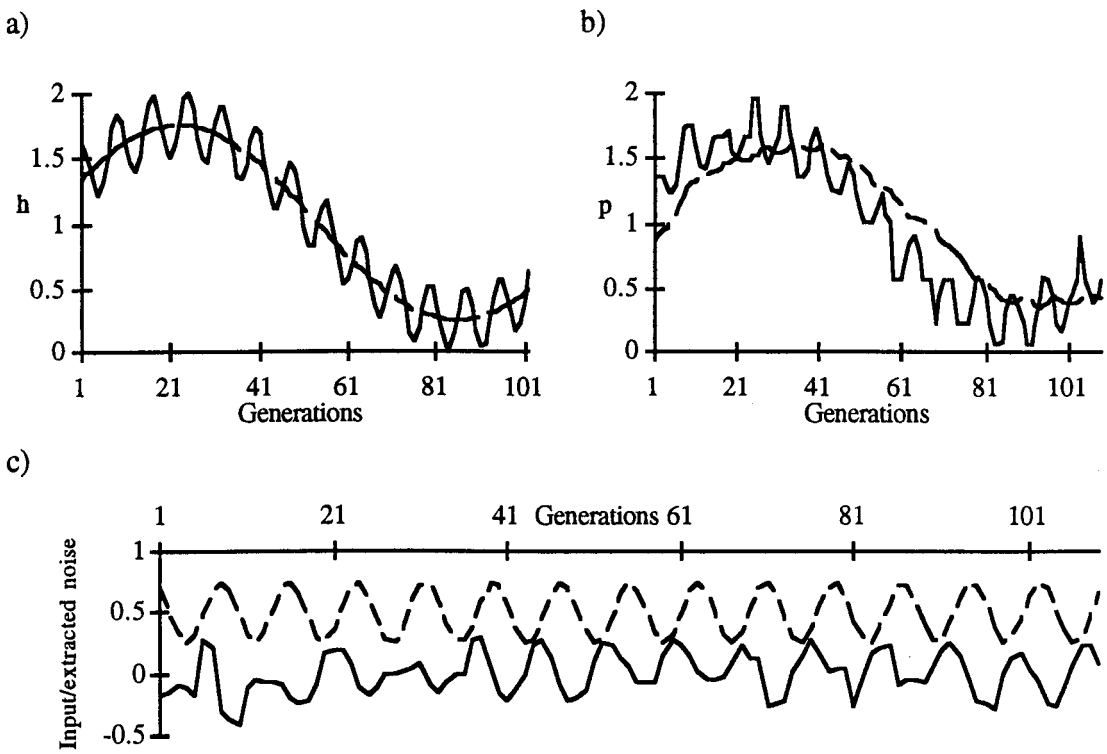

Figure 6-11: Extracting the noise signal: population size 100, mutation rate 0.015/bit; a_1 0.75, a_2 0.25, a_3 0, α 0.05, β 0.8. The traces are of a typical segment of a single, typical run; a) clean signal and its noisy counterpart (the original signal plus a high frequency component) that is used to evaluate individuals; b) the population best (noisy) and mean (filtered) expression for the same period; c) the actual added noise (offset) and the extracted noise component. The delay corresponds to the delay applied to the best signal.

6.4.2 Distinguishing Between Two Frequencies

This simple experiment shows how a suitably tuned filter may be used as part of system to distinguish between two frequencies (figure 6-12). One frequency falls within the pass band of the filter, the other outside it. The evolutionary filter output thus presents a large amplitude signal output for the low frequency input signal, and a rather lower level output for the higher frequency signal. Rectifying and smoothing the output and passing it through thresholding device will give a complete tone discrimination system.

It is possible that by separating the two frequencies futher, a single bit modifier that switches between a high and a low mutation rate may provide a population mean modifier that decodes the binary input signal directly.

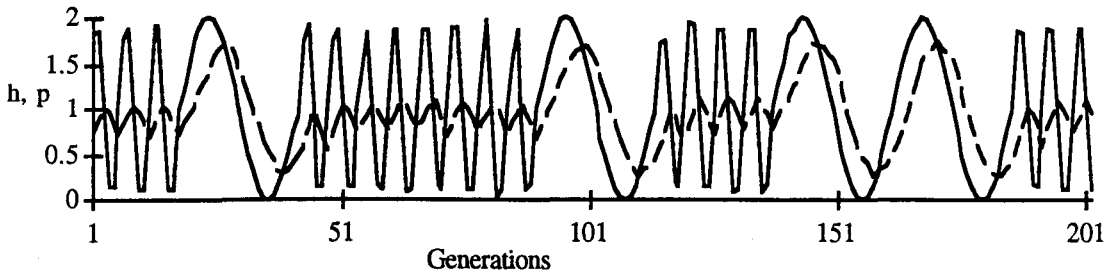

Figure 6-12: Population size 200, mutation rate 0.05/bit, input signal periods 6 and 24 generations; 'pulse width' 24 generations. Trace shows best (solid) and mean (dashed) expression. Clipping of higher frequency signal is a result of the short period.

6.5 Tracking an Input Signal with a Time-Varying Frequency

In this final short experiment, I shall employ an input signal with a time-varying frequency. This represents a more general counterpart to the fluctuating environment with periods of stasis demonstrated by Cobb. The particular test function used is given by a linear chirp test signal for the first part (increasing frequency over time, equation 6.17a), its mirror image for the second part (decreasing frequency, over time, equation 6.17b) to give the *chirprihc* signal:

$$h''_{Gen} = 1.0 + \sin\left(\alpha \times Gen^2\right) \qquad \text{for } Gen \leq \text{t} \qquad (6.17a)$$

$$h''_{Gen} = 1.0 + \sin\left(\alpha \times (2t - Gen)^2\right) \qquad \text{for } Gen > \text{t} \qquad (6.17b)$$

For this experiment, the self-adaptive regime, GA-M, utilises a modifier specific m-rate, set at 0.02/bit, and no other background rate (i.e. the evolved m-rate is the *only* one applied to the targeting bits). The fixed mutation rate has the rate set at 0.03/bit. In each case, the two point crossover rate is set to 0.6 per individual. The evaluation function parameters are set at $\alpha = 0.0005$, t = 500; the resulting frequency vs. generation number plot is shown in figure 6-13.

The rate at which the environmental signal changes during its increasing frequency phase is given by:

$$\frac{dh''_{Gen}}{dGen} = 2\alpha Gen.\cos\left(\alpha \times Gen^2\right) \qquad \text{for } Gen \leq \text{t} \tag{6.17c}$$

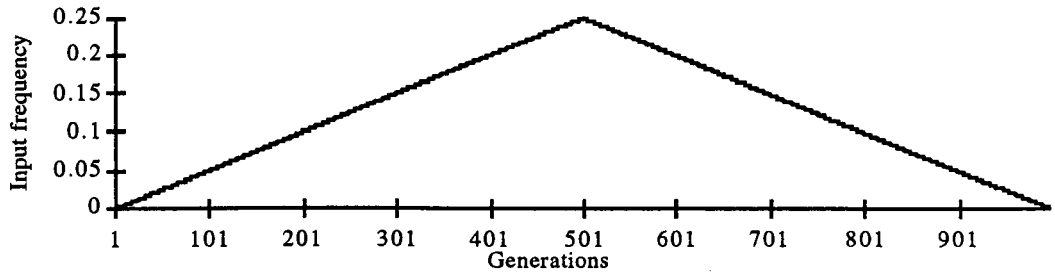

Figure 6-13: Plot of input signal frequency for the chirprihc function with t = *500 and* α = *0.0005.*

From figures 6-14a-d, note how the fixed m-rate regime, SGA appears to track the input signal more reliably than the self-adaptive strategy, GA-M. This is a consequence of the fixed, relatively low mutation rate, compared to the high rate levels achieved early on in the adaptive regime. Figure 6-15 shows the bitwise convergence of the fixed rate regime. Although high levels of convergence are not achieved, the convergence level is commensurate with those achieved by a similar mutation rate in a single frequency test environment (specifically, the convergence level for $\alpha = 0.1$ in the GA-M regime, plotted in figure 6-8b; the equilibrium evolved rate in this case is given in table 6-3 as 0.032). The relatively constant levels of bitwise convergence over increasing frequency suggest that the approximation used in the quantitative genetics theories described earlier (of constant additive genetic variance) hold in this particular experimental setting.

Turning now to the evolved m-rate, figure 6-14e shows that the rate does appear to track the increasing frequency of the input signal. This relationship is shown more clearly in figure 6-16b, where the evolved rates for a range of α (figure 6-16a) are shown. Just how well the modifier tracks the rate at which the input signal frequency increases is shown by a comparison of the gradients of the input frequency curves and the portions of the adaptive rate curves where the rate of increase is roughly linear, figure 6-17. One thing to note, however, is that the evolved rates for any particular frequency value under the chirprihc function are far higher than the equilibrium evolved rate for the corresponding fixed frequency evaluation function.

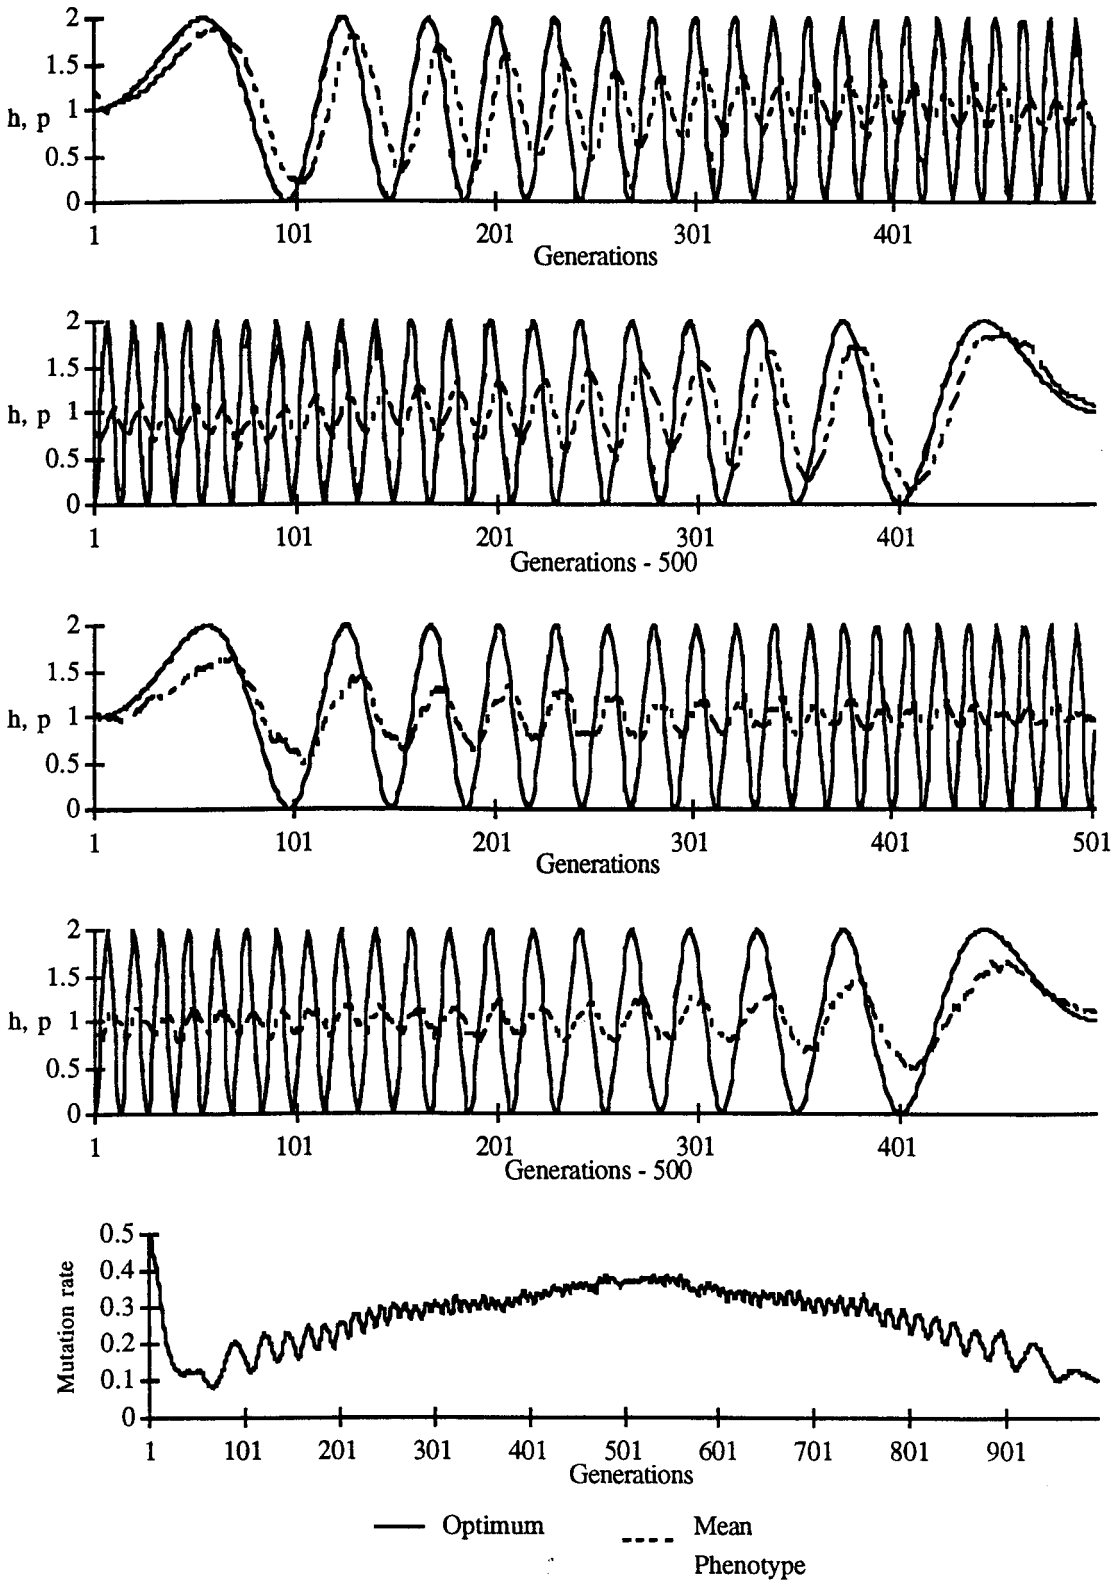

Figure 6-14: Single run traces using an input signal with a time varying frequency - a & b) regime SGA; c &d) regime GA-M; e) the mean evolved m-rate over 10 runs for GA-M.

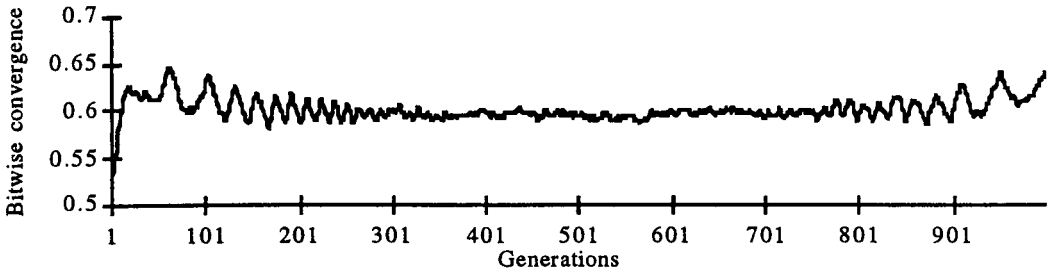

Figure 6-15: Bitwise convergence for SGA with fixed m-rate 0.03/bit, α = 0.0005.

a)

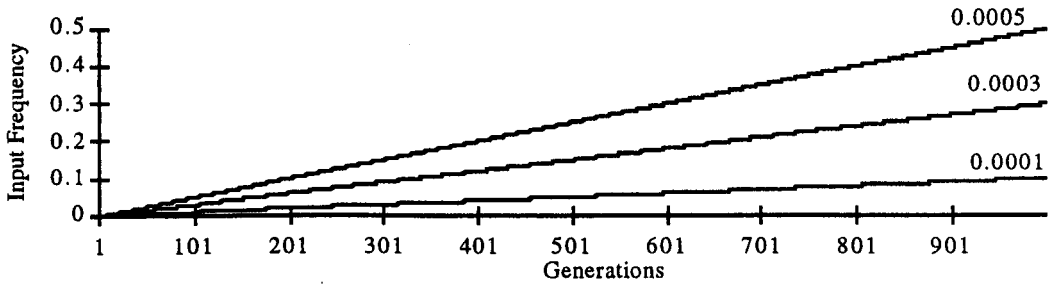

b)

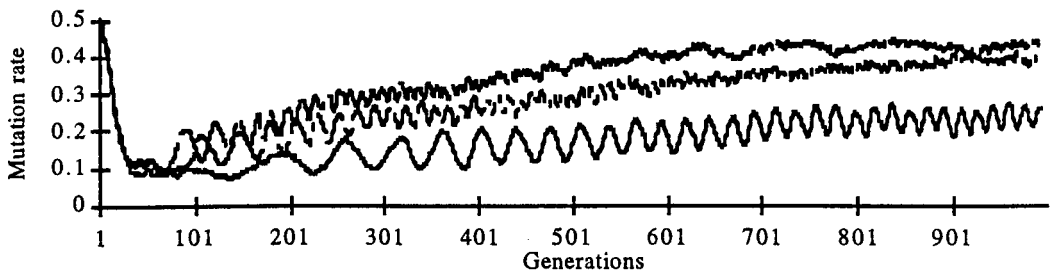

Figure 6-16: a) The input signal frequency over generations for α 0.0001, 0.0003, and 0.0005; b) the mean evolved m-rate taken over 10 runs for those three conditions.

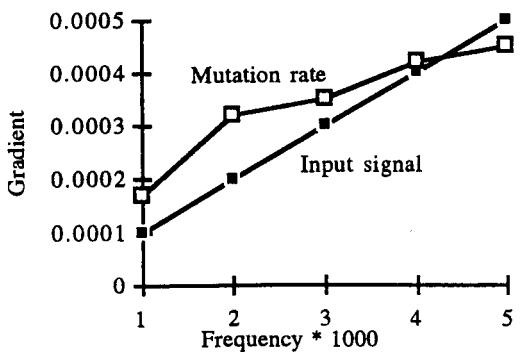

Figure 6-17: Least squares derived gradient of m-rate evolution calculated under chirprihc over generations 150-550 for α 0.0001...0.0005, step size 0.0001, which are also shown.

6.6 Implementing an Evolutionary Digital Filter- The Way Forward for ESP?

In this section, I will speculate on how a digital evolutionary FIR filter may be implemented as an *adaptive filter*. This will be illustrated by way of example (see figure 6-18). The genotype-phenotype map is taken to be the identity and so genotype and phenotype are used equivalently.

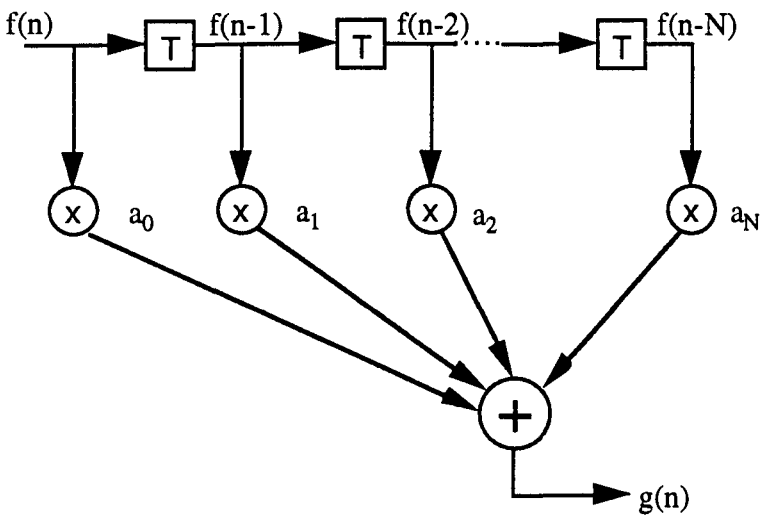

Figure 6-18: Implementing a non-recursive evolutionary digital filter.

By trying to describe an implementation of the digital evolutionary filter, it may be possible to find alternative interpretations of Charlesworth's analysis, or even provide a qualitative description of a model that when suitably analysed will provide an alternative to Charlesworth's approach. In all the implementations that follow, the filter output is given by the FIR filter summation (6.7a). The implementations differ in the way that the coefficients a_i are interpreted.

Implementation 1: the filter coefficients are as described by equation (6.9). This represents a literal interpretation of Charlesworth's analysis with a large population size. Where the ratio of genetic variance to selective variance is fixed, the coefficients will also be held constant. The filter becomes *adaptive* (i.e. the coefficients are allowed to vary over time) if the mutation rate - and hence V_g, then V -is allowed to change (for example, through self-adaptation).

Implementation 2: For a population of size N individuals, each of which has been the optimal individual exactly once within the previous N generations, the coefficients represent the selective value of the *previous* N optima, $f(n-i)$: $0 < i \leq N$, given the *current* optima, $f(n)$. Under rank selection, the coefficients a_i will be fixed according to:

$$a_0 = 0; \tag{6.18a}$$

$$a_i = \frac{1}{N}\left(\eta + 2(1-\eta)\frac{i-1}{N-1}\right) \tag{6.18b}$$

for $0 < i \leq N$, and rank minimum η. Whether this expression can be likened to (6.9) in any way is unclear.

Implementation 3: in this case, I relax the constraint that a_0 is set to zero. The coefficients represent the current proportion of individuals in the population, size $N = 2^L$, (binary genotypes, length L), with phenotype $f(i)$; note the time dependence of these coefficient values:

$$a_i(n) = \frac{a_i(n-1)s_i(n-1)}{\sum_{j=1}^{N} a_{j-1}(n-1)s_{j-1}(n-1)} \tag{6.19}$$

for selective value s, $0 < i \leq N$. Cf. equation (6.1c).

6.7 ESP as a Model of Evolution and an Analytic Tool

In this chapter I have introduced the notion of *Evolutionary Signal Processing*, which describes the implicit filtering properties of a population evolving in a fluctuating environment.

Two quantitative genetics analyses have been described - a continuous case and a discrete case. In the continuous time domain, the evolutionary transfer function describes a first order, low pass Butterworth filter. In the generational case, evolution appears to describe a bandpass FIR filter. A direct comparison of these two models, using the signal processing metaphor, suggests they are incompatible at both low environmental frequencies and in the

magnitude of the responce over all but very high frequencies. Experiments with a GA lend support to the continuous case analysis, but interpreting and applying the generational analysis remains problematic. However, the digital filter approach does suggest a novel way of implementing an evolutionary digital filter and developing this approach further may prove a fruiful line of enquiry.

A a test bed for studying evolution in time varying environments, several evaluation functions based on a sinusoidally varying optimum have been described (specifically by adding two other sinusoidal noise components, or by using a time dependent frequency). The use of self-adaptation has been shown to be a sensible way of tuning the evolutionary filter to a particular input frequency, by automatically locating an appropriate mutation rate. When faced with an input signal with a time dependent increasing frequency, the self-adaptive mutation rate tracks this increasing frequency. Several methodological questions have also been raised. For example, the decomposition of the evolutionary system's behaviour into initial and steady state responses and the unsuitability of evaluation measures, which are better replaced by reports on lag and attenuation.

If the ESP model is to be pursued, then it is of utmost importance that a quantitative relationship between the behaviour of the GA and theoretical models is obtained. The theoretically useful parameters are the variance of the selection function and the population. Selection functions for use in GAs have been analysed in detail by (Blickle & Thiele, 1995), and this represents a good starting point for further analysis with quantitative genetics applications in mind. In the quantitative genetics models, the *additive genetic variance* is the quantity used to describe the available *phenotypic* diversity. Whether a simple measure of the variance of the expressed phenotypic diversity can be related to the genetic variance, or some measure of diversity at the genotypic level is required is not clear. However, for sensible developmental maps, a linear relationship between these latter two quantities should hold if a sensible development function is used. Once a suitable measure of the population variance is identified, the secondary goal is to identify the relationship between the mutation rate and the population variance. If this can be achieved through a simple measure of the variation introduced by mutation, the expression for the mutation-selection balance offered by Anderson, equation (6.1m) may be applied.

Chapter 7 - Evolving Adaptive Individuals in Fluctuating Environments

"We will now proceed to the Third Experiment...This is a most interesting Experiment! It will need time I'm afraid: but that is a trifling disadvantage. Now observe..."

L Carroll, from "The Professor's Lecture" in Sylvie and Bruno Concluded.

7.1 Tracking Fluctuating Environments with Plasticity

In this chapter, I report on how an individual adaptive strategy that employs a form of trial and error learning affects the evolutionary dynamic of an evolving population. A regime which supports the inheritance of characters (IAC) acquired by means of this strategy is also considered. Following on from the use of a self-adaptive mutation rate described in the previous chapter, I implement a modifier gene that self-adaptively controls the amount of 'learning' pursued by an individual in systems with and without IAC. By coevolving the amount of plasticity supported by an individual and the extent to which the characters acquired through this plasticity are transmitted to the next generation, a demonstration of the theoretically predicted trade-off between individual and cultural learning is offered (Boyd & Richerson, 1983; Boyd & Richerson, 1988; Feldman et al., 1996). In addition, self-adaptive mutation and plasticity modifiers are allowed to coevolve, and the resulting dynamic compared to those resulting from the independent evolution of each.

In 7.2, I shall consider the extent to which plasticity is likely to affect the cut-off frequencies of the evolutionary filter under the Evolutionary Signal Processing model. In 7.3, I describe a GA model in which a modifier for adaptive plasticity is evolved under conditions of both simple inheritance, and the inheritance of acquired characteristics. Section 7.4 builds on this by coevolving *two* modifiers on each individual, one controlling the amount of learning pursued by each individual, the other the extent to which the acquired traits are transmitted to the next generation. Using a fluctuating environment with time dependent frequency, the trade-off between individual and cultural learning mechanisms described in chapter 3 is demonstrated within this two modifier model.

7.2 Plasticity in ESP

As mentioned in 6.2.1, (Anderson, 1995) presented a quantitative genetics analysis of learning under randomly fluctuating and constant velocity selective conditions. His intuition was that learning may be modeled, independently of its implementation, by an increase in the variance of selection:

$$V^l = V_s + V_l \tag{7.1}$$

In the continuous model of (Lande & Shannon, 1996), this will have the effect of *reducing* the cut-off frequency of the evolutionary low pass filter - the population will apparently track faster moving environments *less effectively*. This seems counter-intuitive, until one realises that the model predicts the mean *direct* expression of each individual; and in a learning model, the environment is tracked *not* by directly expressed individuals, but by adapted, *acquired* phenotypes (i.e. the developmental map is now one-many). The reduction in selection strength (i.e. increase in selection variance) suggested by Anderson thus describes the effect learning will have on the dynamics of the underlying gene pool. What he does not mention is the effect on the *phenotypes expressed through learning.*

There are two ways we might consider modeling the observed phenotypes acted on by selection - through a *reduction* in the selection variance (since genotypically distinct individuals are now capable of developing *similar* phenotypic expressions through learning); or by *increasing* the population variance to account for the *virtual* population evaluated during learning. I would suggest that the former approach may be used to model the evolution of the expressed mean in any learning regime; and the latter approach (of increasing the population variance) may be used to model the evolution of a population employing learning and the inheritance of acquired characteristics (where genotypes are identical to acquired phenotypes *of the previous generation*).

Each of these two models has a similar effect on the cut-off frequency of Lande & Shannon's low pass evolutionary filter: by *increasing* V, the cut-off frequency of the filter also increases and the *expressed phenotypes* are better able to track more dynamic environments.

If the trade off between individual learning, culture and simple mutation is to be supported, for example by setting the filter cut-off for each of these regimes accordingly, then we may

go further and suggest that the cutoff for each filter (i.e. V under Lande & Shannon's model) must be ordered in the following way:

$$V^m < V^{IAC} < V^l \tag{7.2}$$

Through sleight of hand, we can interpret the effect of learning on the evolution of expressed phenotypes in the following way. Through the introduction of learning, the reduction of V_s applied to expressed phenotypes corresponds to an increase in V, which I shall describe in terms of a multiplier, $l > 1$, applied to corresponding value, V, before learning (cf. the additive increase in selection variance of equation (7.1)). Substituting this into the equation for the gain squared of Lande's filter, (6.6b), we get:

$$\zeta^2(\omega) = \frac{1}{1+\frac{\omega^2}{V^2 l^2}} \tag{7.3a}$$

By incorporating l into the frequency term, this gives us the expression:

$$\zeta(\omega)^2 = \frac{1}{1+\frac{\omega_l^2}{V^2}} \tag{7.3b}$$

within which:

$$\omega_l = \frac{\omega}{l} \tag{7.3c}$$

and hence:

$$\omega_l < \omega \tag{7.3d}$$

since $l > 1$.

In other words, by introducing learning, then as far as the evaluated phenotypes are concerned the environment appears to be changing *more slowly* than for an identical input

signal applied to a population without the benefit of learning (cf. learning smoothing the landscape for a static evaluation function)[65].

In the discrete time domain, for learning with simple inheritance, as far as directly expressed genotypes are concerned the mid-point of the band pass filter is translated *down* the frequency spectrum and its pass band width *reduced*; but for the expressed phenotypes, the corresponding quantities are *increased* (i.e. not only can the population cope with faster moving environments, it can also cope with a wider *range* of environments).

In line with the framework of chapter 3, this ESP analysis allows us to raise a conjecture concerning the evolution of populations in fluctuating environments, specifically:

Conjecture 3: *in a non-stationary environment, the genetic and selection operators define the cutoff frequency of the adaptive filter with respect to the rate at which the valuation landscape changes. Using different development/within generation local search functions, the apparent rate at which the valuation surface changes may be modified. This is directly analogous to the apparent smoothing of the valuation landscape in the static case.*

7.3 Plasticity and Culture in Genetic Algorithms

This experiment follows on from the ones reported in the previous chapter which considered the evolution of mutation rates in a sinusoidally fluctuating environment. Identical fixed frequency (equation 6.15), noisy (equation 6.16) and temporally dependent frequency (equation 6.17) input signals were used.

7.3.1 Modeling Plasticity.

Plasticity, herein taken as a simple form of adaptive development or learning (i.e. within generation, local adaptive search), allows any particular genotype to realise one of several

---

[65] For the analysis at the genetic level, where the increase in selective variance *reduces V*, the environment is apparently 'speeded up' and the filter response effectively translated down the frequency spectrum.

possible phenotypes. Plasticity is modeled through a probabilistic mapping between each genotypic locus (bit position) and its phenotypic correlate over a fixed number of (learning) trials. In each learning trial, every 'phenotypic' locus is set to a value equivalent to the inherited allele with a probability (1-*rate)*, its complement otherwise, where *rate* is an evolved learning rate (1-*rate*) carried independently by each individual[66]. The evaluation value returned for each individual is then the lowest error score achieved over all that individual's trials. Note the 'innate' evaluation of the individual (i.e. the raw evaluation of the directly expressed individual) is only measured when no bits happen to be flipped in a particular trial. Whilst this scheme may seem to offer cost free plasticity, the probabilistic nature of the learning regime means there is an implicit cost between generations through the imperfect transmission of the parent's adaptive phenotype (discounting transmission errors arising from recombination). Cultural inheritance (through imagining the model at a solely phenotypic level, equating the inherited state to an initial phenotypic state rather than as a genetic message), or the inheritance of characters acquired through plasticity in a developmental model, is achieved by transmitting the genotype of the developed individual that gave rise to the lowest evaluation. IAC is applied to the fittest learned phenotype and is achieved by copying the phenotypic bits directly onto that individual's genome. In both of these learning regimes, the mutation rate is fixed, although in section 7.3 I will consider the simultaneous evolution of plasticity and mutation rates for the simple inheritance case.

7.3.2 To Compare SGA and Plastic GA With and Without the Inheritance of Acquired Characteristics (IAC).

In this experiment, I compare the performance of a plastic GA (using the scheme described above) with and without IAC (GA-P and GA-IAC respectively) with SGA from the previous

---

[66] Where *rate* is set to 0.5, this resembles the Hinton and Nowlan learning scheme for a population of individuals bearing all and only plastic alleles. That particular scheme may be modified to support tunable, switched learning capable of IAC by adding a learning rate modifier and using a coding scheme as follows: each locus is coded for by 2 bits. The first bit determines whether or not the allele is plastic, the second, targeting bit determines its initial state (i.e. its state at the beginning of a generation before learning). When the allele is plastic, learning is achieved according to flipping the bit away from its initial value with probability *rate*. IAC allows acquired targeting bit states to be transmitted through inheritance to the next generation.

experiment in the same sinusoidally fluctuating environment of equation 6.15. For each individual, 10 learning trials are allowed; the learning rate is an individually carried, evolved probability between 0.0 and 0.5[67] (the same rate for all bits, but an independent rate for each individual), represented as 5 contiguous binary coded bits. In both GA-P and GA-IAC regimes, there is a background mutation rate of 0.001 applied to *all* bits. Population sizes for these two learning regimes were set at 100, half the size of the population in the mutation only conditions. This follows Cobb's approach in which adaptive regimes were employed by populations smaller than simple, fixed regimes. The motivation there was that adaptive regimes should perform at least as well despite the smaller population size. A more pragmatic reason is that adaptive regimes are computationally more expensive, and some of this cost may be offset by means of reducing the population size[68].

Comparing generation 300, time averaged evaluation data, figure 7-1, both plastic regimes offer best of generation individuals that perform at least as well as for the SGA in all environments, and considerably better in the faster changing ones. In addition, mean online performance levels are higher in all environments for the plastic cases, GA-IAC outperforming GA-P. Note too that the mean online fitness for GA-P is similar across all but the slowest moving environment, whereas the mean fitness for GA-IAC deteriorates. The time averaged values taken over the last 150 generations (not shown) are very similar - each regime attains its equilibrium level of performance rather more quickly than the mutation only regimes.

How does this translate in terms of actual tracking ability (figure 7-2)? For the cases of plasticity, there are three useful measures that may be taken:- the best and mean phenotypic expression of the population, and the population mean direct evaluation of the genotype without plasticity. As for the mutation only cases (SGA, GA-M), the best individual of each generation tracks the target well, with the population capable of IAC closely following the environment as a whole both at the expressed and innate level. Specifically, closer inspection reveals that for GA-IAC in the slowly changing environments there is a lag of a single generation between environment and mean expressed value and two generations between

---

67 Setting the range between 0.0 and 1.0 allows each individual to choose the polarity of the binary alleles. For the sake of simple measurements, the rate was thus restricted to half the unit range.

68 An interesting control experiment would be to compare the performance of a simple population, size *N*, with an adaptive (learning) population of size *N/trials*.

environment and mean innate value; as the environmental rate increases, the population expressed mean tracks without any lag and the innate mean lag is reduced to a single generation. Similarly for GA-P, the population mean expressed phenotype lags the environment by a single generation in slow environments, not at all in faster ones, but the directly evaluated value lags considerably further (in the order of 10 generations for $\alpha = 0.1$), greater in fact than for the case of SGA. This is in accord with a quantitative genetics analysis of Anderson (Anderson, 1995) in which plasticity is modeled by a relaxation of the selection strength in both static and dynamic environments. In any particular run, the low rate at which *genetic* variation is introduced into the gene pool means the population may converge to a point from which learning alone is sufficient to track the moving target. In other cases, however, some degree of gene pool tracking is supported, although as has been mentioned, the lag is considerable. This problem may be addressed by increasing the background m-rate.

Where a noisy input signal is applied, tracking of the input signal occurs at the phenotypic level (i.e. the mean phenotypic expression as well as by the best phenotype), but the *direct* genotypic expression does represent a badly filtered version of the signal as in the no-plasticity case described in the previous chapter. Just how the directly expressed mean is affected by the mutation rate is left for further work.

a) Generation 300 offline fitness

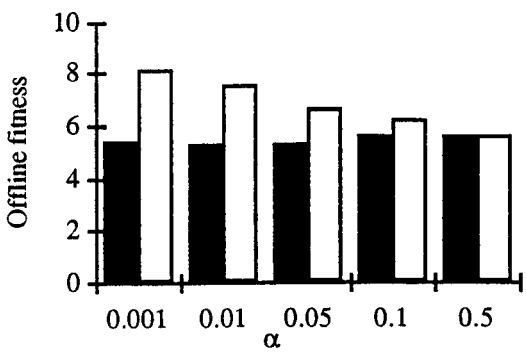

b) Generation 300 online fitness

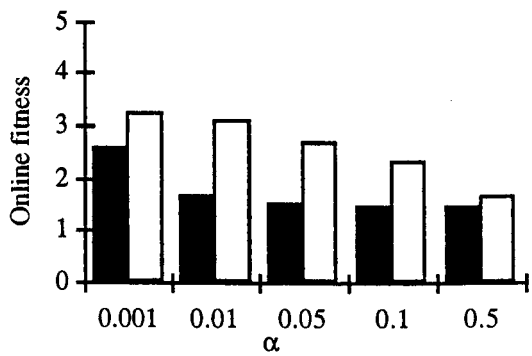

Figure 7-1: 'Fitness' measures against environmental rate, α, for GA-P (solid), GA-IAC (empty); a) Mean offline fitnesses at 300 generations; b) mean online fitness.

As with the evolved mutation rate, regime GA-M discussed in chapter 6, higher environmental frequencies support increased amounts of introduced variation, whether through an increased learning or mutation rate. Looking at the evolution of the learning rates, for regime GA-P, figure 7-3a, the rate assumes a relatively stable value for high and low environmental velocities, but for intermediate rates of change the rate itself varies as the

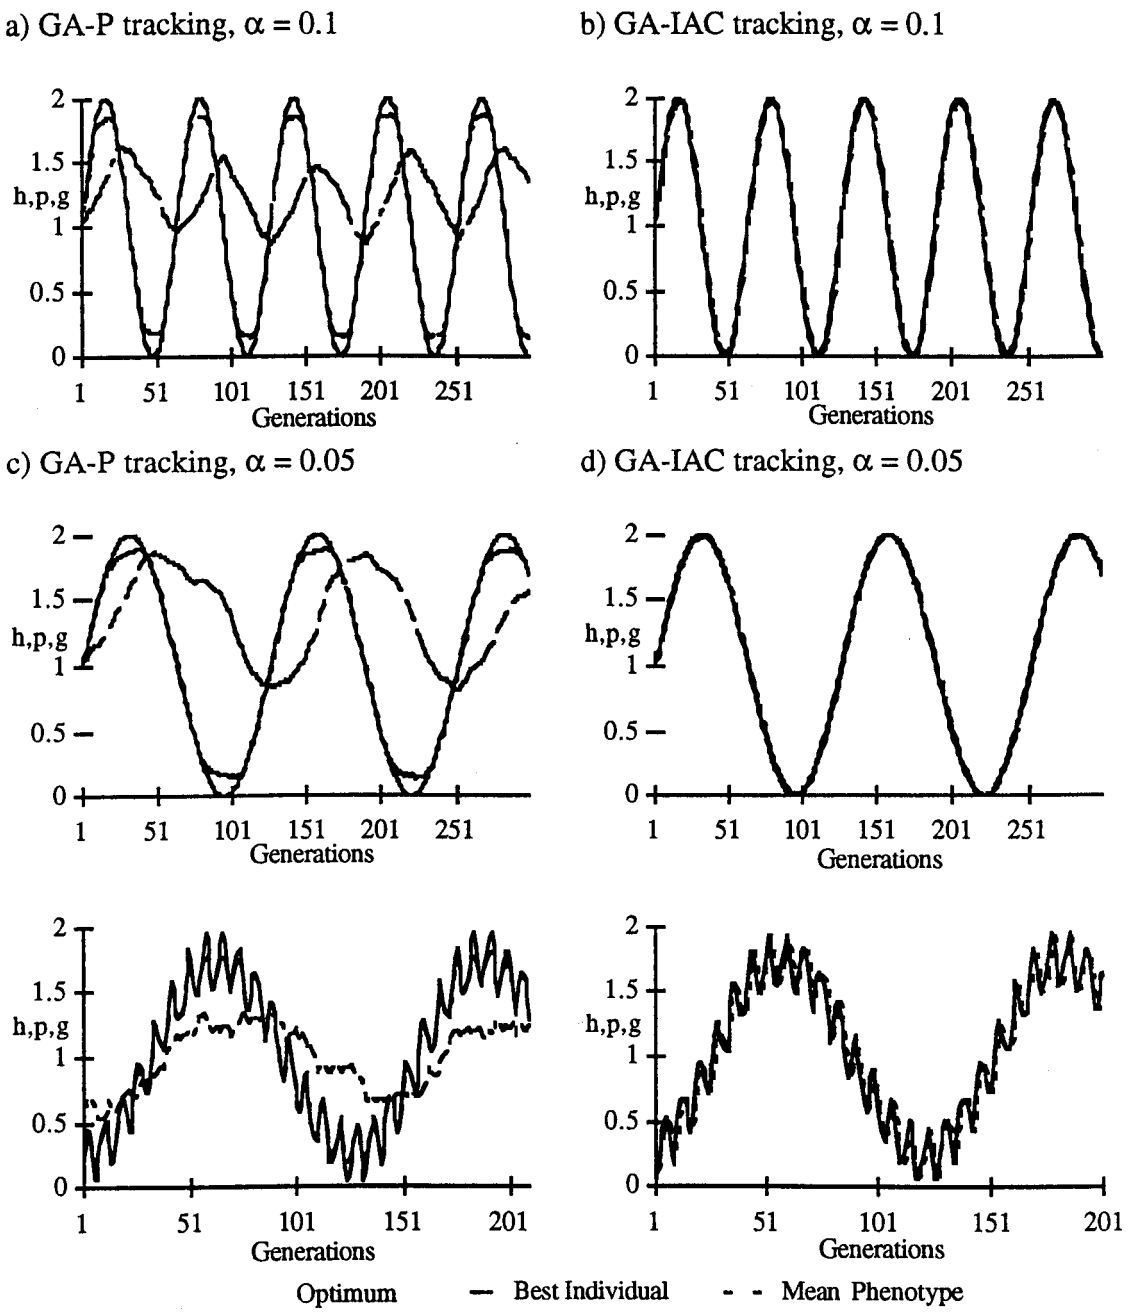

Figure 7-2: Tracking a moving environment during a single, typical run over 300 generations- a) and b), α = 0.1; c) and d), α = 0.05. Cases e and f show a single run trace over 200 generations for the noisy function, run 1, given in table 6-7. In all cases, the trajectory of the environmental value over time is very closely followed by the best of each generation. The solid line represents the environmental target, almost ideally tracked by the best of each current population. The dashed lines represent the mean phenotypic value of the population during the current generation - a, c, e) GA-P; b, d, f) GA-IAC.

population as a whole tries to track the environment. In a numerical simulation involving environments that move at a constant velocity, (Anderson, 1995) reports a similar increase in the amount of learning supported for increasing environmental rates. Where IAC is allowed, figure 7-3b, stable, lower learning rates are quickly adopted across the range of environments, although for intermediate α there is an initial period of increasingly damped oscillation. For $\alpha = 0.1$ at least, the rate evolved under GA-IAC is little affected by the strength of selection (not shown). Increasing the background m-rate for the simple inheritance case allows higher plasticity rates in the faster changing environments, figure 7-3c.

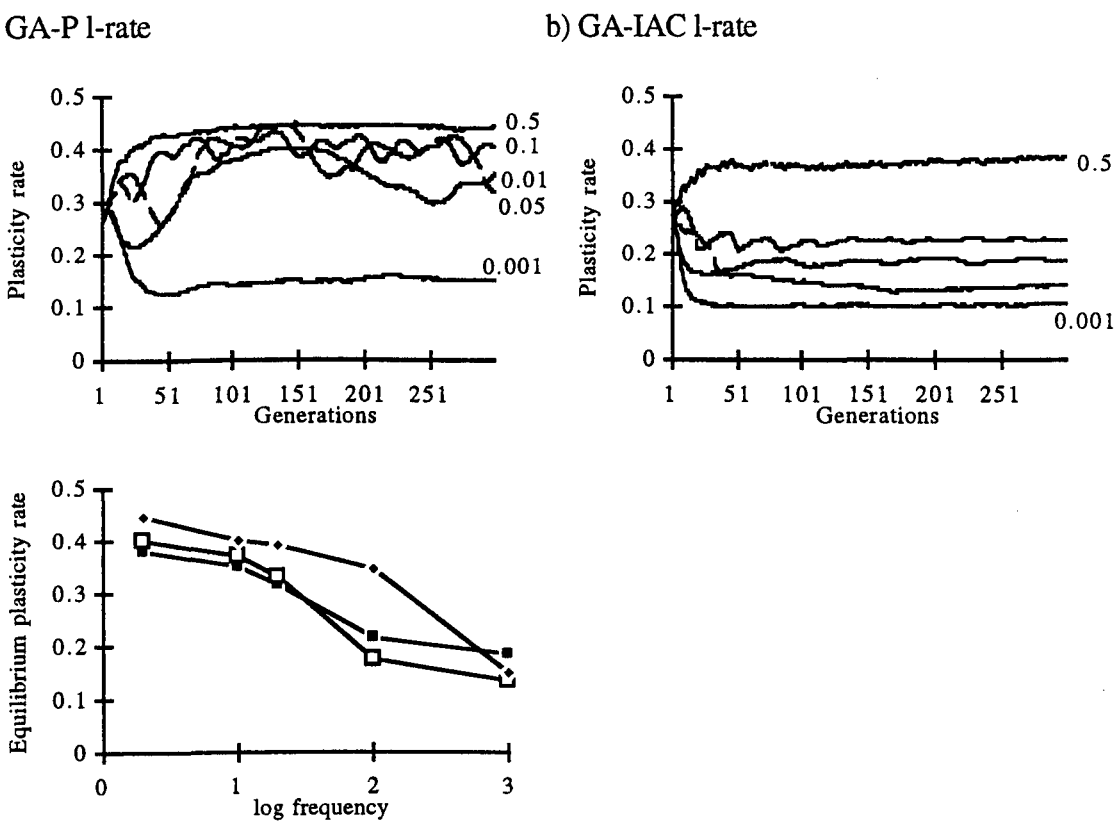

Figure 7-3: The evolution of plasticity rates over the range of environmental velocities for a) GA-P; b) GA-IAC. Generally, the higher the value of α, the greater the plasticity rate; c) plasticity rate vs. log frequency for regime GA-P over a range of background m-rates (0.001 (solid square), 0.02 (empty square) and 0.045 (solid diamond) /bit).

Since the developmental map is essentially simple (phenotypic states have a direct genotypic coding), GA-IAC may be thought of as an example of directed mutation as opposed to the weaker adaptive mutation of regime GA-M. That is, only currently *useful* 'mutations' are allowed to occur and GA-IAC is effectively allowed to evolve a 'useful mutation' rate, rather than a 'random mutation' rate. It thus makes sense to compare the evolved rates for these two regimes: in particular, GA-IAC supports a range of switching rates of comparable high order across all environments, whereas the evolved mutation rate of regime GA-M spans two orders of magnitude and only compares with the rates of GA-M in the fastest moving environment.

Considering the degree of convergence for the plastic populations over the range of environmental rate (figure 7-4), GA-P populations maintain a higher degree of convergence than do GA-IAC populations, as a result of the low rate at which genetic variation is introduced. For intermediate to large α, the IAC population displays a degree of convergence on a par with that of the evolved mutation rate case, regime GA-M. For low values of α, GA-P converges to a similar degree as GA-M. At higher fluctuation rates, the plastic algorithm maintains a consistently high level of convergence - the high degree of plasticity supplies the phenotypic variation necessary for effective tracking, and the low mutation rate restricts the amount of genetic variation. It falls to further work to make a direct comparison of the convergence of an SGA regime and a GA-P regime with similar fixed mutation rates across all environmental rates.

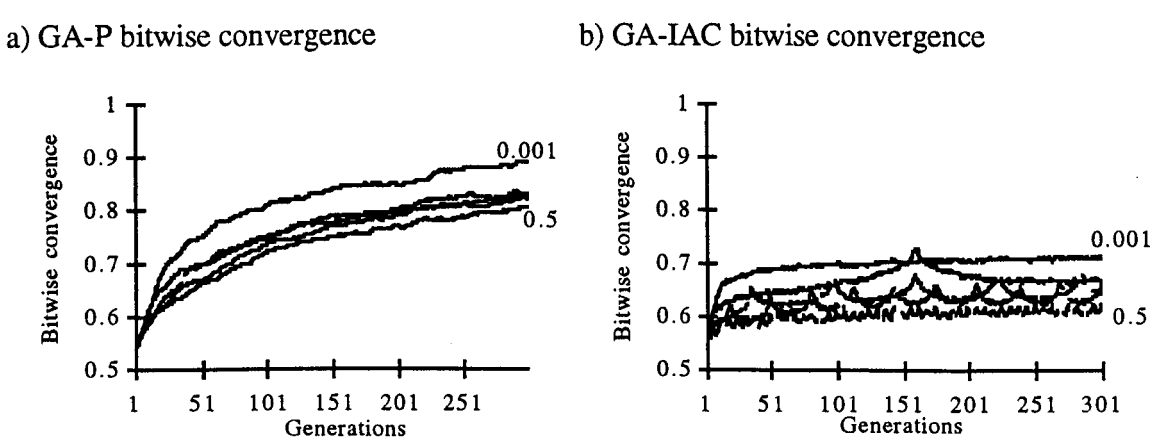

Figure 7-4: Bitwise convergence for a) GA-P; b) GA-IAC.

7.3.3 Plasticity in Environments with Time Dependent Input Frequency.

In figure 7-3, the evolved plasticity rates are shown for input signals of a fixed frequency. In this section, I shall use the chirprihc signal given in equation 6.17, to see whether the modifiers are capable of tracking the input signal frequency; note that for this experiment $\alpha = 0.001$. The frequency of the input signal over generations is given in figure 7-5a.

Any costs arising from learning with simple inheritance are likely to arise form the indeterminate nature of the learning style applied. As such, any learning cost is likely to be quite small. In figure 7-5b, the plot of learning rate over time for the learning and simple inheritance regime GA-P shows how for all but the slowest regimes, the learning rate quickly establishes itself at a high rate that then remains more or less constant. For the case of IAC, the learning rate is an increasing function of the signal frequency. This suggests that if IAC is supported, more and more 'learning' is required to ensure that appropriate states are acquired if inappropriate ones are inherited.

Figure 7-5 (overleaf): a) input frequency over time; b) evolved rate for regime GA-P; c) evolved rate for regime GA-IAC; d and e) tracking behaviour of GA-P and GA-IAC respectively, over 200 generations in frequency increasing phase.

—— Optimum — — Best Individual ---- Mean Phenotype

a)

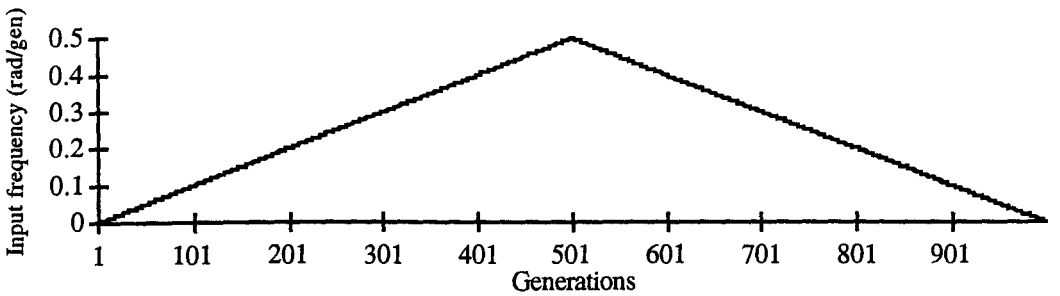
Input frequency (rad/gen)
0.5
0.4
0.3
0.2
0.1
0
1
101
201
301
401
501
601
701
801
901
Generations

b)

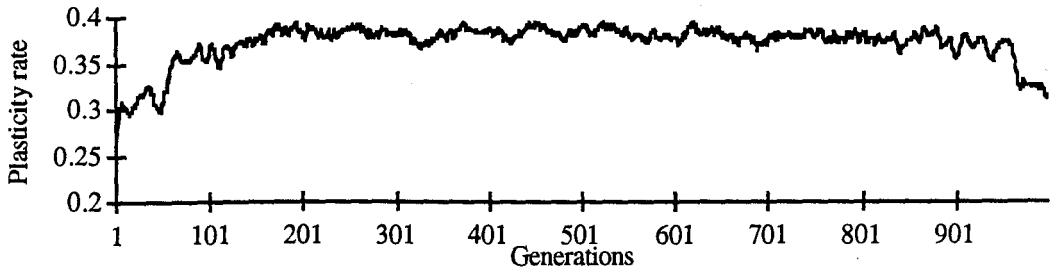
Plasticity rate
0.4
0.35
0.3
0.25
0.2
1
101
201
301
401
501
601
701
801
901
Generations

c)

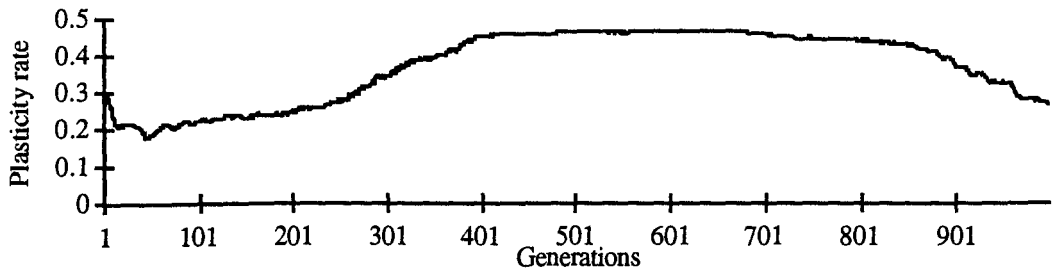
Plasticity rate
0.5
0.4
0.3
0.2
0.1
0
1
101
201
301
401
501
601
701
801
901
Generations

d)

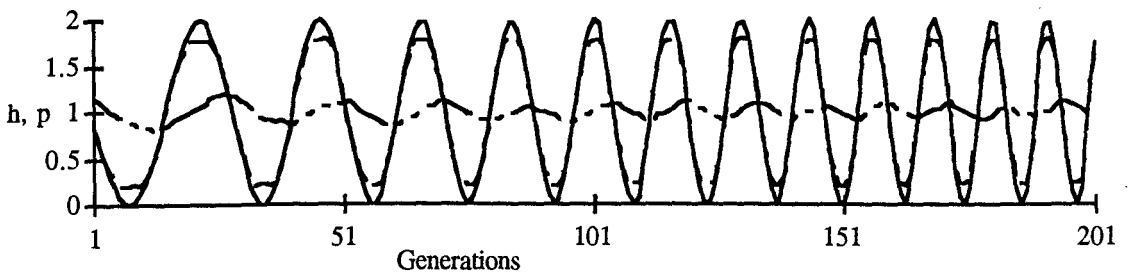
h, p
2
1.5
1
0.5
0
1
51
101
151
201
Generations

e)

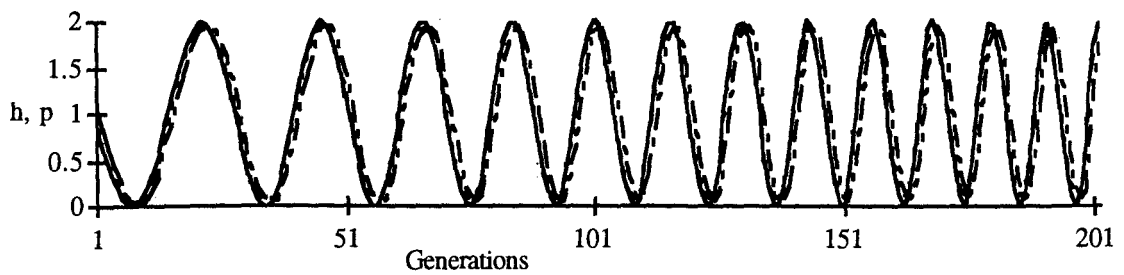
h, p
2
1.5
1
0.5
0
1
51
101
151
201
Generations

7.4 Coevolving Rates

With the addition of a second, evolvable rate parameter to the genome of each individual, it is possible to set the rate at which IAC occurs, in a regime denoted GA-C: the rate I gives the probability that each *acquired* bit will be transmitted, otherwise that individual's *inherited* state is transmitted. This experiment has a theoretical population genetics counterpart in a simple analysis by (Boyd & Richerson, 1988).

A second pairing of evolved rates that suggests itself is that of learning and mutation rates, GA-PM. For GA-P, it was found that the low background mutation rate occasionally resulted in the population failing to track the environment genetically and converging to a point from which plasticity had to provide all the necessary variation. By evolving the mutation rate, the population will be able to adapt the amount of variation introduced by that mechanism.

The results of running GA-PM and GA-C ten times over 300 generations on the standard range of fixed input frequencies are shown in figure 7-6; comparing this with figure 7-1, GA-C is seen to perform as well as GA-IAC in offline fitness terms, although underperforming slightly with regard to online performance. GA-PM outperforms GA-P for offline fitness, with again a slight underperformance by the online fitness measure. Again, the plasticity with simple inheritance strategy offers similar mean fitness across all but the slowest environments, whereas mean performance falls off for increasing α where IAC is allowed.

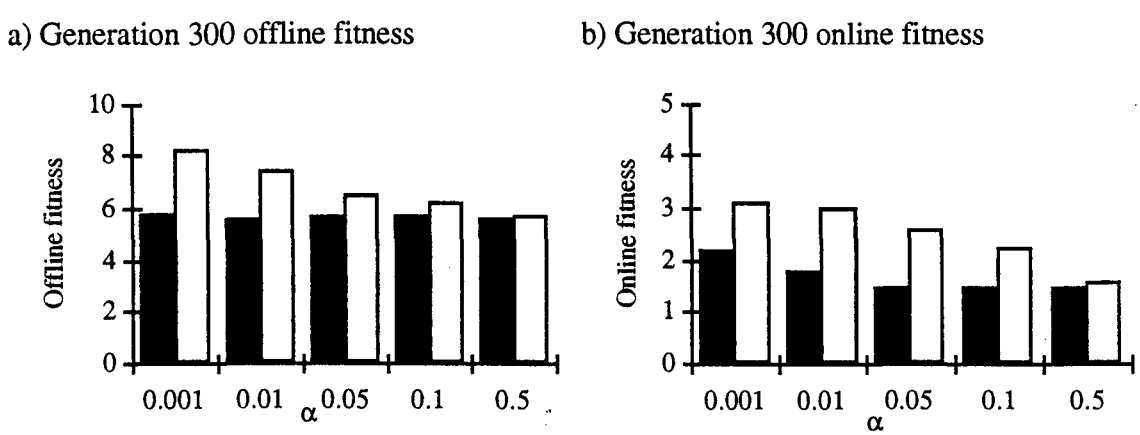

Figure 7-6: Fitness values against environmental rate, a, *for GA-PM (solid), GA-C (empty); a) Mean offline fitnesses at 300 generations; b) mean online fitness.*

7.4.1 Coevolved Mutation and Plasticity Rate Dynamics.

For GA-PM, figure 7-7b, in the slower fixed rate environments, the mutation rate starts off at a high level and then drifts randomly - the learning rate (figure 7-7a) is such that plasticity is capable of tracking the optimum under its own steam. As the environmental period decreases, the learning rate saturates at its maximal value and the mutation rate climbs to a high level, although this is not necessarily an adaptive measure. In the reported model, the selection pressure on the mutation rate modifier is so weak that the modifier is not driven down, but rather is free to change randomly. This is demonstrated even more vividly in figure 7-7c, where the input signal frequency changes over time. In this case, the mutation rate rarely settles (in fact, it doesn't seem to find a stable equilibrium value either within, or across, runs). What appears to be happening is that tracking is purely a function of plasticity, and the environmental frequency appears to have little effect on the dynamics of the modifiers once the plasticity rate achieves its high value. This suggests that the mean direct expression of the population fails to track the optimum at all, as is indeed the case (figure 7-7d).

7.4.2 Coevolved Plasticity and Inheritance Rate Dynamics.

Turning to a comparison of the coevolved rates for GA-C, figure 7-8, IAC is universally supported to a considerable extent in all but the fastest moving environments, with an increase in learning rate for increasing α. For high α, the culture rate begins to drop and the learning rate increases greatly - for faster moving environments, individual learning would appear to predominate over directly transmitted states.

This is clearly demonstrated in figure 7-8c, where the behaviour of the dual modifier rates are shown for the chirprihc function with $\alpha = 0.001$. After the initial period of settling, the plasticity rate climbs with increasing input frequency, but the inheritance rate falls. As input frequency decreases, the inheritance rate begins to climb and now the individual learning rate falls from its peak value.

a) GA-PM plasticity rate

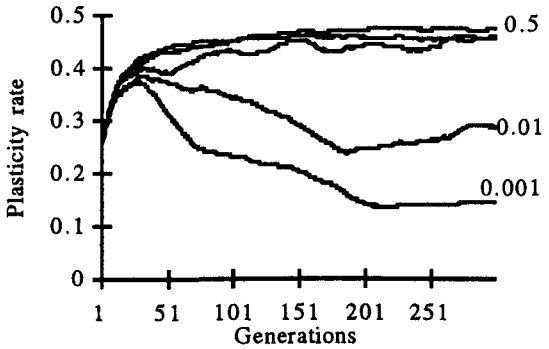

b) GA-PM m-rate

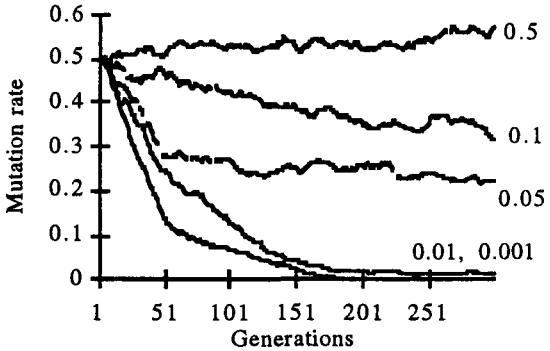

c) Plasticity and m-rates for GA-PM under chirprihc

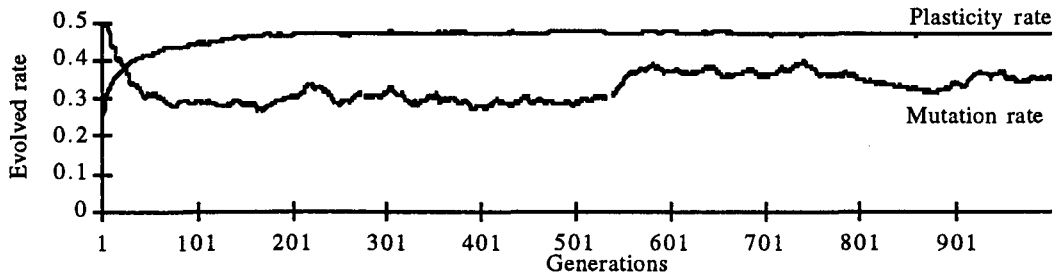

d) Typical within run tracking behaviour for GA-PM

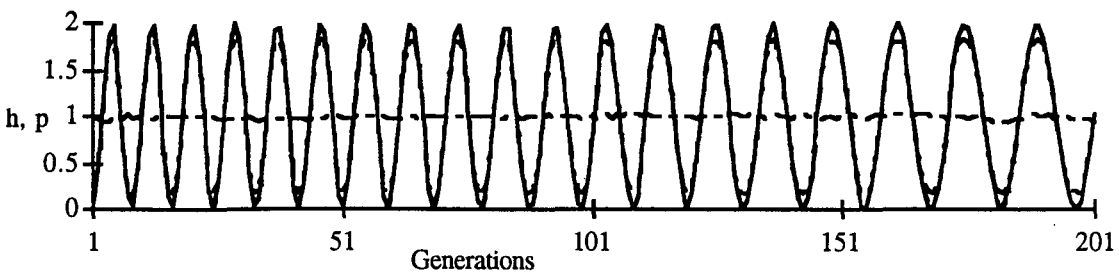

Figure 7-7: Evolved rates for GA-PM (α 0.001, 0.01, 0.05, 0.1, 0.5, 1, 5): a) GA-PM l-rate (increasing for increasing α); b) GA-PM m-rate (increasing for increasing α); c) plasticity and m-rates under chirprihc; d) trace over 200 generations during decreasing frequency phase of chirprihc - directly expressed mean fails to track at all.

a) GA-C plasticity rate

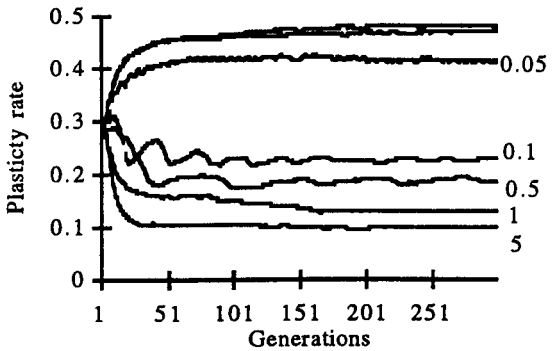

b) GA-C culture/inheritance rate

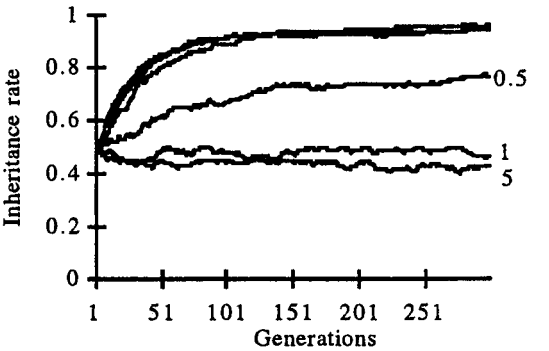

c) Plasticity and inheritance rates for GA-C

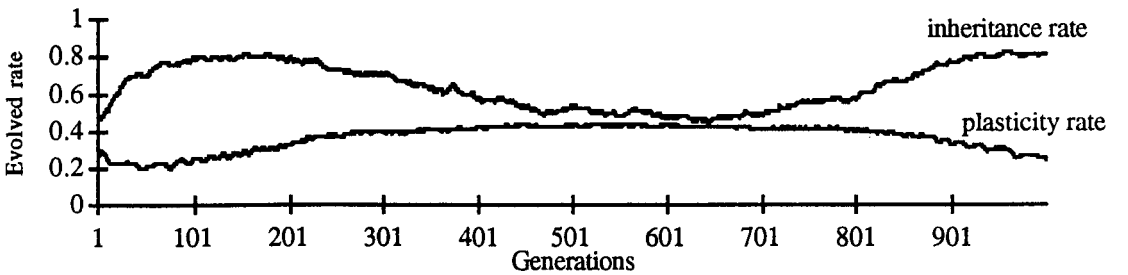

d) Tracking behaviour

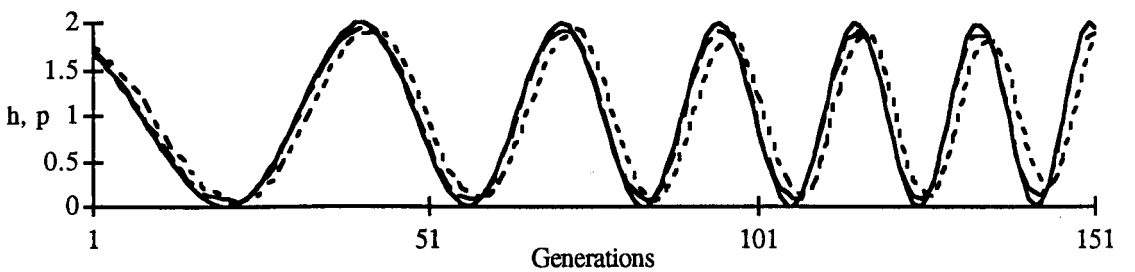

e) Tracking behaviour cont...

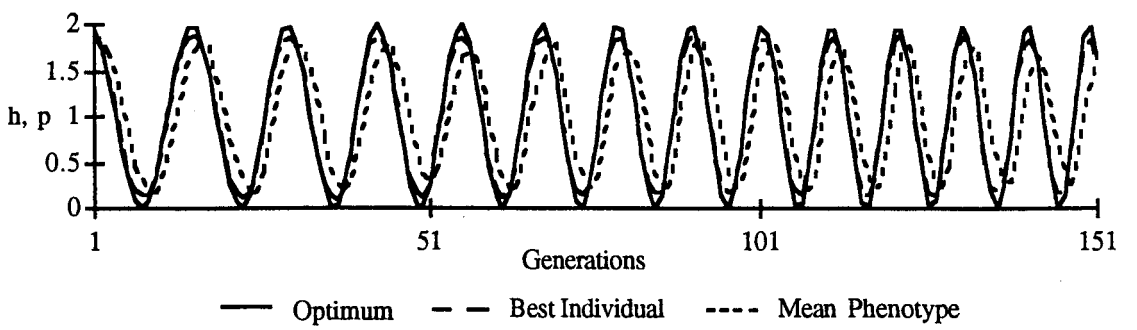

Figure 7-8: a) GA-C plasticity rate (increasing for increasing α); b) culture rate for GA-C (decreasing for increasing α); c) Inheritance and plasticity rate for GA-C under chirprihc; d) GA-C tracking over generations 50..200 during frequency increasing phase; e) 200..350.

Figure 7-9 demonstrates the trade off again, only this time in an environment that attains higher rates of change.

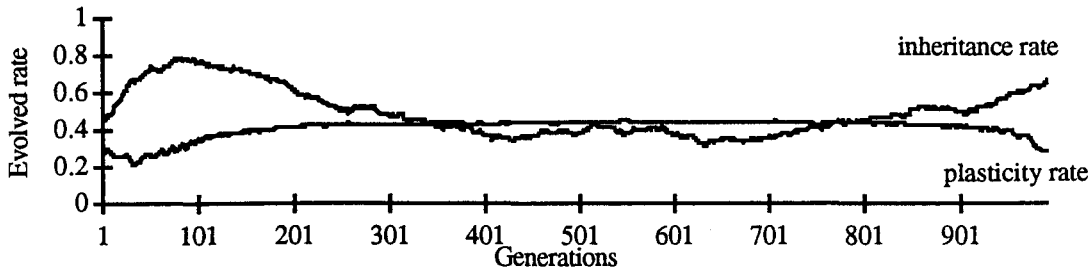

Figure 7-9: Inheritance and learning rates for GA-C under chirprihc, α 0.002.

7.5 Summary

Population members were allowed to evolve individual plasticity and inheritance rates. Generally, as the environmental period decreased, the evolved rate increased. Where plasticity with simple inheritance is introduced, the expressed population mean tracks the optimum more closely, although the underlying genes lag ever more. With IAC, lag is reduced to a single generation and the underlying gene pool is itself able to track the optimum very closely.

One problem identified during the experiments reported herein was the suitability of various 'traditional' measures of fitness for sensibly reporting the behaviour of populations evolving in fluctuating environments. Where only steady state, equilibrium behaviour is of interest, time averaged measures that take into account the early generations may skew the result. Where knowledge of the settling time behaviour is also required, then traditional on/offline measures may be suitable.

Treatments of learning and culture, in which a distinction is made between genotype and expressed phenotype, may be modeled by changes in selection strength and genetic variation. Consequences for the evolutionary filter view include the following effects on cut-off frequencies:

- with individual learning, there are essentially two filters in operation, one providing an output phenotypic mean expression, the other the directly expressed genotypic mean. The

large amounts of phenotypic variation provide a high cut-off frequency for the phenotypic mean and so mean phenotypic tracking of rapidly fluctuating environments is possible; the reduction in selection pressure on the genetic basis of selected phenotypes lowers the cut off frequency of the 'genetic filter' and so rapid fluctuations of the optimum are not mirrored by the genotypic mean.

- in a cultural or social learning system (e.g. (Boyd & Richerson, 1988; Feldman et al., 1996)), the distinction between phenotype and genotype is essentially removed, in that traits acquired through learning in one generation may be transmitted directly to the next, although there is a single generation lag between the mean acquired trait and the mean inherited trait. The net effect of high phenotypic variance of the 'virtual population' induced through learning, which under the inheritance of acquired characteristics translates to high genetic variance under strong selection (as individuals learn, then transmit, the same 'Good Tricks'), serves to set the single filter cut off frequency to an intermediate value.

What this means is that the simple evolutionary system has a low cut off frequency and only passes slowly changing signals. Cultural algorithms have an intermediate cut-off frequency and adaptive plasticity alone has a high cut off frequency, although all are low pass filters. (Boyd & Richerson, 1988; Feldman et al., 1996) offer theoretical results demonstrating these properties, and the genetic algorithm model discussed above provides some experimental support. The instability of systems where the environmental period is 1 or 2 generations (Feldman et al., 1996) may be related to the sampling theorem which states that the sampling frequency should be at least twice the cut off frequency (6.2.3.1); this deserves further study.

Comparing single strategies (SGA, GA-M, GA-P and GA-IAC) over environmental rate, GA-IAC appears to offer the best response, at least in terms of the fitness measures reported here. GA-P offers the next best strategy followed by the mutation rate regimes with a suitably set mutation rate. However, GA-P does appear to guarantee mean performance across α, whereas for all other strategies, mean fitness deteriorates with increasing environmental rate.

Additionally, IAC itself may be costly both in terms of local search, and also in terms of 'reverse engineering' an acquired phenotype in order to generate the genotype that represents it. Theoretical results from (Boyd & Richerson, 1988) suggest that in environments with an intermediate period (10s of generations) cultural inheritance/IAC offers an optimal trade-off

between individual learning and traits discovered through mutation. In rapidly fluctuating environments, where there is little correlation between adapted individuals of one generation and the next, an individual learning strategy is most appropriate. It is claimed that the coevolutionary dynamics of the learning and inheritance modifiers displayed in figures 7-8d and 7-9 (and also the learning rate dynamic for regime GA-IAC, figure 7-5c) demonstrate this trade off. For the claim to be valid, I assume that given high values of individual learning (the high learning rate) cultural inheritance is appropriately modeled by a high inheritance rate. Consequently, as the inheritance rate falls, the extent to which acquired traits are transmitted directly to the next generation falls; that is, an individual's evaluation is influenced more by traits acquired by the individual during its own lifetime.

The corresponding trade-off between plasticity and mutation for slower changing environments is revealed in figures 7-3a and b. However, where cost-free plasticity and mutation rates are co-evolved (regime GA-PM) the mutation rate dynamic suggests that the responsible modifier is placed under little, if any selection pressure, where there is a significant degree of plasticity. A more realistic scenario may be realised by introducing an explicit cost dependent on the amount of plasticity supported by the modifier.

Chapter 8 - Conclusions

Your Light May Go Out

A student of Tendai, a philosophical school of Buddhism, came to the Zen abode of Gasan as a pupil. When he was departing a few years later, Gasan warned him: 'Studying the truth speculatively is useful as a way of collecting preaching material. But remember that unless you meditate constantly your light of truth may go out.'

In 'Zen Flesh, Zen Bones', compiled by P Reps, Pelican 1971.

8.1 Contributions...and Criticism

The major contributions of this thesis result from taking a systems level view of the evolutionary process. By adopting such a position, the scope of the thesis has necessarily been a broad one. During the writing of the thesis, I was frequently forced into subtle changes of interpretation of the perceived system dynamics in order to accommodate conclusions drawn from other conceptual and behavioural considerations.

In chapter 1, I stated the intention that this thesis was offered as a contribution not only to the EC community, but also to those interested in a wider general evolutionary theory. Since evolutionary theory has been developed most pointedly by evolutionary biologists, I have tried throughout to relate ideas from that discipline to those of EC. To this end, a canonical evolutionary system based on a biological model was described in 2.3.

In my overview of genetic algorithms, and in contrast to the established order, I suggested that the behaviour of the genetic operators may be described in terms of the neighbourhoods they induce, rather than their effects on the schemata present within a population (2.4.1.1). This was further developed in 3.2.2.3. The notion of a schema was retained as a useful way of establishing the size of a given neighbourhood, and also as a way of describing the potential volume of the search space coverable by a population (2.5.2). Although not pursued in detail herein, the neighbourhood concept relates closely to the idea of molecular quasi-species. If EAs are to be applied widely in the discovery of robust designs, the quasi-species idea is likely to prove a powerful way of guaranteeing the tolerance limits/sensitivity of a

design nominally specified by some 'master sequence'. Various adaptive strategies for introducing variation were described in 2.6, and work on the behaviour of such systems in fluctuating environments was reviewed in 2.8.

The widespread terminological confusion also identified in chapter 2 within the evolution sciences in general, and the EA literature in particular, motivated the framework of chapter 3, which is offered as a way of describing the behaviour of a canonical evolutionary algorithm (3.2). In particular, landscape microstructure is viewed in terms of both operator and population dependencies. The distinction drawn between valuation, evaluation and selective value landscapes (3.3) is shown to have far reaching consequences for the way we interpret the evolutionary dynamics of a population evolving over these structures. Evolutionary search from an initially random population is considered as a 3 phase process (3.4): an initial period of random search, phase 1 (exploration), is followed by the second phase (exploitation) dynamics in which effective operators exploit certain regions of the valuation landscape. This culminates in a third phase of the search (equilibrium), specifically dynamic equilibrium of the population as a stable quasi-species (which under certain circumstances can be forced into a state of complete population convergence). This view of evolution, coupled with the work on landscape structure, lead to the conjecture that: *for a population with a given set of selection and genetic operator parameters, evolved populations are driven towards regions of the valuation landscape of a characteristic local ruggedness for the particular genetic and selective operators applied.* After a review of currently used measures of ruggedness, the runtime operator correlation coefficient was suggested as an appropriate measure of local ruggedness for testing this conjecture.

In chapter 4, I describe how such landscapes may be transformed by either transforming the underlying landscape structure (e.g. by self-adaptive mutation rates, 4.3); or the developmental map through either careful choice of the genotype-phenotype map, or by the introduction of within generation local search (4.4 and 4.5 respectively). For the cases of within generation local search, both adaptive and maladaptive, landscape transformation is shown to depend on the style of selection used (rank based or evaluation proportionate) and the evaluation transforming effect of the local search operator (4.5). Generally, one may think of the adaptive strategy *smoothing* the *valuation* landscape, and *maladaptive* search as *sharpening* it.

The analyses of landscape transformation are illustrated in chapter 5 with a set of experiments over the well known NK landscapes (5.3), and the novel *Tower of Babel* evolvable

evaluation function (5.4). By using GAs, it is possible to define a huge range of possible evaluation functions and collect vast amounts of data. The trick comes in finding test functions that ideally a) illustrate particular properties; b) are amenable to analysis; c) are computationally cheap; and measures that a) are interesting (e.g. by representing a system parameter), and b) understandable. By adding to the growing corpus of historical results using the same family of functions, the NK landscapes themselves become better understood and this in turn allows for greater understanding of the evolutionary systems in which they are applied.

Under the assumption that populations evolve towards regions of the valuation surface to which they are adapted, transformation of the valuation surface through local search should result in a population applying such measures settling in regions of the evaluation surface of a different ruggedness to a population of directly expressed individuals. This is the essence of my second conjecture (which represents a consequence of the first), specifically that: *for a population with a given set of selection and genetic operator parameters, the local ruggedness of the region of the valuation surface to which simple populations and populations capable of within generation local search converge is the same.* Again, the runtime operator correlation coefficient is suggested as an appropriate measure for investigating this claim further.

My particular contribution to the NK literature concerns the equilibrium position of a phase 3 population. This is shown to be within regions of predicted landscape ruggedness, as estimated by the spread of evaluations in a particular region of the search space: populations capable of individual learning identify more rugged regions of the *evaluation* surface (with higher optima) than simple populations; and populations susceptible to induced faults settle in comparatively smoother regions of the evaluation landscape.

In the experiments using the Tower of Babel evaluation function, I show how the cohesiveness of the steady state quasi-species plays a role in defining the evaluation surface to which the population is adapted. Evolvable evaluation functions provide one way into the domain of open-ended, incremental evolution, representing as they do a correlate to the notion of evolvable development functions.

In chapter 6, I introduce the notion of *Evolutionary Signal Processing*, which is based on the observation that mathematical analyses of populations evolving in sinusoidally fluctuating environments exhibit the same structure as expressions describing filters in the signal processing domain (6.2). In the steady state case, evolution is described by a low pass filter

whose cut off frequency is set according to the ratio of the additive genetic variance to the variance of selection. This observation is incorporated into the framework by a third conjecture, specifically that: *in a non-stationary environment, the genetic and selection operators define the cutoff frequency of the adaptive filter with respect to the rate at which the valuation landscape changes. Using different development/within generation local search functions, the apparent rate at which the valuation surface changes may be modified. This is directly analogous to the apparent smoothing of the valuation landscape in the static case.* A discrete case analysis, also obtained from the quantitative genetics literature, is shown to identify a conflicting response; this analysis also presents considerable problems for its sensible interpretation. Finally, an implementation of a digital evolutionary filter is suggested.

Experiments with a GA suggest that low pass filtering effect of evolution holds true for this mode of simulation (6.3). Although similar experiments have been run elsewhere previously, the GA community's fixation on reporting population best and mean 'fitnesses' almost to the exclusion of anything else meant this effect had gone unnoticed. The use of time averaged measures was also shown to be misleading for a periodic evaluation function, since the population response could clearly be separated into a period of settling (phases 1 & 2 of the search) followed by the steady state, dynamic equilibrium. Tuning of the evolutionary filter was achieved through the application of a self-adaptive mutation rate (6.3.2). This was shown to also respond to an input signal with a time varying frequency (6.5). Noise rejection capabilities of the evolutionary filter were demonstrated in 6.4.1.

Self-adaptive mutation rates are used in the ESP approach to tune the evolutionary filter in the following way:

- the self-adaptive population is presented with a sinusoidal input signal containing a single, fixed frequency component;
- the population is evolved until a steady state mutation rate is achieved;
- since this steady state mutation rate is itself likely to describe a sinusoid (of low amplitude), a time averaged mean value (taken over an integer number of periods of the rate) is obtained;
- it is suggested that this mean value describes a mutation rate that sets the evolutionary filter cut off frequency such that the input frequency lies within the pass band of the evolutionary filter.

Finally, in chapter 7, I showed how a GA using populations of plastic individuals supporting simple inheritance and the inheritance of acquired characteristics exhibited a trade-off identified in the biological literature between the relative adaptedness of these two strategies. In rapidly changing environments, individual learning is the evolutionarily stable strategy; in environments fluctuating with an intermediate period, a strategy that supports IAC is better adapted.

It is possible that closer parallels may be drawn between adaptive plasticity as a dynamic buffer against genetic and environmental perturbations, compared to the passive buffer of canalisation. This relies on decomposing the development function into a fixed maturational component (notwithstanding the potential for different fixed reaction norms) and a more flexible learning function. Coping with instability is likely to provide a significant set of constraints and preferences on designs realised through artificial evolution.

8.2 Recommendations

Reflecting on the GA work in sinusoidally fluctuating environments, it is possible to identify the strengths and weaknesses of the approach followed, and from this make the following recommendations as regards further work in such environments:

- the distinction between settling periods and the steady state response of the population suggests that measures of settling time and steady state behaviour should be clearly separated;
- evaluation measures are not necessarily the most informative; a better approach is to record the amplitude and phase response of the system in the steady state regime;
- all measures should be taken as a time average over an integer nuumber of periods;
- although phase information is lost, where noisy evaluation functions are used, a comparison of the Fourier coefficients of the input and output signals is likely to prove most informative. This should be related to the signal to noise ratio of the input signal.

8.3 Further Work

The *neighbourhood view of operator behaviour* suggests further work along two axes - a) the formal application of molecular quasi-species analyses of Eigen et al. to the behaviour of GAs; b) characterisation of the ruggedness of landscape in terms of ruggedness over population and operator neighbourhoods.

- Experimental validation of the resulting theory of quasi-species in GAs will require the development of appropriate measures of population dispersal/convergence, at least one of which should include a distance metric that is operator specific. My own reported results on bitwise convergence may have been clearer if a separate measure had been applied to 'targeting' and modifier regions of the genotypic bitstring.

- The development of appropriate metrics for describing 'operator ruggedness' are *essential* if the two conjectures made in the first half of the thesis are to be established. The presented results in terms of a local measure of the *range* of valuations represents a qualitative demonstration of the conjectures, but much remains to be done if a rigorous quantitative corroboration is to be forthcoming.

- It is suggested that the runtime operator correlation coefficient represents an appropriate measure of local ruggedness for testing the first two conjectures more thoroughly. This measure should also be capable of identifying the extent of canalisation of a trait, and as such may be useful for predicting the robustness of an evolved solution to a problem.

The *models of learning* that I applied were very simplistic. I suggest that future work should consider the design of local search techniques *with reference to the transformations that result.* Further work on the reuse of ideas from image processing in the design of local search operators that induce neighbourhoods akin to pixel masks with well known image filtering properties (4.7) is also recommended as potential area of future work.

None of the evaluation functions applied in this thesis present 'real world' problems. Further work should aim to demonstrate that evolution is capable of discovering solutions of a

specifiable sensitivity. It is suggested that initial test problems should be well-understood examples selected from the general class of parameterised optimisation problems.

Although much emphasis has been placed on the use of non-stationary evaluation functions, this thesis has addressed neither situations involving *within generation change*, where the state of the selective environment changes during the lifetime of an individual, nor steady state (as opposed to generational) GA execution models. Both these conditions should be addressed if a more complete understanding of evolution in fluctuating environments is to be obtained. In particular, the within generation change model may be used to a) study the evolution of learning subject to implicit costing through continual evaluation during an individual's lifetime; and b) the evolution of reaction norms: for example, by using a self-adaptive individual that codes for two phenotypes, these may be switched adaptively at the start of a generation according to a comparison made between a self-adaptive 'reaction norm' modifier and the initial state of the environment.

The development of the ideas of *Evolutionary Signal Processing* may also be pursued in several distinct ways, both theoretical and practical. In terms of general evolutionary theory, correlates in the other disciplines that employ evolutionary metaphors (e.g. evolutionary economics) should be identified. In terms of GAs in particular, the relationship between operator parameters and filter parameters must be identified to allow for the accurate tuning of the evolutionary filter. Whilst the evolution of mutation rates has been shown to be proportional to the rate at which the environment changes for a particular string length, the relationship between self-adaptive mutation rate and environmental rate has not been demonstrated with respect to *different length* individuals. Whether the relationship is robust as length increases is left as an open question. The behaviour of the self-adaptive mutation rate facing noisy evaluation functions should be investigated more thoroughly with respect to the signal to noise ratio of the input. In particular, the threshold levels at which the population tracks the noise, as opposed to signal, component should be identified.

If ESP is to provide an alternative to traditional techniques, issues of computational complexity and behavioural robustness of evolutionary filters must be addressed.

Bibliography

Ackley, D, & Littman, M. (1991) "Interactions Between Learning and Evolution." In CG Langton, C Taylor, JD Farmer & S Rasmussen (ed.) *Artificial Life II*, pp. 487-509. Addison Wesley. SFI Studies in the Sciences of Complexity, vol. X.

Altenberg, L, & Feldman, MW. (1987) "Selection, Generalized Transmission and the Evolution of Modifier Genes. I. The Reduction Principle." *Genetics,* 117:559-572.

Altenberg, L. (1995a) "Genome Growth and the Evolution of the Genotype-Phenotype Map." In W Banzhaf & FH Eeckman (ed.) *Evolution and Biocomputation: Computational Models of Evolution*, pp. 205-259. Springer Verlag.

Altenberg, L. (1995b) "The Schema Theorem and Price's Theorem." In LD Whitley & M Vose (ed.) *Foundations of Genetic Algorithms 3 (FOGA3)*, pp. 23-49. Morgan Kaufmann.

Altenberg, L. (1997a) "Fitness Distance Correlation Analysis: An Instructive Counter Example." In T Bäck (ed.) Proceedings of *Seventh International Conference on Genetic Algorithms (ICGA97).* Morgan Kaufmann.

Altenberg, L. (1997b) "NK Fitness Landscapes." In T Bäck, D Fogel & Z Michaelewicz (ed.) *Handbook of Evolutionary Computation*, Oxford University Press.

Anderson, RW. (1995) "Learning and Evolution: A Quantitative Genetics Approach." *Journal of Theoretical Biology,* 175:89-101.

Anderson, RW. (1997) "The Baldwin Effect." In T Bäck, D Fogel & Z Michaelewicz (ed.) *Handbook of Evolutionary Computation*, Oxford University Press.

Angeline, PJ, & Pollack, JB. (1994) "Coevolving High Level Representations." In CG Langton (ed.) Proceedings of *Artificial Life III*, pp. 55-71. Addison Wesley. SFI Studies in the Sciences of Complexity, vol. XVI.

Angeline, PJ. (1995) "Adaptive and Self-Adaptive Evolutionary Computations." In M Palaniswami, Y Attikouzel, R Marks, D Fogel & T Fukuda (ed.) *Computational Intelligence: A Dynamic Systems Perspective*, pp. 152-163. IEEE Press.

Arslan, T, & Horrocks, DH. (1995) "The Design of Analogue and Digital Filters Using Genetic Algorithms." In (ed.) *Digest of IEE 15th Saraga Colloquium on Analogue and Digital Filters and Filtering Systems*, pp. 7/1-7/5. IEE.

Asoh, H, & Mühlenbein, H. (1994) "On the Mean Convergence Time of Evolutionary Algorithms without Selection and Mutation." In Y Davidor, H-P Schwefel & R

Männer (ed.) *Parallel Problem Solving From Nature 3*, pp. 88-97. Springer Verlag. LNCS, vol. 866.

Aßelmeyer, T, Ebeling, W, & Rosé, H. (1996) "Smoothing Representation of Fitness Landscapes - The Genotype-Phenotype Map of Evolution." *Biosystems,* 39(1):63-76.

Atlan, H, & Koppel, M. (1990) "The Cellular Computer DNA: Program or Data." *Bulletin of Mathematical Biology,* 52(3):335-348.

Atmar, W. (1994) "Notes on the Simulation of Evolution." *IEEE Transactions on Neural Networks,* 5(1):130-147.

Bäck, T. (1992a) "The Interaction of Mutation Rate, Selection, and Self-Adaptation Within a Genetic Algorithm." In R Männer & B Manderick (ed.) Proceedings of *Parallel Problem Solving From Nature 2*, pp. 85-94. Elsevier.

Bäck, T. (1992b) "Self Adaptation in Genetic Algorithms." In FJ Varela & P Bourgine (ed.) *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pp. 263-271. MIT Press/Bradford Books.

Bäck, T. (1994) "Evolutionary Algorithms: Comparison of Approaches." In R Paton (ed.) *Computing with Biological Metaphors*, pp. 227-243. Chapman & Hall.

Bäck, T, & Hoffmeister, F. (1994) "Basic Aspects of Evolution Strategies." *Statistics and Computing,* 4:51-63.

Bäck, T, & Schwefel, H-P. (1996) "Evolutionary Computation: An Overview." In(ed.) Proceedings of *Third IEEE Conference on Evolutionary Computation*, pp. 20-29. IEEE Press.

Bäck, T, & Schütz, M. (1996) "Intelligent Mutation Rate Control in Canonical Genetic Algorithms." In W Ras & Z Michaelewicz (ed.) *Foundation of Intelligent Systems 9th International Symposium (ISMIS96)*, pp. 158-167. Springer Verlag.

Baher, H. (1990) *Analog & Digital Signal Processing*. John Wiley & Sons Ltd.

Baldwin, JM. (1896) "A New Factor in Evolution." *American Naturalist,* XXX:441-451, 536-553.

Baluja, S. (1994) Population-Based Incremental Learning: A Method for Integrating Genetic Search Based Function Optimisation and Competitive Learning. Computer Science Technical Report, CMU-CS-94-163, Carnegie Mellon University. June 1994.

Barnett, L. (1997) *Tangled Webs: Evolutionary Dynamics on Fitness Landscapes with Neutrality.* School of Cognitive Sciences, University of Sussex,

Bedau, MA, & Seymour, R. (1995) "Adaptation of Mutation Rates in a Simple Model of Evolution." *Complexity International,* 2

Behera, N, & Nanjundiah, V. (1996) "The Consequences of Phenotypic Plasticity in Cyclically Varying Environments: A Genetic Algorithm Study." *Journal of Theoretical Biology,* 178:135-144.

Belew, RK. (1989) "When Both Individuals and Populations Search: Adding Simple Learning to the Genetic Algorithm." In JD Schaffer (ed.) Proceedings of *Third International Conference on Genetic Algorithms (ICGA89)*, pp. 34-41. Morgan Kaufmann.

Belew, RK. (1990) "Evolution, Learning, and Culture: Computational Metaphors for Adaptive Algorithms." *Complex Systems,* 4:11-49.

Belew, RK, McInerney, J, & Schraudolph, NN. (1991) "Evolving Networks: Using the Genetic Algorithm with Connectionist Learning." In CG Langton, C Taylor, JD Farmer & S Rasmussen (ed.) *Artificial Life II*, pp. 511-547. Addison Wesley. SFI Studies in the Sciences of Complexity, vol. X.

Belew, RK, & Mitchell, M. (1996) *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison Wesley.

Bergman, A, Goldstein, DA, Holsinger, K, & Feldman, MW. (1995) "Population Structure, Fitness Surfaces, and Linkage in the Shifting Balance Process." *Genetical Research,* 66:85-92.

Blickle, T, & Thiele, L. (1995) A Comparison of Selection Schemes Used in Genetic Algorithms. TIK-Report, Nr.11, December 1995 v2, Swiss Federal Institute of Technology.

Bowler, PJ. (1973) "Bonnet and Buffon: Theories of Generation and the Problem of Species." *Journal of the History of Biology,* 6(2):259-281.

Boyd, R, & Richerson, PJ. (1983) "The Cultural Transmission of Acquired Variation: Effects on Genetic Fitness." *Journal of Theoretical Biology,* 100:567-596.

Boyd, R, & Richerson, PJ. (1985) *Culture and the Evolutionary Process*. University of Chicago Press.

Boyd, R, & Richerson, PJ. (1988) "An Evolutionary Model of Social Learning: the Effects of Spatial and Temporal Variation." In TR Zentall & BG Galef Jr (ed.) *Social Learning: Psychological and Biological Perspectives*, pp. 29-48. Lawrence Erlbaum Associates.

Brown, D, & Rothery, P. (1993) *Models in Biology: Mathematics, Statistics and Computing*. John Wiley & Sons.

Bulmer, MG. (1985) *The Mathematical Theory of Quantitative Genetics*. Oxford University Press.

Byerly, HC, & Michod, RE. (1991) "Fitness and Evolutionary Explanation." *Biology and Philosophy,* 6:1-22.

Charlesworth, B. (1993) "Directional Selection and the Evolution of Sexual Recombination." *Genetical Research,* 61:205-224.

Cobb, HG. (1990) An Investigation into the Use of Hypermutation as an Adaptive Operator in Genetic Algorithms Having Continuous, Time-Dependent Nonstationary Environments. Navy Center for Applied Research in AI. December 11, 1990.

Cobb, HG, & Grefenstette, JJ. (1993) "Genetic Algorithms for Tracking Changing Environments." In S Forrest (ed.) Proceedings of *Fifth International Conference on Genetic Algorithms (ICGA93)*, pp. 523-530. Morgan Kaufmann.

Collins, RJ, & Jefferson, DR. (1992) "The Evolution of Sexual Selection and Female Choice." In FJ Varela & P Bourgine (ed.) *Toward a Practise of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL92)*, MIT Press/Bradford Books.

Coyne, J, & Paton, R. (1994) *Genetic Algorithms and Directed Adaptation. Poster presented at AISB workshop on Evolutionary Computation.*

Csanyi, V. (1981) "General Theory of Evolution." *General Systems,* XXVI:73-94.

Csanyi, V, & Kampis, G. (1991) "Modelling Biological and Social Change: Dynamical Replicative Network Theory." In E Laszlo (ed.) *The New Evolutionary Paradigm*, pp. 77-92. Gordon & Breach Science Publishers.

Culberson, JC. (1994) "Mutation-Crossover Isomorphisms and the Construction of Discriminating Functions." *Evolutionary Computation,* 2(3):279-311.

Dasgupta, D, & McGregor, DR. (1992) "Nonstationary Function Optimisation using the Structured Genetic Algorithm." In R Männer & B Manderick (ed.) Proceedings of *Parallel Problem Solving From Nature 2*, pp. 145-154. Elsevier Science Publishers.

Dasgupta, D. (1994) Optimisation in Time Varying Environments Using Structured Genetic Algorithms. Technical Report, IKBS-17-93, Dept. of Computer Science, University of Strathclyde. December 1993 (updated 6/4/94).

Davis, L. (1989) "Adapting Operator Probabilities in Genetic Algorithms." In J Schaffer (ed.) Proceedings of *Third International Conference on Genetic Algorithms*, pp. 61-69. Morgan Kaufmann.

Davis, L. (1991) *Handbook of Genetic Algorithms.* van Nostrand Rheinhold.

de Jong, KA. (1993) "Genetic Algorithms are NOT Function Optimizers." In LD Whitley (ed.) *FOGA2*, pp. 5-17. Morgan Kauffman.

Eiben, AE, Raue, P-E, & Ruttkay, Z. (1994) "Genetic Algorithms with Multi-Parent Recombination." In Y Davidor, H-P Schwefel & R Manner (ed.) *Parallel Problem Solving from Nature 3*, pp. 78-87. Springer Verlag. vol. 866.

Eigen, M, McCaskill, J, & Schuster, P. (1988) "Molecular Quasi-Species." *Journal of Physical Chemistry,* 92:6881-68791.

Eigen, M, McCaskill, J, Schuster, P. (1989) "The Molecular Quasi-Species." In I Prigogine & SA Rice (ed.) *Advances in Chemical Physics LXXV*, pp. 149-263. John Wiley & Sons.

Faber, M, & Proops, JLR. (1991) "Evolution in Biology, Physics and Economics: A Conceptual Analysis." In PP Saviotti & JS Metcalfe (ed.) *Evolutionary Theories of Economic and Technological Change: Present Status and Future Prospects*, pp. 58-87. Harwood Academic Publishers.

Feldman, MW, & Laland, KN. (1996) "Gene-Culture Coevolutionary Theory." *Trends in Ecology and Evolution,* 11(11):453-457.

Feldman, MW, Aoki, K, & Kumm, J. (1996) Individual versus Social Learning: Evolutionary Analysis in a Fluctuating Environment. Working Paper, 96-05-30, Santa Fe Institute.

Fogarty, T. (1989) "Varying the Probability of Mutation in the genetic Algorithm." In J Schaffer (ed.) Proceedings of *Third International Conference on Genetic Algorithms*, pp. 104-109. Morgan Kaufmann.

Fogel, D. (1994) "Evolutionary Programming: An Introduction and Some Current Directions." *Statistics and Computing,* 4:113-129.

Fogel, LJ, Owens, AJ, & Walsh, MJ. (1966) *Artificial Intelligence through Simulated Evolution.* John Wiley & Sons.

Fontanari, JF, & Meir, R. (1990) "The Effect of Learning on the Evolution of Asexual Populations." *Complex Systems,* 4:401-414.

French, RM, & Messinger, A. (1994) "Genes, Phenes and the Baldwin Effect: Learning and Evolution in a Simulated Population." In RA Brooks & P Maes (ed.) Proceedings of *Artificial Life IV.* MIT Press.

Gavrilets, S. (1997) "Evolution and Speciation on Holey Adaptive Landscapes." *Trends in Ecology and Evolution,* 12(8):307-312.

Gillespie, JH. (1981) "Mutation Modification in a Random Environment." 35(3):468-476.

Gitchoff, P, & Wagner, G. (1996) "Recombination Induced HyperGraphs: A New Approach to Mutation-Recombination Isomorphism." *Complexity,* 2:37-43.

Goldberg, DE, & Segrest, P. (1987) "Finite Markov Chain Analysis of Genetic Algorithms." In JJ Grefenstette (ed.) Proceedings of *Genetic Algorithms and their Applications:*

Proceedings of the Second International Conference on Genetic Algorithms (ICGA87). Lawrence Erlbaum Associates.

Goldberg, DE, & Smith, RE. (1987) "Nonstationary Function Optimization Using Genetic Algorithms with Dominance and Diploidy." In JJ Grefenstette (ed.) *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms*, pp. 59-68. Lawrence Erlbaum Associates.

Goldberg, DE. (1989) *Genetic Algorithms in Search, Optimization & Machine Learning.* Addison Wesley.

Gomulkiewicz, R, & Kirkpatrick, M. (1992) "Quantitative Genetics and the Evolution of Reaction Norms." *Evolution,* 46(2):390-411.

Grefenstette, JJ. (1992) "Genetic Algorithms for Changing Environments." In R Männer & B Manderick (ed.) Proceedings of *Parallel Problem Solving from Nature*, pp. 137-144.

Happel, R, & Stadler, PF. (1995) Canonical Approximation of Fitness Landscapes. Working Paper, 95-07-068, Sante Fe Institute.

Harik, GR, Lobo, FG, & Goldberg, DE. (1997) The Compact Genetic Algorithm.

Hart, WE. (1994) Adaptive Global Optimization with Local Search. PhD, Department of Computer Science and Engineering, University of California, San Diego.

Hart, WE, Kammeyer, TE, & Belew, RK. (1994) "The Role of Development in Genetic Algorithms." In LD Whitley & MD Vose (ed.) *Foundations of Genetic Algorithms 3*, pp. 315-332. Morgan Kaufmann.

Hart, WE. (1995) Evolutionary Pattern Search Algorithms. SAND95-2293, Sandia National Laboratories.

Harvey, I. (1992) "Species Adaptation Genetic Algorithms: The Basis for a Continuing SAGA." In FJ Varela & P Bourgine (ed.) *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pp. 346-354. MIT Press/Bradford Books.

Harvey, I. (1993) "The Puzzle of the Persistent Question Marks: A Case Study of Genetic Drift." In S Forrest (ed.) Proceedings of *Fifth International Conference on Genetic Algorithms (ICGA93)*, pp. 15-22. Morgan Kaufmann.

Harvey, I, & Thompson, A. (1996) "Through the Labyrinth Evolution Finds a Way: A Silicon Ridge." In T Higuchi, M Iwata & W Liu (ed.) *Evolvable Systems: From Biology to Hardware*, pp. 406-422. Springer Verlag. LNCS, vol. 1259.

Hinterding, R, Michaelewicz, Z, & Eiben, AE. (1997) "Adaptation in Evolutionary Computation: A Survey." In(ed.) Proceedings of *Fourth IEEE International Conference on Evolutionary Computation.* IEEE Press.

Hinton, GE, & Nowlan, SJ. (1987) "How Learning Can Guide Evolution." *Complex Systems,* 1:497-502.

Hirst, AJ. (1996a) "Evolving Adaptive Computer Systems." In T Furuhashi (ed.) Proceedings of *First Online Workshop on Soft Computing (WSC1)*, pp. 257-262. Nagoya University.

Hirst, AJ. (1996b) "Evolving Adaptive Systems - On the Role of Learning and Fault Tolerance in Evolution." In T Furuhashi (ed.) Proceedings of *Second Online Workshop on Evolutionary Computing (WEC2)*, pp. 57-60. Nagoya University.

Hirst, AJ. (1996c) "Notes on the Evolution of Adaptive Hardware." In IC Parmee (ed.) Proceedings of *Adaptive Computing in Engineering Design and Control (ACEDC'96)*, pp. 212-219. University of Plymouth.

Hirst, AJ. (1996d) "Search Space Neighbourhoods as an Illustrative Device." In T Furuhashi (ed.) Proceedings of *First Online Workshop on Soft Computing (WSC1)*, pp. 49-54. Nagoya University.

Hirst, AJ. (1997a) "Evolutionary Signal Processing: A Preliminary Study." In P Husbands & I Harvey (ed.) Proceedings of *Fourth European Conference on Artificial Life (ECAL97)*, pp. 425-431. MIT Press/Bradford Books.

Hirst, AJ. (1997b) "On The Structure and Transformation of Landscapes." In D Corne & J Shapiro (ed.) Proceedings of *AISB97 Workshop on Evolutionary Computation.* Springer Verlag.

Hirst, AJ. (1997c) "Plasticity and Culture in Cyclically Fluctuating Environments." In D Corne & J Shapiro (ed.) Proceedings of *AISB97 Workshop on Evolutionary Computation.* Springer Verlag.

Ho, M-W, & Saunders, PT, ed. (1984) *Beyond Neo-Darwinism: An introduction to the new evolutionary paradigm.* Academic Press Inc.

Holland, JH. (1962) "Outline for a Logical Theory of Adaptive Systems." *Journal of the Association for Computing Machinery,* 9(3):297-314.

Holland, JH. (1992) *Adaptation in Natural and Artificial Systems.* MIT Press.

Holliday, R. (1987) "The Inheritance of Epigenetic Defects." *Science,* 238:163-170.

Hordijk, W. (1997) "A Measure of Landscapes." *Evolutionary Computation*, 4(4): 335-360.

Houck, CR, Joines, JA, Kay, MG, & Wilson, JR. (1997) "Empirical Investigation of the Benefits of Partial Lamarckianism." *Evolutionary Computation,* 5(1):31-60.

Hull, D. (1980) "Individuality and Selection." *Annual Review of Ecology and Systematics,* 11:311-332.

Jablonka, E, & Lamb, MJ. (1989) "The Inheritance of Acquired Epigenetic Variations." *Journal of Theoretical Biology,* 139:69-83.

Jablonka, E, Oborny, B, Molnar, I, Kisdi, E, Hofbauer, J, & Czaran, T. (1995) "The Adaptive Advantage of Phenotypic Memory in Changing Environments." *Philosophical Transactions of the Royal Society of London B,* 350:133-141.

Jennison, C, & Sheehan, N. (1995) "Theoretical and Empirical Properties of the Genetic Algorithm as a Numerical Optimizer." *Journal of Computational and Graphical Statistics,* 4(4):296-318.

Jones, T. (1994) A Model of Landscapes. Technical Report, SFI TR 95-02-21, Sante Fe Institute. February 1, 1994.

Jones, T, & Forrest, S. (1995) "Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms." In LJ Eshelman (ed.) Proceedings of *Sixth International Conference on Genetic Algorithms*, pp. 23-49. Morgan Kaufmann.

Julstrom, BA. (1995) "What Have You Done For Me Lately? Adapting Operator Probabilities in A Steady State Genetic Algorithm." In LJ Eshelman (ed.) Proceedings of *Sixth International Conference on Genetic Algorithms*, pp. 81-87. Morgan Kaufmann.

Kargupta, H, & Goldberg, DE. (1994) Decision Making in Genetic Algorithms: A Signal-to-Noise Perspective. IlliGAL Report, 94004, Dept of General Engineering, University of Illinois at Urbana-Champaign. August 1994.

Kauffman, SA. (1993) *The Origins of Order.* Oxford University Press.

Kiss, GR. (1992) "Variable Coupling of Agents to their Environment: Combining Situated and Symbolic Automata." In E Werner & Y Demazeau (ed.) *Decentralized AI 3*, Elsevier/North Holland.

Kondrashov, AS. (1995) "Modifiers of Mutation-Selection Balance: General Approach and the Evolution of Mutation Rates." *Genetical Research,* 66:53-69.

Koza, J. (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press.

Koza, JR. (1994) "Genetic Programming as a Means for Programming Computers by Natural Selection." *Statistics and Computing,* 4:87-112.

Lachmann, M, & Jablonka, E. (1996) "The Inheritance of Phenotypes: an Adaptation to Fluctuating Environments." *Journal of Theoretical Biology,* 181:1-9.

Lande, R, & Shannon, S. (1996) "The Role of Genetic Variation in Adaptation and Population Persistence in a Changing Environment." *Evolution,* 50(1):434-437.

Levins, R. (1968) *Evolution in changing environments : some theoretical explorations.* Princeton University Press.

Lewis, RW. (1981) "Evolution: A System of Theories." *General Systems,* XXVI:95-109.

Littman, M, & Ackley, DH. (1991) "Adaptation in Constant Utility Non-Stationary Environments." In RK Belew & L Booker (ed.) Proceedings of *Fourth International Conference on Genetic Algorithms (ICGA91)*, pp. 136-142. Morgan Kaufmann.

Lloyd, EA. (1994) *The Structure and Confirmation of Evolutionary Theory*. Princeton University Press.

Lund, HH. (1994) Adaptive Approaches Towards Better GA Performance in Dynamic Fitness Landscapes. Technical Report, PB-487, DAIMI, Aarhus University.

Lund, HH, & Parisi, D. (1994) Simulations with an Evolvable Fitness Formula. Technical Report, Institute of Psychology, CNR, Rome.

Lund, HH. (1995) "Genetic Algorithms with Dynamic Fitness Measures." In JT Alander (ed.) Proceedings of *First Norwegian Workshop on Genetic Algorithms (1NWGA)*. University of Vaasa.

Mahner, M, & Kary, M. (1997) "What Exactly Are Genomes, Genotypes and Phenotypes? And What About Phenomes?" *Journal of Theoretical Biology,* 186:55-63.

Maley, C. (1997) "Mutation Rates as Adaptations." *Journal of Theoretical Biology,* 186:339-348.

Manderick, B, de Weger, M, & Spiessens, P. (1991) "The Genetic Algorithm and the Structure of the Fitness Landscape." In RK Belew & LB Booker (ed.) Proceedings of *Fourth International Conference on Genetic Algorithms (ICGA91)*, pp. 143-150. Morgan Kaufmann.

Mani, GS. (1991) "Is There a General Theory of Biological Evolution?" In PP Saviotti & JS Metcalfe (ed.) *Evolutionary Theories of Economic and Technological Change: Present Status and Future Prospects*, pp. 31-57. Harwood Academic Publishers.

Mayley, G. (1996a) "The Evolutionary Cost of Learning." In P Maes, MJ Mataric, J-A Meyer & SW Wilson (ed.) *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, pp. 458-467. MIT Press/Bradford Books.

Mayley, G. (1996b) "Landscapes, Learning Costs and Genetic Assimilation." *Evolutionary Computation,* 4(3 - Special Issue (Evolution, Learning and Instinct: 100 Years of the Baldwin Effect)):213-234.

Mayley, G. (1996c) No Pain, No Gain: Landscapes, Learning Costs and Genetic Assimilation. CSRP 409, COGS. February, 1996.

Mayley, G. (1997) "Guiding or Hiding: Explorations into the Effects of Learning on the Rate of Evolution." In P Husbands & I Harvey (ed.) Proceedings of *Fourth European Conference on Artificial Life (ECAL97)*, pp. 135-144. MIT Press/Bradford Books.

Maynard Smith, J. (1978) "Optimization Theory in Evolution." *Annual Review of Ecology and Systematics,* 9:31-56.

Mayr, E. (1972) "Lamarck Revisited." *Journal of the History of Biology,* 5(1):55-94.

McFarland, D. (1977) "Decision Making in Animals." *Nature,* 269:15-21.

Menczer, F. (1994) Changing Latent Energy Environments: A Case for the Evolution of Plasticity. Technical Report, CS94-336, UCSD Computer Science and Engineering Dept.

Menczer, F, & Belew, RK. (1996) "Latent Energy Environments." In RK Belew & M Mitchell (ed.) *Adaptive Individuals in Evolving Populations: Models and Algorithms,* pp. Chapter 13. Addison Wesley.

Miller, BL, & Goldberg, DE. (1996) "Genetic Algorithms, Selection Schemes, and the Varying Effects of Noise." *Evolutionary Computation,* 4(2):113-131.

Miller, GF. (1995) Artificial Life as Theoretical Biology: How to do real science with computer simulation. CSRP 378, COGS, University of Sussex.

Mitchell, M, & Forrest, S. (1994) "Genetic Algorithms and Artificial Life." *Artificial Life,* 1(3):267-289.

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

Motta, E, & Zdrahal, Z. (1996) "Parametric Design Problem Solving." In(ed.) Proceedings of *10th Knowledge Acquisition for Knowledge Based Systems Workshop.*

Mühlenbein, H, & Schlierkamp-Voosen, D. (1995) "Analysis of Selection, Mutation and Recombination in Genetic Algorithms." In W Banzhaf & FH Eeckman (ed.) *Evolution as a Computational Process,* pp. 188-214. Springer Verlag. LNCS, vol. 899.

Mühlenbein, H, & Voigt, H-M. (1995) "Gene Pool Recombination in Genetic Algorithms." In IH Osman & JP Kelly (ed.) Proceedings of *Metaheuristics International Conference.* Kluwer Academic Publishers.

Mühlenbein, H. (1995) "Genetic Algorithms." In E Aarts & JK Lenstra (ed.) *Local Search in Combinatorial Optimization,* pp. Chapter 6. John Wiley & Sons.

Murakawa, M, S Yoshizawa, & T Higuchi. (1997) "Adaptive Equalization of Digital Communication Channels Using Evolvable Hardware." In T Higuchi, M Iwata & W Liu (ed.) Proceedings of *Evolvable Systems: From Biology to Hardware,* pp. 379-389. Springer Verlag.

Nelson, RN. (1995) "Recent Evolutionary Theorizing About Economic Change." *Journal of Economic Literature,* XXXIII:48-90.

Nix, AE, & Vose, MD. (1991) "Modeling Genetic Algorithms with Markov Chains." *Annals of Mathematics and Artificial Intelligence,* 5:79-88.

Noble, J. (1997) The Scientific Status of Artificial Life. Poster presented at Fourth European Conference on Artificial Life.

Nolfi, S, Elman, JL, & Parisi, D. (1994) "Learning and Evolution in Neural Networks." *Adaptive Behaviour,* 3:5-28.

Nolfi, S, & Parisi, D. (1995) Learning to Adapt to Changing Environments in Evolving Neural Networks. Technical Report, 95-15, Dept. of Neural Systems and A-Life, Institute of Psychology, CNR - Rome. November 1995.

Olby, RC. (1963) "Charles Darwin's Manuscript of "Pangenesis"." *British Journal for the History of Science,* 1(3):251-263.

Packard, NH. (1988) "Intrinsic Adaptation in a Simple Model for Evolution." In C Langton (ed.) *Artificial Life*, pp. 141-155. Addison Wesley. SFI Studies in the Sciences of Complexity, vol. VI.

Parisi, D, & Nolfi, S. (1996) "The Influence of Learning on Evolution." In RK Belew & M Mitchell (ed.) *Plastic Individuals in Evolving Populations*, Addison Wesley. SFI Series,

Parmee, I. (1996) "Towards an Optimal Engineering Design Process Using Appropriate Adaptive Search Strategies." *Journal of Engineering Design,* 7(4):341-362.

Pask, GA. (1965) "The Cybernetics of Evolutionary Processes and of Self Organizing Systems." In(ed.) Proceedings of *3rd International Conference on Cybernetics*, pp. 27-74.

Provine, WB. (1986) *Sewall Wright and Evolutionary Biology*. University of Chicago Press.

Radcliffe, NJ, & Surry, PD. (1994) "Formal Memetic Algorithms." In T Fogarty (ed.) Proceedings of *Evolutionary Computation: AISB Workshop*. Springer Verlag.

Ray, TS. (1995) "An Evolutionary Approach to Synthetic Biology, Zen and the Art of Creating Life." In C Langton (ed.) *Artificial Life, An Overview*, MIT Press.

Roughgarden, J. (1979) *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. Macmillan.

Rowe, J, & East, I. (1995) "Direct Replacement: A Genetic Algorithm without Mutation which Avoids Deception." In X Yao (ed.) *Progress in Evolutionary Computation. Selected Papers*, pp. 41-48. Springer Verlag. LNCS, vol. 956.

Rowe, J. (1997) *Personal Communication.*

Saviotti, PP, & Metcalfe, JS. (1991) "Present Development and Trends in Evolutionary Economics." In & JS Metcalfe PP Saviotti (ed.) *Evolutionary Theories of Economic and Technological Change: Present Status and Future Prospects*, pp. 1-30. Harwood Academic Publishers.

Scheiner, SM. (1993) "Genetics and Evolution of Phenotypic Plasticity." *Annual Review of Ecology and Systematics,* 24:35-68.

Schlicting, CD, & Pigliucci, M. (1995) "Gene Regulation, Quantitative Genetics and the Evolution of Reaction Norms." *Evolutionary Ecology,* 9:154-168.

Schwefel, H-P. (1981) *Numerical Optimisation of Computer Models.* Wiley.

Schwefel, H-P. (1994) "On the Evolution of Evolutionary Computation." In JM Zurada, RJ Marks II & CJ Robinson (ed.) *Computational Intelligence - Imitating Life*, pp. 116-124. IEEE Press.

Silverberg, G. (1988) "Modelling Economic Dynamics and Technical Change: Mathematical Approaches to Self-Organisation and Evolution." In G Dosi, C Freeman, R Nelson, G Silverberg & L Soeye (ed.) *Technical Change and Economic Theory*, pp. Chapter 24. Pinter.

Simpson, GG. (1953) "The Baldwin Effect." *Evolution,* 7:110-117.

Smith, JE, & Fogarty, TC. (1996) "Adaptively Parameterised Evolutionary Systems: Self-Adaptive Recombination and Mutation in a Genetic Algorithm." In HM Voigt, C Ebeling, I Rechenberg & HP Schwefel (ed.) Proceedings of *Parallel Problem Solving From Nature IV*, pp. 441-450. Springer Verlag.

Smith, JE, & Fogarty, TC. (1997) "Operator and Parameter Adaptation in Genetic Algorithms." *Soft Computing,* 2

Sniegowski, PD, & Lenski, RE. (1995) "Mutation and Adaptation: The Directed Mutation Controversy in Evolutionary Perspective." In (ed.) *Annual Review of Ecology and Systematics*, pp. 553-578. vol. 26.

Sober, E. (1984) *The Nature of Selection.* Bradford Books.

Spears, WM. (1995) "Adapting Crossover in Evolutionary Algorithms." In JR McDonnell, RG Reynolds & DB Fogel (ed.) *Evolutionary Programming IV: Proceedings of the Fourth Annual Conference on Evolutionary Programming*, pp. 367-384. MIT Press/Bradford Books.

Srinivas, M, & Patnaik, LM. (1994) "Adaptive Probabilities of Crossover and Mutation in Genetic Algorithms." *IEEE Transactions on Systems, Man and Cybernetics,* 24(4):656-667.

Stadler, PF. (1995) Landscapes and Their Correlation Functions. Working Report, 95-07-67, Santa Fe Institute.

Stadler, PF, & Happel, R. (1995) Random Field Models for Fitness Landscapes. Working Paper, 95-07-069, Sante Fe Institute.

Stadler, PF, & Wagner, GP. (1996) The Algebraic Theory of Recombination Spaces. Working Paper, 96-07-046, Sante Fe Institute.

Stark, H & Tuteur, FB. (1979) *Modern Electrical Communications: Theory and Systems.* Prentice Hall.

Stephens, CR, & Waelbroeck, H. (1996) Analysis of the Effective Degrees of Freedom in Genetic Algorithms. Preprint, ICN-UNAM-96-08, Institute for Nuclear Sciences, National Autonomous University of Mexico.

Stephens, CR, Garcia Olmedo, I, Mora Vargas, J, & Waelbroeck, H. (1997) Self-Adaptation in Evolving Systems. ICN-UNAM-97-10, 1st July, 1997.

Sterelny, K, Smith, KC, & Dickinson, M. (1996) "The Extended Replicator." *Biology and Philosophy,* 11:377-403.

Symonds, N. (1994) "Directed Mutagenesis - A Current Perspective." *Journal of Theoretical Biology,* 169(4):317-322.

Syswerda, G. (1989) "Uniform Crossover in Genetic Algorithms." In JD Schaffer (ed.) Proceedings of *Third International Conference on Genetic Algorithms and Their Applications (ICGA89)*, pp. 2-9. Morgan Kaufmann.

Syswerda, G. (1993) "Simulated Crossover in Genetic Algorithms." In LD Whitley (ed.) *Foundations of Genetic Algorithms 2 (FOGA2)*, pp. 239-255. Morgan Kaufmann.

Taylor, CE, Jefferson, DR, Turner, SR, & Goldman, SR. (1988) "RAM: Artificial Life for the Exploration of Complex Biological Systems." In C Langton (ed.) *Artificial Life*, pp. 275-295. Addison Wesley. SFI Studies in the Sciences of Complexity, vol. VI.

Thompson, A. (1995) "Evolving Fault Tolerant Systems." In(ed.) Proceedings of *1st IEE/IEEE International Conference on GAs in Engineering Systems (GALESIA95)*, pp. 524-529. IEE. IEE Conference Publications, vol. 414.

Thompson, A. (1996) "Evolutionary Techniques for Fault Tolerance." In Proceedings of *UKACC International Conference on Control (CONTROL'96).* IEE Conference Publication.

Thompson, A, Harvey, I, & Husbands, P. (1996) "Unconstrained Evolution and Hard Consequences." In E Sanchez & M Tomassini (ed.) *Towards Evolvable Hardware: The Evolutionary Engineering Approach*, pp. 136-165. Springer Verlag. LNCS, vol. 1062.

Turney, P, Whitley, D, & Anderson, R. (1996) "Evolution, Learning and Instinct: 100 Years of the Baldwin Effect." *Evolutionary Computation,* 4(3):iv-viii.

Turney, P. (1996) "How to Shift Bias: Lessons from the Baldwin Effect." *Evolutionary Computation,* 4(3):271-295.

Tuson, A, & Ross, P. (1996a) Co-evolution of Operator Settings in Genetic Algorithms. DAI Research Report, 789, Department of Artificial Intelligence, Edinburgh University.

Tuson, A, & Ross, P. (1996b) Cost Based Operator Adaptation: An Investigation. DAI Research Report, 787, Department of Artificial Intelligence, Edinburgh University.

van Nimwegen, E, Crutchfield, JP, & Mitchell, M. (1996) Finite Populations Induce Metastability in Evolutionary Search. 96-08-54, Sante Fe Institute. August 5, 1996.

van Valkenberg, ME. (1982) *Analog Filter Design.* Harcourt Brace Jovanovich Publishers.

Vavak, F, Fogarty, T, & Jukes, K. (1996) "A Genetic Algorithm with Variable Range of Local Search for the Adaptive Control of Dynamic Systems." In HM Voigt, C Ebeling, I Rechenberg & H-P Schwefel (ed.) Proceedings of *Parallel Problem Solving From Nature IV*, pp. 376-385. Springer Verlag.

Via, S. (1993) "Adaptive Phenotypic Plasticity: Target or By-Product of Selection in a Variable Environment." *American Naturalist,* 142:352-365.

Via, S, Gomulkiewicz, R, De Jong, G, Scheiner, SM, Schlicting, CD, & Van Tienderen, PH. (1995) "Adaptive Phenotypic Plasticity: Consensus and Controversy." *Trends in Ecology and Evolution,* 10(5):212-217.

von Neumann, J, ed. (1961) *The General and Logical Theory of Automata.* Pergamon Press. John von Neumann - Collected Works, vol. V.

Vose, MD. (1993a) A Critical Examination of the Schema Theory. Computer Science Dept, University of Tennessee.

Vose, MD. (1993b) "Modeling Simple Genetic Algorithms." In LD Whitley (ed.) *Foundations of Genetic Algorithms 2 (FOGA2)*, pp. 63-73. Morgan Kaufmann.

Vrba, ES, & Eldredge, N. (1984) "Individuals, Hierarchies, and Processes: Towards a More Complete Evolutionary Theory." *Paleobiology,* 10(2):146-171.

Wagner, GP, & Altenberg, L. (1996) "Perspective - Complex Adaptations and the Evolution of Evolvability." *Evolution,* 50(3):967-976.

Wagner, GP, & Booth, G. (1997) "A Population Genetic Theory of Canalization." *Evolution,* 51(2):329-347.

Weber, M. (1996) "Evolutionary Plasticity in Prokaryotes: A Panglossian View." *Biology and Philosophy,* 11:67-88.

Weinberger, ED. (1990) "Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference." *Biological Cybernetics*, 63: 325-336.

Weinberger, ED. (1991) "Local Properties of Kauffman's N-k Model: A Tuneable Rugged Landscape." *Physical Review A*, 44(10): 6399-6413.

Weinberger, ED. (1996) NP Completeness of Kauffman's N-k model, a Tuneably Rugged Fitness Landscape. Working Paper, SFI 96-02-003, Sante Fe Institute.

West-Eberhard, MJ. (1989) "Phenotypic Plasticity and the Origins of Diversity." *Annual Review of Ecology and Systematics,* 20:249-78.

Whitley, D. (1994) "A Genetic Algorithm Tutorial." *Statistics and Computing,* 4:65-85.

Whitley, D, Gordon, VS, & Mathias, K. (1994) "Lamarckian Evolution, The Baldwin Effect and Function Optimization." In Y Davidor, H-P Schwefel & R Manner (ed.) Proceedings of *Parallel Problem Solving from Nature III*, pp. 6-15. Springer Verlag. 866.

Whitley, LD. (1995) "Modeling Hybrid Genetic Algorithms." In G Winter, J Periaux, M Galan & P Cuesta (ed.) *Genetic Algorithms in Engineering and Computer Science*, pp. 191-201. John Wiley.

Williams, MB. (1970) "Deducing the Consequences of Evolution: A Mathematical Model." *Journal of Theoretical Biology,* 29:343-385.

Woodcock, G, & Higgs, PG. (1996) "Population Evolution on a Multiplicative Single-Peak Fitness Landscape." *Journal of Theoretical Biology,* 179:61-73.

Wright, S. (1988) "Surfaces of Selective Value Revisited." *The American Naturalist,* 131(1):115-123.

Zirkle, C. (1945) "The Early History of the Idea of the Inheritance of Acquired Characteristics and of Pangenesis." *Transactions of the American Philosophical Society,* 35:91-150.